Quest® Enterprise Reporter 3.2.1

Configuration Manager User Guide

Contents

Product Overview

Key Features of Enterprise Reporter

Organizations worldwide are struggling to keep up with corporate policies, changing government regulations, and industry standards. Generating reports that prove compliance, and deciding what data to include is a time consuming and difficult process. In order to meet compliance requirements or initiate IT best practices, organizations must know exactly what is in the IT infrastructure at any moment in time, how it is configured, and who has access to it. Quest presents Enterprise Reporter as a solution to these problems.

Enterprise Reporter provides a unified solution for data discovery and report generation. Using the Enterprise Reporter Configuration Manager, administrators can easily configure and deploy discoveries to collect and store data. Once the data has been collected, the Report Manager allows users to produce reports that help organizations to ensure that they comply with industry regulations and standards, adhere to internal security policies, monitor hardware and software requirements, and fulfill many other reporting requirements.

Using the Configuration Manager, you can:

- Configure your collection environment to minimize network traffic and optimize performance.

- Create discoveries to collect data that will be made available to the Report Manager:

    - information about your Active Directory® environment

    - information about files and folders from domains, OUs, computers, NetApp® and EMC® filers, shares, and DFS shares

    - information about the computers in your environment

    - data from specified SQL Server® computers, instances, and databases

    - general and registry information from selected computers

    - high-level summary information on file storage

    - high-level summary information and permissions in your Exchange® environment

    - information about your Azure® subscriptions, licenses, and service plans

    - information about your Azure Active Directory® environment

    - information about your Azure® resources

    - information about files and folders in your OneDrive® environment

    - information about your Office 365® Exchange Online environment

    - information about your Microsoft Teams in Office 365®

- Schedule discoveries to run automatically.

- Track the progress of discoveries, and pinpoint any errors in the collection.

Using the Report Manager, you can:

- Run reports on the data you have collected.

- Make predefined reports available to reporting users by publishing them.

- Create your own customized reports.

- Customize the appearance of your reports.

- Schedule reports to run when you need them.

- Publish reports to Knowledge Portal.

- Use the File Storage Analysis summary reports, with meaningful charts and graphs and the ability to drill down for more detailed information, to answer challenging administrative questions about file storage.

- Use the Exchange® summary reports, with meaningful charts and graphs and the ability to drill down for more detailed information, to answer challenging administrative questions about your Exchange® environments.

- Use the Exchange® reports to monitor and update the access permissions of accounts in an efficient and timely manner to ensure mailbox information security.

- Use the OneDrive® reports to answer questions about file and folder permissions in your OneDrive® environment.

- Use the Azure® reports to answer questions about your Azure® subscriptions, licenses, and settings.

- Use the Azure Active Directory® reports to answer questions about your Azure Active Directory® environment.

- Use the Azure® resource reports to answer questions about your Azure® resources.

- Us the Office 365 Exchange® Online reports to answer questions about your Exchange® Online mailbox, mailbox folders, and public folders and their permissions.

- Use the Microsoft Teams reports to answer questions about your Microsoft Teams in Office 365®.

Components of Enterprise Reporter 3.2.1

Enterprise Reporter has two components that work together to collect and report on your corporate IT data. Additionally, you can install the Database Wizard to help manage your Enterprise Reporter database and use the Knowledge Portal to make reports available online.

See also:

- The Configuration Manager

- The Report Manager

- The Database Wizard and Database Content Wizard

- The Encryption Key Manager

- Knowledge Portal

The Configuration Manager

In order to create reports, you must first gather the data. The Configuration Manager manages the process of creating and managing the discoveries that collect data, and the computers that perform the actual collecting. The Configuration Manager is intended for use by administrators who are responsible for managing data collection.

There are several tasks that you perform using the Configuration Manager:

- Configure clusters and nodes
- Create and run discoveries
- View errors and statistics for each discovery

See also:

- An Overview of the Configuration Manager
- Clusters
- Shared Data Locations
- Nodes
- Discoveries
- Role of the Server

An Overview of the Configuration Manager

Figure 1 outlines the process of using the console to configure and perform data collections.

Figure 1. Using the Configuration Manager

Setup

1 Create Cluster
- Logical collection of one or more nodes
- Minimalize bandwith by configuring one or more clusters per geographical location
- Scalable to represent needed processing power for performing discoveries

2 Create Nodes Within Clusters
- Computers which do the work of discovering data
- Configure nodes to control the number of tasks they can process

Configure

3 Create Discoveries
- Defining the objects that will be discovered
- Assigned to cluster

Run

4 Run Discoveries
- Run as often as you need to maintain data freshness
- Create jobs to discover data
- Monitoring and troubleshooting the discovery jobs

Clusters

You must configure at least one cluster. A cluster is a logical collection of one or more computers (nodes) on which discoveries are executed. A discovery must be assigned to a cluster. A cluster can access an optional shared data location for discovery data. This reduces network traffic, and the processing load on the server.

> **i** | **TIP:** In order to reduce network traffic and avoid delays in communication, a cluster should serve a single geographic location.

Shared Data Locations

Each discovery cluster can have an optional shared data location that is used by each of the cluster's nodes. Data collected from discoveries is stored in this shared data location, and then added to the SQL database maintained by the Enterprise Reporter server for report generation. When new data is discovered, it is compared to the data currently held in the Enterprise Reporter database. The difference between the existing data in the database and the new data from the discovery is added.

A shared data location may only be used by nodes within its assigned cluster. Since the Enterprise Reporter server receives only the difference between previously collected data and new data, the shared data locations cannot be shared among clusters.

> **i** | **NOTE:** Not all discovery types utilize the shared data location. For example, Active Directory, Exchange, NTFS, Azure Active Directory, and Office 365 discoveries do not use the shared data location to optimize on the number of collection tasks generated to collect the data in a more efficient manner.
>
> Ensure the optional shared data location is secured as it contains temporary files of collected data. For more information, see Things to Consider Before Creating a Cluster on page 26.

Nodes

A node is a computer assigned to a cluster and is responsible for processing discoveries. A node may only be assigned to a single cluster. A discovery consists of one or more tasks, each of which collects information from a target. A node may collect data from more than one target and may process more than one discovery simultaneously. The maximum concurrent tasks of each node defaults to 0 to allow the node to determine how many tasks it will process based on the available CPU. This setting allows the node to process discoveries with optimal performance. For more information, see Clusters on page 11, Discoveries on page 12, and Improving the Performance of a Node on page 34.

Discoveries

Discoveries are created to collect data. A discovery contains a number of targets, and is assigned to a cluster. The Enterprise Reporter server distributes the work among the nodes in that cluster. For more information, see Defining the Data Collection (Discoveries) on page 50.

Role of the Server

The Enterprise Reporter server is the central component of the Enterprise Reporter application. It directs the collection of data, maintains the report data in a SQL database (the central data store), organizes the computers running discoveries (nodes) into logical collections called clusters, assigns discoveries to the nodes using load-balancing, and executes report schedules.

A Simple Example of Enterprise Reporter Configuration

A global corporation has decided to use Enterprise Reporter to keep track of the various SQL databases throughout their enterprise. They have offices located in New York City, London, and Tokyo, and each of these offices has a unique domain hosting several SQL databases.

Each office will host 3 discovery nodes on various computers in their respective domains. These nodes are collected into clusters. A total of 3 clusters are created; one cluster for each office. Additionally, each cluster has its own shared data location located within its domain. When a discovery is created, it is assigned to the cluster in the office.

In each office, when a discovery process is executed, the server assigns the discovery to the nodes within the assigned cluster. All of the required data is written to the local shared data location and then uploaded to the central data store controlled by the Enterprise Reporter server.

The Report Manager

Once you have collected data, you use the Report Manager to generate reports. The Report Manager is a robust, flexible console that lets you create, modify and run reports. The Report Manager is intended for use by users who need to produce reports.

Reports in the Report Manager can be read-only, so that the settings cannot be changed, or they can be modifiable. For a modifiable report, you have control over what data is in the report, and how it is organized and laid out. Report definitions can be exported from one console, and imported into another. You can report on data collected from all clusters in your deployment.

The Database Wizard and Database Content Wizard

The Database Wizard is a stand-alone utility you can use to create and manage your Enterprise Reporter database. For more information, see Managing Your Database Using the Database Wizard chapter in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

You can launch the Database Content Wizard from the Database Wizard.

The Enterprise Reporter Database Content Wizard allows a user to perform the following tasks on the information stored in Enterprise Reporter SQL Server® Databases.

- Configuration information (Clusters and Nodes, Discoveries, and Reports)
    - Transfer
    - Backup
    - Restore
- Collected data
    - Clean
    - Merge

The Encryption Key Manager

Enterprise Reporter makes use of FIPS 140-2 compliant encryption to secure user credentials and includes an encryption key management tool. The Enterprise Reporter Encryption Key Manager can be started from the Windows Start menu. This tool allows you to perform the following tasks related to the Enterprise Reporter encryption key.

- Generating an encryption key
- Importing an encryption key from a backup file
- Exporting an encryption key to a backup file
- Resetting Enterprise Reporter user credentials

Knowledge Portal

Knowledge Portal is a cross product online reporting platform. Once Knowledge Portal is deployed, you can configure Enterprise Reporter for publishing. Reports can then be published, allowing users to generate reports using a web browser instead of the Report Manager.

For more information, see the Quest Enterprise Reporter Report Manager User Guide in the Technical Documentation.

Enterprise Reporter Architecture

Figure 2 shows how the components of Enterprise Reporter are related.

Figure 2. Enterprise Reporter Architecture
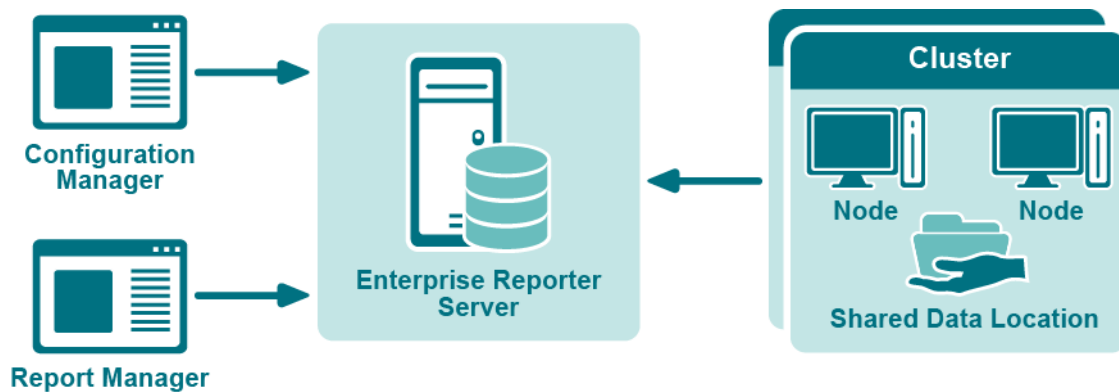
Summarizing the Workflow

The Enterprise Reporter workflow has three distinct phases:

Table 1. Phases of Enterprise Reporter Workflow

Phase	Console	Description	Frequency
Configuration	Configuration Manager	Set up clusters and nodes	Set up initially Modifications as your environment changes
Data collection	Configuration Manager	Create discoveries	Once for every set of targets
		Execute discoveries	As often as you need to maintain the desired data freshness
Reporting	Report Manager	Generate reports based on collected data	As needed

These phases are not linear. You may run useful reports for a period of time, and then decide that you need to add another node, create a new discovery, or add another managed computer. You can move around between them as needed.

Configuring the Configuration Manager

Starting the Configuration Manager

When you open the Configuration Manager, your first step is to connect to a server. Connecting to a server gives you access to its associated clusters, nodes, and discoveries. You need to know the name of the server, and the port number. The server name is the name of the computer where the server is installed. The port number was configured during the server installation. For more information, see Viewing Information About Your Enterprise Reporter Configuration and Installing Enterprise Reporter in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

> **i** | **NOTE:** If UAC is enabled, you must have elevated permissions to open the Configuration Manager.

> **i** | **NOTE:**
> To start the Configuration Manager, you must be a discovery administrator.
>
> For more information, see Installing and Configuring the Configuration Manager and Role Based Security in Enterprise Reporter in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

If this is your first time opening the Configuration Manager, you need to provide a license. For more information, see Licencing Enterprise Reporter in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

To connect to a server

1   Click the **Start** menu and select **Quest | Configuration Manager**.

2   Type the name of the server.

    - OR -

    Click **Browse**, and locate the computer where the server is installed.

    Once you have connected to a server, the server name is stored in the list for future use.

3   If necessary, type in the port number.

4   Click **Log In**.

Finding Answers and Getting Help

When using Enterprise Reporter, you can press F1 from any screen for help. The documentation included with Enterprise Reporter will open and display the section of documentation most relevant to the screen you are viewing.

The Help menu provided in Enterprise Reporter also addresses frequently asked questions related to the screen you are viewing. Clicking on a any of the topics provided will display the related section of documentation.

See also:

• Checking Online Support

• Help Finding and Creating Custom Reports

Checking Online Support

Enterprise Reporter also provides a tool for searching knowledge articles and community discussions.

To search Online Support

1   Click the link for **Online Support** in the header bar of the console.

2   Select whether to search **Knowledge Articles**, **Community Discussions**, or both.

3   Enter a keyword or phrase in the text box.

4   Press **Enter** or click the search icon.

5   Click the link of any result to open the result in the Online Support portal.

    i | **NOTE:** Signing in to the Online Support portal is required to view premium knowledge articles.

An Overview of Enterprise Reporter Communications and Credentials Required

There are many communication channels in Enterprise Reporter, involving different sets of credentials. This allows for controlled access to your environment, but you must understand where each set of credentials are used, and what permissions they need.

Figure 3 outlines where and for what each of the credentials are used, and the following tables explain the necessary permissions. For information on managing the credentials used in the Configuration Manager, see Using the Credential Manager on page 36.

Figure 3. Credentials used to communicate in the Configuration Manager
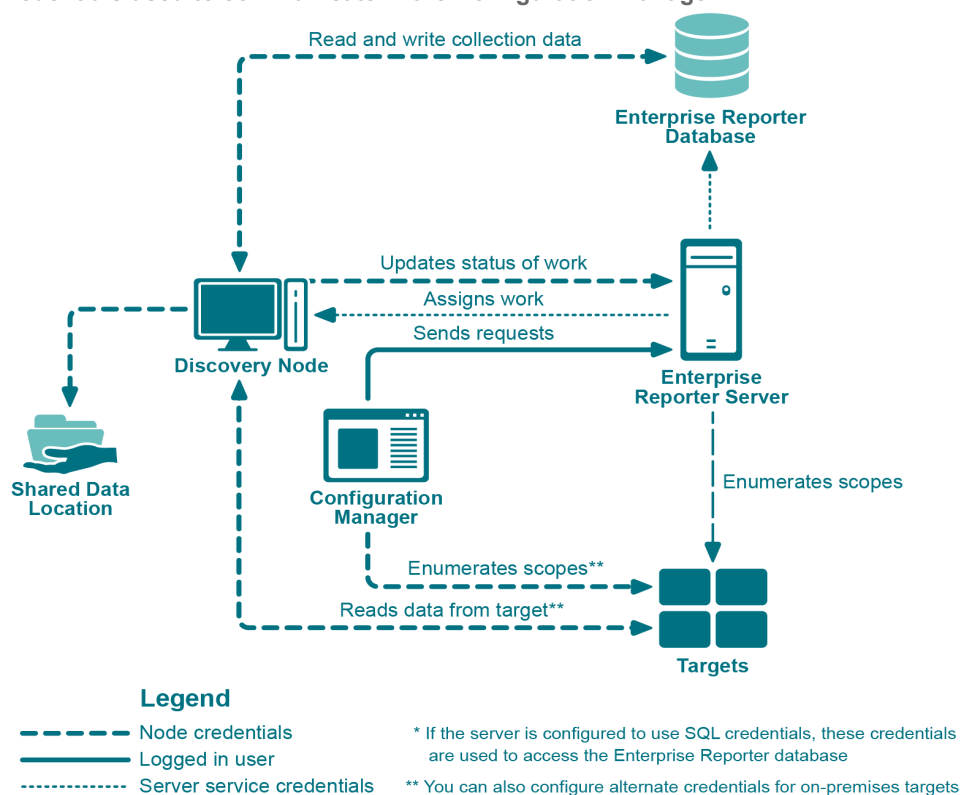
See also:

- Node Credential and Alternate Credential Details for On-Premises Discoveries

- Detailed Permissions for Enterprise Reporter Discoveries

- Permissions for Enterprise Reporter Discoveries on NAS Devices

- Permissions for Enterprise Reporter Tenant Applications

- Logged In User Details

- Server Service Credential Details

Node Credential and Alternate Credential Details for On-Premises Discoveries

Node credentials are provided when a discovery node is created, and you can modify them as needed. By default, the node's credentials are used to enumerate scopes and access on-premises targets.

If you want to use different credentials for a particular discovery, you can configure them in the Discovery Wizard. By using these alternate credentials, you can target anything on-premises for which you have credentials, in any domain. You can minimize the permissions given to node credentials, and use alternate credentials for scoping and collecting your on-premises discoveries.

The following table outlines the use of the node and alternate credentials, and how to properly configure your environment to ensure successful data collection:

Table 2. Node Credentials and Alternate Credentials in Configuration Manager

From	To	Permission Details	Configuration
Discovery Node	Enterprise Reporter Server	Provide server with job status, errors, statistics and logs.	Configured during node creation, or when you edit the node properties to change the credentials. The node credentials must have local administrator access to the host computer and be a member of the group "Reporter_Discovery_Nodes".
Discovery Node	Shared Data Location (if the cluster is configured to use one)	Read and write to the shared data location during data collection.	The shared data location is configured during the creation of a cluster. Ensure the node has read and write access to this file share. For more information, see Things to Consider Before Creating a Cluster on page 26.

Table 2. Node Credentials and Alternate Credentials in Configuration Manager

From	To	Permission Details	Configuration
Discovery Node	Enterprise Reporter Database	There are two options for communicating with the database: 1. You can use the same service credentials that the node service uses. 2. You can specify SQL credentials only for use when the database is accessed. The credentials you choose must be able to read and write to the database.	The account must be in the Reporter_Discovery_Nodes security group. (Note that if you use the same account as the Enterprise Reporter server it is already permissioned appropriately). For more information, see Role Based Security in Enterprise Reporter and Configuring the Database in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation. If you use SQL authentication to connect with the database, you must manually permission the SQL user, either by adding them to the database role Discovery_Nodes_Role or by permissioning specific tables in the database.
Discovery Node	Targets	Read access on all targets. For on-premises discoveries, all domains with which the credentials have a forest or domain level trust will be enumerated. If required, you can configure alternate credentials for specific discoveries, instead of using the default node credentials.	The targets are defined as part of a discovery. The discovery tasks are assigned to a particular node based on availability, so all nodes in a cluster should have access to all targets defined in all discoveries assigned to the node's cluster. For on-premises discoveries, ensure the node credentials or alternate credentials have read access to the target. In addition, a trust is required between the node computer and the targets. For more information on Azure and Office 365 Discoveries, see Detailed Permissions for Enterprise Reporter Discoveries on page 20.

Detailed Permissions for Enterprise Reporter Discoveries

The following table outlines the permissions required for Enterprise Reporter discoveries.

Table 3. Detailed Permissions required for Enterprise Reporter discoveries

Discovery Type	Permissions Required for Discovery Credential
Active Directory	An account with Active Directory read permissions is required to collect domain information, trusts, sites, domain controllers, and Active Directory computers, users, groups, and organizational units.
	The account being a member of the Built-in Domain Users group is sufficient to assign read permissions.
Azure Active Directory	An identity with read permission for the discovery target tenant. Read permissions are required for collection of tenant information, Azure Active Directory users, groups, group members, roles, and service principals.
	If additional credentials are being specified to minimize Azure throttling limitations, these credentials must have the same permissions as stated above.
	Also refer to credentials required to create and consent to the Enterprise Reporter Azure application required for this discovery. For more information, see Using the Tenant Application Manager on page 47.
Azure Resource	An identity with read permissions for the discovery target tenant. Read permissions are required for collection of subscription, Resource groups, and resources.
	If additional credentials are being specified to minimize Azure throttling limitations, these credentials must have the same permissions as stated above.
	Also refer to credentials required to create and consent to the Enterprise Reporter Azure Resource application required for this discovery. For more information, see Using the Tenant Application Manager on page 47.
Computer	An account with local administrator access on the scope computers to collect computer information, local groups and users, printers, services, policies, and event logs.
Exchange	To collect from Exchange targets, the credential account must have a mailbox on the target organization with access to read the permissions on the targets through EWS.
	To collect from Exchange 2007 targets, the credentials must be a member of the Exchange Organization Administrators Group.
	To collect from Exchange 2010, Exchange 2013, 2016, or Mixed Modes, the credentials must be a member of the Organization Management Group.
	To collect from Exchange 2016 or Exchange 2019, the credentials must have an administrator role with an assigned "ApplicationImpresonation" role.
Exchange Online	An account with access to the discovery target tenant.
	Read permission is required for collection of all Exchange Online information including mailboxes, mailbox delegates, public folders, mail-enabled users, mail contacts, distribution groups, group members, and permissions.
	If additional credentials are being specified to minimize Azure throttling limitations, these credentials must have the same permissions as stated above.

Table 3. Detailed Permissions required for Enterprise Reporter discoveries

Discovery Type	Permissions Required for Discovery Credential
File Storage Analysis	An account with local administrator access on the scoped computer is required to collect file, folder, share, and home drive analysis data.
	For permissions required when collecting NAS devices, see Permissions for Enterprise Reporter Discoveries on NAS Devices on page 22.
Microsoft SQL	An account with local administrator access on the SQL Server is required.
	Additionally, the account must have read access to the scoped database to collect database information.
Microsoft Teams	An identity with read permissions for the discovery target tenant. Read permissions are required for collection of Microsoft Teams information including teams, members, channels, applications, and drives.
	If additional credentials are being specified to minimize Azure throttling limitations, these credentials must have the same permissions as stated above.
	Also refer to credentials required to create and consent to the Enterprise Reporter Microsoft Teams application required for this discovery. For more information, see Using the Tenant Application Manager on page 47.
NTFS	If collecting through the administrator share, an account with local administrator access to the scoped computer is required.
	If collecting through a network share, an account with read permissions to the scoped shares is required.
	For permissions required when collecting NAS devices, see Permissions for Enterprise Reporter Discoveries on NAS Devices on page 22.
OneDrive	An account with access to the discovery target tenant. Administrator permissions are required for collection of all drives including drive information, configuration settings, files, folders, and permissions. A SharePoint administrator role is recommended.
	Additionally, the discovery credentials must have site collection administrator rights to each drive that is being collected.
	If additional credentials are being specified to minimize Azure throttling limitations, these credentials must have the same permissions as stated above.
	Also refer to credentials required to create and consent to the Enterprise Reporter Azure application required for this discovery. For more information see Using the Tenant Application Manager on page 47.
Registry	An account with local administrator access to the scoped computer is required to collect registry information.

Permissions for Enterprise Reporter Discoveries on NAS Devices

The following table outlines the permissions required for Enterprise Reporter discoveries.

Table 4. Permissions required for Enterprise Reporter discoveries on NAS Devices

Discovery Type	Permissions Required for Discovery Credential
NetApp Cluster Mode	Multiple virtual machines belong to a single cluster. All of these virtual machines can be specified as discovery targets. These virtual machines must be part of a domain.
	The NAS configuration must point to the cluster (name or IP address) with credentials that have read access to the cluster. These would typically be administrator credentials.
NetApp 7 Mode	In NetApp 7 mode, data can be collected on the storage controller or vFilers that are derived from the storage controller. Credentials with read access to the controller and vFiler are required.
NetApp Storage Controller	In NetApp 7 mode, data can be collected on the storage controller or vFilers that are derived from the storage controller. Credentials with read access to the controller and vFiler are required.
NetApp Filer	The vFiler can be a discovery target. In this case, the NAS configuration must point to the storage controller from which the vFilers are derived and the credentials must have read access to the storage controller.
Dell Fluid FS	The discovery target can be any Fluid FS VM. The NAS configuration must be the machine name or IP where Dell Enterprise Manager is installed and credentials must have access to Dell Enterprise Manager.
EMC Isilon	The discovery target can be any Isilon virtual machine. The NAS configuration must be the machine or IP that hosts the OneFS administration site and the credentials must have read access to it. By default, the connection is established using https and, if the connection is not deemed to be secure, the discovery will fail.

Permissions for Enterprise Reporter Tenant Applications

Enterprise Reporter requires Azure applications for the collection of Azure and Office 365 objects and attributes. These applications must be registered in the Azure portal and consent must be granted for delegated permissions. To manage tenant applications used by Enterprise Reporter please refer to in the System | Configuration | Application Tenant Management section in the Enterprise Reporter Configuration Manager User Guide.

OneDrive Azure Application Permissions

For the OneDrive discovery, an application with the name Quest Enterprise Reporter One Drive Discovery will be created. To create this application in your tenant, you must specify an account with administrative access to create applications. The account must have the Global Administrator role to be able to create and consent to the application.

Once created, the application must also be delegated permissions and an administrator must consent to the application's permissions using the Microsoft consent wizard. For the Quest Enterprise Reporter One Drive Discovery application, the following delegated permissions are required:

- Microsoft Graph: Read user files

- Office 365 SharePoint Online: Read user files

- Windows Azure Active Directory: Access the directory as signed-in user
- Windows Azure Active Directory: Read directory data

Azure Active Directory Application Permissions

For the Azure Active Directory discovery, the Exchange Online discovery, and the collection of group members for the OneDrive discovery, an application with the name Quest Enterprise Reporter Azure Discovery will be created. To create this application in your tenant, you must specify an account with administrative access to create applications. The account must have the Global Administrator role to be able to create and consent to the application.

Once created, the application must also be delegated permissions and an administrator must consent to the application's permissions using the Microsoft consent wizard. For the Quest Enterprise Reporter Azure Discovery application, the following delegated permissions are required:

- Microsoft Graph: Access directory as the signed in user
- Microsoft Graph: Read all users' basic profiles
- Microsoft Graph: Read all users' full profiles
- Microsoft Graph: Read directory data
- Microsoft Graph: Read identity risky user information
- Microsoft Graph: Read your organization's security events
- Azure Active Directory Graph: Access the directory as signed-in user
- Azure Active Directory Graph: Read all groups

Azure Resource Application Permissions

For the Azure Resource discovery, an application with the name Quest Enterprise Reporter Azure Resource Discovery will be created. To create this application in your tenant, you must specify an account with administrative access to create applications. The account must have the Global Administrator role to be able to create and consent to the application.

Once created, the application must also be delegated permissions and an administrator must consent to the application's permissions using the Microsoft consent wizard. For the Quest Enterprise Reporter Azure Resource Discovery application, the following delegated permissions are required:

- Windows Azure Service Management API: Access Azure Service Management as organization users
- Windows Azure Active Directory: Access the directory as signed-in user
- Windows Graph: Read all users' basic profiles

Microsoft Teams Application Permissions

For the Microsoft Teams discovery, an application with the name Quest Enterprise Reporter Microsoft Teams Discovery will be created. To create this application in your tenant, you must specify an account with administrative access to create applications. The account must have the Global Administrator role to be able to create and consent to the application.

Once created, the application must also be delegated permissions and an administrator must consent to the application's permissions using the Microsoft consent wizard. For the Quest Enterprise Reporter Microsoft Teams Discovery application, the following delegated permissions are required:

- Microsoft Graph: Read all users' basic profiles
- Windows Azure Active Directory: Access the directory as signed-in user
- Windows Azure Active Directory: Read all groups

Logged In User Details

The following table outlines the use of the logged in user credentials, and how to properly configure your environment to ensure successful data collection:

Table 5. Logged In User Credentials in Configuration Manager

From	To	Permission Details	Configuration
Configuration Manager	Enterprise Reporter Server	Must be a member of the Reporter_Discovery_Admins group in order to log in to the console. Configuration Manager will send configuration and set up requests to the server.	Configuration is dependent on your deployment's security group setup. See the System | Information page to determine the type of security in place. For more information, see Configuring the Database and Security Groups in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.
Configuration Manager	Targets	Must be able to enumerate the targets during scope selection, unless alternate credentials are provided for the discovery. All domains with which the credentials have a forest or domain level trust will be enumerated.	On each target, grant the user read access.

See also:

- Server Service Credential Details

Server Service Credential Details

Server service credentials are provided during the installation of the server. The following table outlines the use of the service account credentials, and how to properly configure your environment to ensure successful data collection:

Table 6. Server Service Credentials in Configuration Manager

From	To	Permission Details	Configuration
Enterprise Reporter Server	Enterprise Reporter Database	Must be in the local administrators group for the service to start properly. Must be able to read and write to the database. If the server is configured to use SQL authentication, the SQL credentials will be used to access the database, not the service account.	Configured automatically during installation. If you change the service credentials for the Quest Enterprise Reporter Server service, you need to ensure that a SQL login exists for that account, and create one if none exists. The login must be added to the database roles. For more information, see Configuring the Database and Security Groups in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.
Enterprise Reporter Server	Discovery Node	Must be in the local administrators group for the server service to fully function. Must be in the group Reporter_Discovery_Nodes for the service to fully function. Must be able to write to the Admin$ share to deploy the node. Controls the actions of the node.	On the node host, grant the service account local administrator rights.

Setting Up Your First Collection Computers

Before you can collect and report on data, you must set up the computers that will perform the collections. The minimum deployment is a single cluster with a single node, with the node residing on the same computer as the Enterprise Reporter server.

See also:

- Configuring Clusters and Nodes for Effective Data Collection
- Things to Consider Before Creating a Cluster
- Creating Your First Cluster and Node

Configuring Clusters and Nodes for Effective Data Collection

A cluster is a logical grouping of the physical computers (nodes) that will be collecting the data. Each physical computer in a cluster is a node, and each node may belong to only one cluster. You will be assigning collection jobs to a cluster, and the collection tasks are then spread across the nodes. To help make collections more scalable, all of the computers in the cluster share a data store, where the results of a data collection are stored.

Clusters provide scalability and performance benefits—you can have as few or as many clusters as your network demands.

Figure 4 outlines a three cluster implementation of Enterprise Reporter. The server and database are located in New York, with clusters in three other cities. Each cluster contains 3 nodes.

Figure 4. A typical enterprise deployment of Enterprise Reporter

> **i** | **TIP:** To maximize performance, and minimize network traffic, clusters should be physically close to the computers hosting the data you are collecting.

> **i** | **NOTE:** In order to collect data, you need to create a cluster — even if you are only planning to use a single computer to perform the collections.

Things to Consider Before Creating a Cluster

Make sure you are clear on the following before creating the cluster:

Do you want to use a shared data location?

As data is collected, it is compared to previously collected data on either the SQL Server® or the shared data location, depending on how you configure your cluster. If you have a lightly loaded SQL Server® that is physically close to your nodes, you may find that performance is improved by choosing not to use a shared data location. On the other hand, if network traffic is high and your SQL Server® is under a heavy load or physically distant from the nodes in the cluster, a shared data location will produce faster results.

- If you are using one, where is the shared data location?

  Create the shared folder, and give read and write access to the credentials you will be using for the nodes. As long as it is accessible, the shared data location can be located on any computer in your environment. Be cautious and ensure that the location is secured while it is accessible as it contains temporary files of collected data. For maximum benefit, locate the data source physically close to the nodes in the cluster.

  > **i** | **NOTE:** Not all discovery types utilize the shared data location. For example, Active Directory, For example, Active Directory, Azure Active Directory, Exchange, and NTFS discoveries do not use the shared data location to optimize on the number of collection tasks generated to collect the data in a more efficient manner.

- What are the first nodes that you want to add?

  A cluster is not functional until you add a node and enable it. The node computer is the computer that will resolve the targets of the collection, and perform the actual collection. You can add as many nodes as you want during the initial creation of a cluster.

- How many tasks do you want to run at the same time on each node?

  Discoveries are executed as tasks on the nodes of the assigned cluster. To maximize the performance of a node, you can limit the number of tasks that can run concurrently. The default setting of zero to allow the node to determine how many concurrent tasks it will process is a good place to start; you can later

experiment with limiting this number to improve performance. For more information, see Improving the Performance of a Node on page 34.

- What credentials are you going to use?

  Discoveries are performed by a service that runs on the node computer. The credentials used by the node must have read and write access to the Enterprise Reporter database, and have the read permissions to access the targets, and collect the necessary data, and permission to write to the shared data location. You can optionally choose to use SQL credentials for the database connection. For more information, see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17.

  > **ⓘ** | **NOTE:** If you use credentials that are different than the Enterprise Reporter server credentials, they need to be granted access to the database.
  >
  > For more information, see Configuring the Database and Security Groups in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

  > **ⓘ** | **NOTE:** These credentials are also used to access the target computers during a discovery, unless alternate credentials are specified for a discovery. For more information, see Step 1. Create the Discovery on page 54.

- Access to Admin$ on the node host

  To ensure the success of node deployment, the node installation files are copied into Admin$ (\\computername\admin$). The service account must have read and write access to this share. Once the node has been successfully installed, you can remove this access if required.

Creating Your First Cluster and Node

The Create Cluster wizard walks you through this process. You can create a cluster without a node, and add the nodes later, but you will not be able to run a discovery without at least one node enabled.

To create your first cluster and node

1. On the Manage Discovery Clusters pane, click **Create Cluster**.
2. Enter a name for the cluster.

   A default name, First Cluster, is provided, but you should change this to something meaningful, such as the location of your cluster.
3. Optionally, provide a description.
4. To use a network share, browse to and select your shared data location, and click **OK**.

   - OR -

   Select **No network share specified**.

   For more information, see Things to Consider Before Creating a Cluster on page 26.
5. Optionally, modify the connection timeouts.

   When you first create a cluster, it is recommended that you leave the default settings. Change the timeout settings only if you are getting timeout error messages. For more information, see Troubleshooting Connection Timeouts on page 121.
6. Click **Next**.

   If you do not want to add any nodes at this time, skip to step 13.
7. Click **Add** to configure a computer to serve as a node for this cluster.
8. Browse to the computer where the node is to be created and click **OK**.

   If you do not change the default entry, the first node is created on the current computer.

9   Select a node credential from the Credential Manager.

    If the account you want is not on the list, click Add and enter the account, then select it from the list. For more information, see Using the Credential Manager on page 36.

    > **i** | **NOTE:** This account must be a member of the Reporter_Discovery_Nodes group to get access to the encryption key used to encrypt all passwords. This account must have local administrator access to the node computer and write access to the shared data location for the cluster. These credentials also require read access the target computers during a discovery, unless alternate credentials are specified for a discovery. For more information, see Step 1. Create the Discovery on page 54.

10 Optionally, select **Specify a separate SQL Server Authentication credential for the database**, then select (or add) the SQL Server account in the Credential Manager.

    > **i** | **NOTE:** This account must have read and write access to the Enterprise Reporter database and must be in the Discovery_Admins group.

11 Optionally, set tasks to be **System - managed** or specify a **Maximum number of tasks**. For more information, see Nodes on page 12 and Improving the Performance of a Node on page 34.

12 Click **OK**.

    You can add more than one node at a time to the same cluster. Repeat steps 7through 12until you have added all of your nodes.

    By default, nodes are enabled after they are created. If you prefer to manually enable the nodes, clear the Enable Nodes check box. At least one node must be enabled in order for your cluster to be functional.

13 Click **Finish**.

    If you chose to create a cluster without any nodes, click Yes.

    The new node appears on the Discovery Nodes tab. Your node will be enabled, unless you cleared the Enable the nodes check box. For a listing of possible node statuses, see What does the status of a node or cluster indicate? on page 35.

    > **i** | **NOTE:** When a node is deployed and enabled, the cluster is also enabled. If you deployed the node without enabling it, you have to manually enable the cluster. For more information, see Enabling a Cluster on page 30.

    > **i** | **NOTE:** If a node fails to deploy, you must delete the node and recreate it. For information, see Node Issues on page 124.

Modifying your Deployment

For some situations, a single cluster with one node may be adequate. Other deployments may range from two or more nodes in a cluster to many clusters with many nodes in each.

See also:

- When Do You Add a Cluster?
- Modifying a Cluster
- Deleting a Cluster
- Disabling a Cluster
- What To Do if a Cluster is Disabled
- Enabling a Cluster
- Managing Nodes

When Do You Add a Cluster?

Typically, clusters are geographically based. Set up a cluster for each geographical location. You could also set up clusters to match your security structure, and group nodes into clusters based on the credentials you want to use for collections. For details on adding a cluster, see To create your first cluster and node on page 27.

Modifying a Cluster

You can change the name of a cluster, its description, and its associated shared data location. You may also want to change the timeout settings for the cluster if you are getting timeout error messages as you work in the Configuration Manager. For more information, see Troubleshooting Connection Timeouts on page 121. Another troubleshooting tool available to you is to change the level of logging for the nodes in the cluster. For more information, see Changing the Node Logging Level on page 126.

> **i** | **NOTE:** Before you change the shared data location, make sure no jobs are running. See Viewing a Cluster's Queue on page 116 and Canceling a Task or Discovery on page 118 for more information.

To modify a cluster

1   In the Manage Discovery Clusters pane, select the cluster.

2   On the Cluster Details tab in the bottom pane, change the name, description, database timeout settings, node logging level, or shared data location.

3   Click **Apply**.

Deleting a Cluster

Before you can delete a cluster, all nodes in the cluster must first be removed. Nodes cannot be removed until all jobs have either finished processing or been canceled. For more information, see Removing a Node on page 31 and What does the status of a node or cluster indicate? on page 35.

> **i** | **NOTE:** A discovery is assigned to a cluster. If you delete a cluster, the discovery cannot run. Re-create the discovery and assign it to another cluster.

To delete a cluster

1   In the Manage Discovery Clusters pane, select the cluster.

    Since the cluster has no nodes, it is in the Disabled state. If it is not disabled, you still have nodes deployed and must remove them.

2   Click **Delete Cluster**.

3   In the confirmation dialog box, click **Yes**.

Disabling a Cluster

Disable a cluster whenever you want to take all of the nodes in that cluster offline, but you know you will be using the cluster again. For example, if you need to perform maintenance tasks, you can disable a cluster. No work will be assigned to the cluster, but you can quickly bring it back online by enabling it.

To disable a cluster

1   In the Manage Discovery Clusters pane, select the cluster.

2   Click **Disable Cluster**.

What To Do if a Cluster is Disabled

A cluster that is not online is indicated in red. A disabled cluster cannot accept any jobs, so you should either troubleshoot the problem and enable the cluster, or re-create the discovery and assign it to another cluster. Care should be used when re-assigning discoveries, as you may affect the network load by increasing the distance between the data and the nodes.

Enabling a Cluster

If you have disabled a cluster, you need to enable it before it can do any work. When you enable a cluster, it enables all nodes in the cluster. This makes the cluster available for collections.

To enable a cluster

1   In the Manage Discovery Clusters pane, select the cluster.

2   Click **Enable Cluster**.

Managing Nodes

You may want to change the credentials your node is using. Occasionally, nodes may need to be disabled or stopped in order to perform regular system maintenance, or as part of troubleshooting issues with your Enterprise Reporter deployment.

See also:

• Modifying Node Credentials

• Modifying the Location of a Node's Temporary Files

• Removing a Node

• Why Disable a Node?

• Stopping a Node

• Enabling a Node

• Starting a Node

Modifying Node Credentials

You can modify the node credentials using the Configuration Manager. To ensure that your change goes smoothly, the new credentials should be permissioned as outlined in Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17. If you modify credentials that are used by a discovery node, the node service must be restarted before the changes take effect. Enterprise Reporter will attempt to restart the node; however, if the restart fails, you may need to manually start the service on the computer hosting the node. If there are jobs currently running on the node, they will be canceled. To prevent this, either change the credentials during a down time, or cancel the discoveries yourself and restart them once the change takes effect.

Occasionally a credential change will fail, for example because the node host computer is not available on the network or has a firewall configured, or because you have provided invalid credentials. In this case, the node will indicate that it failed to start in the bottom pane of the Manage Discovery Clusters view. You can manually restart the service using the Services console on the remote computer. For more information, see Troubleshooting Credential Change Failures on page 121.

> **i**   **NOTE:** The node account must be a member of the group Reporter_Discovery_Nodes.
>
> **NOTE:** If you are performing a password update on credentials, see Changing Passwords Using the Credential Manager on page 37 for more information.

i | **NOTE:** Nodes can potentially lock an Active Directory® password, so remember to change the password on the node whenever you update the password on the node's account.

To modify the credentials a node uses

1   Select **Manage Discovery Clusters**, and click the cluster containing the node.

    It is recommended that you cancel any running discoveries before changing your password.

2   In the bottom pane, select the **Discovery Nodes** tab.

3   Select the node, and click **Edit Node**.

4   To modify the service credential, select a user account from the Credential Manager.

5   To modify the credential that is used to connect to the Enterprise Reporter database:

    ▪   Click Database Credential

    ▪   Select the type of authentication

    ▪   If using SQL Server® Authentication, select a SQL account from the Credential Manager.

    If you are using Windows® authentication to connect to the database, you must use the same account as the service credential.

6   Click **OK**.

    A progress dialog box appears.

7   Verify that your changes were processed. If any errors occur, you will need to troubleshoot the issue and manually make any changes.

8   Click **Close**.

Modifying the Location of a Node's Temporary Files

Temporary data collection files are created while discoveries are running. These files are stored in the user's temp folder and can be large. If storage space is limited, these files may be stored in an alternate location by adding a DiscoveryTempFolder key to the appSettings portion of the ReporterNode.Exe.config file. This alternate location is optional and if it is absent or inaccessible, the default location is used.

i | **NOTE:** Ensure the user's temporary folder or alternate location is secured as it does contain temporary data collection files.

To modify the location of a node's temporary files

1   Open the ReporterNode.exe configuration file.

2   Add the following key to the appSettings portion of the file.
    Specify the path where the temporary files should be created.

    <appSettings>

          <add key="DiscoveryTempFolder" value="c:\reporter\somefolder" />

3   Save and close the configuration file.

Removing a Node

A node may need to be removed for a number of reasons, including:

•   You may be replacing the computer or hard drive hosting the node service.

•   Your node deployment failed. In this case, you must delete the node, and then recreate it. For more information, see Node Issues on page 124.

•   You want to delete a cluster.

- You may no longer need the node.

  > **i** | **NOTE:** You must disable a node before you can delete it.

  > **i** | **NOTE:** When you remove a node, it uninstalls the node from the host computer. If the removal fails for any reason, you can use the Control Panel on the host computer to uninstall the Quest Enterprise Reporter Node.

To delete a node:

1 On the Manage Discovery Clusters pane, select the cluster.

2 In the bottom pane, select **Discovery Nodes** tab and select the node.

3 Click **Disable Node**.

4 Click **Remove Node**.

5 In the confirmation dialog box, click **Yes**.

   The node's state changes to Undeploying until it is removed.

Why Disable a Node?

Disable a node whenever you want to take the node offline, but you know you will be using the node again. For example, if you need to perform maintenance tasks on the node computer, you can disable a node. No work will be assigned to the node, but you can quickly bring it back online by enabling it. You also need to disable a node before you can stop it (see Stopping a Node on page 32).

To disable a node

1 On the Manage Discovery Clusters pane, select the cluster.

2 In the bottom pane, select the **Discovery Nodes** tab and select the node.

   You can multi-select nodes to disable all nodes in the cluster.

3 Click **Disable Node**.

   If a node is actively processing a task, you will be prompted to either allow the task to finish processing, or cancel the task before disabling.

   Once the action is complete, the status of the node will change to Disabled. For a listing of possible node statuses, see What does the status of a node or cluster indicate? on page 35.

Stopping a Node

If you want to stop the service on the node host computer, you can use the Stop Node button. For example, this is useful when you need to change the password for your node credentials.

> **i** | **NOTE:** You must disable a node before stopping it. For more information, see Why Disable a Node? on page 32.

To stop a node

1 On the Manage Discovery Clusters pane, select the cluster.

2 In the bottom pane, select the **Discovery Nodes** tab and select the disabled node.

3 Click **Stop Node**.

   Occasionally the server will be unable to stop a node. In this case, you can use the Services console on the node host computer to stop the node.

Enabling a Node

Occasionally, a node may go offline and need to be enabled. Or, you may have chosen to create the node without enabling it. In order to be used for discoveries, a node must be enabled.

To enable a node

1   On the Manage Discovery Clusters pane, select the cluster.

2   In the bottom pane, select the **Discovery Nodes** tab and select the node.

   You can multi-select nodes to enable all nodes in the cluster.

3   Click **Enable Node**.

Starting a Node

When you start a node, it is immediately enabled, and available for processing discoveries.

To start a node

1   On the Manage Discovery Clusters pane, select the cluster.

2   In the bottom pane, select the **Discovery Nodes** tab and select the stopped node.

3   Click **Start Node**.

Improving the Performance of Your Discoveries

When a discovery is run on a cluster, the Enterprise Reporter server assigns work to its nodes. You can add nodes to a cluster at any time. Each node can only belong to one cluster. You can increase the performance of your discoveries by ensuring that nodes are configured to optimize the maximum number of concurrent tasks or by adding new nodes.

Each collection task is assigned to a node, balancing the distribution across the nodes until all the nodes are processing as many tasks as they are able. If no nodes are available to process the task, the task must wait until a node becomes available. Node performance is based on a combination of memory, processor speed and network bandwidth. If your network, computer memory, or processor speed are less than you would like, lean toward adding nodes. If your node is under used, set the maximum concurrent tasks to 0 to optimize node performance instead of adding more nodes. If a node starts to slow down when node performance is already optimized (by setting the maximum concurrent tasks to 0), adding a node will increase performance. For more information, see Nodes on page 12 and Improving the Performance of a Node on page 34.

See also:

•   Adding a Node

•   Improving the Performance of a Node

Adding a Node

Since a node must belong to a cluster, you must first create a cluster. For more information on creating clusters, see To create your first cluster and node on page 27.

> **i** | **NOTE:** The credentials you provide are the credentials used to access the targets of your discovery. Ideally, this should be a service account that has elevated privileges, not a user account.
>
> If you use credentials that are different than the Enterprise Reporter server credentials, they need to be granted access to the database.

> **i** | **NOTE:** The account must be a member of the group Reporter_Discovery_Nodes.

> **NOTE:** Performance counters must be enabled on the node machine as the performance information is required by the Enterprise Reporter node performance optimization functionality.

To create a node

1   On the Manage Discovery Clusters pane, select the cluster to which the node will be added.

2   In the bottom pane, on the Discovery Nodes tab, click **Add Node**.

3   Browse to the computer where the node is to be created and click **OK**.

4   Select a service account from the Credential Manager.

   If the account you want is not on the list, click Add and enter the credential, then select it from the list. For more information, see Using the Credential Manager on page 36.

   > **i** | **NOTE:** These credentials are also used to access the target computers during a discovery, unless alternate credentials are specified for a discovery. For more information, see Step 1. Create the Discovery on page 54.

5   To use SQL credentials to connect to the database, select **Database Credential**, then choose SQL Authentication and select the SQL account from the Credential Manager.

6   Set tasks to be **System - managed** to optimize performance. For more information, see Nodes on page 12 and Improving the Performance of a Node on page 34.

   > **i** | **NOTE:** CPU load is always used to determine how many concurrent tasks are assigned. Optionally, enter any number greater than 0 to set a **Maximum number of tasks** that the node will never exceed.

7   Click **Add**.

8   Repeat steps 3 through 7 to add additional nodes.

9   Click **OK**.

   The new node appears on the Discovery Nodes tab. By default, nodes on this pane are sorted by status. Your node will be enabled, unless you cleared the Enable nodes check box. For a listing of possible node statuses, see What does the status of a node or cluster indicate? on page 35.

   > **i** | **NOTE:** If a node fails to deploy, you must delete the node and recreate it. For more information, see Node Issues on page 124.

Improving the Performance of a Node

It is recommended to set tasks to be **System-managed** to optimize node performance. CPU load is always used to determine how many concurrent tasks are assigned. You may optionally enter any number greater than 0 to set a **Maximum number of tasks** that the node can process concurrently.

Modifying the maximum number of tasks on a node

1   On the Manage Discovery Clusters pane, select the cluster to which the node to modify belongs.

2   On the Discovery Nodes tab in the bottom pane, select the desired node.

3   Click **Edit Node**.

4   Set tasks to be **System - managed** to optimize performance.

- OR -

Set the **Maximum number** of tasks to any number greater than 0 to specify the number of tasks the node can process concurrently and click **OK**.

A progress dialog box is displayed.

5   Once the change is successfully made, click **Close** in the progress dialog box.

What does the status of a node or cluster indicate?

As you deploy, enable and disable nodes and clusters, the Configuration Manager gives you feedback. This feedback is visible in the Status column of the Manage Discovery Clusters pane. By default, clusters and nodes are grouped and sorted by Status.

> **i** | **NOTE:** You can change the sort order and grouping of your nodes and clusters.

The following table outlines each status of a cluster:

Table 7. Statuses of a Cluster in Configuration Manager

Status	Meaning
Disabled	The cluster is disabled and no new jobs will be processed. Any jobs currently running when the node was disabled will continue to process until they either complete or are canceled by the user.
Enabled	The cluster has at least one enabled node and is available to process jobs.

The following table outlines each status of a node:

Table 8. Statuses of a Node in Configuration Manager

Status	Meaning
Deploying	The node is currently being installed on the node computer.
Deployment Failed	The node could not be successfully deployed. For information on troubleshooting, see Node Issues on page 124.
Enabled	The node is online and available to process jobs.
Disabled	The node is disabled and no new jobs will be processed. If you attempt to disable a node while it is actively processing a task, you will be prompted to either cancel it or wait to disable the node until the task has completed.
Failed to Start	The node is still stopped, as the server was unable to start it.
Failed to Stop	The node is still running, as the server was unable to stop it. It will not accept new tasks from the server.
Faulted	The nodes regularly communicate with the server to confirm their health. A faulted node has not had contact within an acceptable time frame.
Incompatible Version	The node is not the same version of the software as the Enterprise Reporter server to which it is connecting.
Initializing	The node service is currently being configured, and required components are being downloaded to the node.
Removal Failed	The node service could not be deployed (or undeployed) successfully from the node computer.

Table 8. Statuses of a Node in Configuration Manager

Status	Meaning
Starting	The node service is in the process of starting up.
Stopped	The node service has been stopped. It is unavailable to process jobs until it is restarted.
Stopping	The node service is in the process of stopping.
Undeploying	The node service is currently being uninstalled from the node computer.
Upgrading	The node is currently being upgraded on the node computer.

Using the Credential Manager

Credentials are used in different places in Enterprise Reporter. For example, nodes and report schedules both use credentials. The Credential Manager is a central store for accounts and passwords used throughout the system. This makes it easy to keep passwords up to date, and allows you to enter the credential details once, and access them repeatedly.

> **i** | **NOTE:** Credentials added in the Report Manager are only available to the user who added them, while accounts added in the Configuration Manager are available to all Configuration Manager users on the same Enterprise Reporter server.

Accounts are not verified when you add them, and they must already exist in order to be used by Enterprise Reporter. For each account, you can add a description. This is particularly useful for differentiating between similar accounts, such as similarly named service credentials, or your SQL Server® default "sa" accounts. The combination of the account name and the description must be unique.

Credential passwords are encrypted with FIPS 140-2 compliant algorithms and an encryption key that has a user-supplied password. The encryption key is secured in the Windows Credential Manager (not to be confused with the Enterprise Reporter Credential Manager). The encryption key is used to encrypt all passwords that are used in Enterprise Reporter for collections and reporting. Additionally, Enterprise Reporter security groups and roles are used to validate who has access to the encryption key and the key cannot be retrieved without proper authorizations.

The Credential Manager will display a red key icon next to each account that requires a password.

See also:

- Changing Passwords Using the Credential Manager
- Changing Account Names Using the Credential Manager

To open the Credential Manager

- On the System | Configuration page, go to Credential Manager and click **Manage credentials**.

  - OR -

  Click the ellipsis anywhere credentials are required.

To use a credential from the Credential Manager

1   Select the account from the list.

2   Click **OK**.

To add a credential for use in the system

1   Open the Credential Manager.

2   Click **Add.**

3   Type the account name.

You can enter any account that you want to use in Enterprise Reporter, including Windows® accounts, SQL Server® accounts, and Azure accounts.

If you are entering credentials in the Configuration Manager, remember that other users may have entered the same credential, so if necessary, verify that you are adding a unique account.

4   Type the password for the account.

5   Optionally, type a description.

   The combination of the account name and the description must be unique.

6   Click **OK**.

To edit a credential

1   Open the Credential Manager.

2   Select an account.

3   Click **Edit**.

4   Optionally, modify the password by selecting **Edit Password**.

5   Make all changes.

6   Click **OK**.

7   If changing the password on an Office 365 Azure account, any Enterprise Reporter tenant applications using that account must also be reconfigured. For more information, see To reconfigure application for a tenant on page 48.

To delete a credential

1   Open the Credential Manager.

2   Select an account from the list.

   You can only delete accounts that are not currently in use.

3   Click **Delete**.

4   Click **OK**.

Changing Passwords Using the Credential Manager

When passwords are changed in Active Directory® and Azure Active Directory, they need to be updated everywhere they are in use in Enterprise Reporter. It is possible that the account could be locked if you do not make this change. The Credential Manager makes this easy by providing a central store for accounts. You can change the password, and it is updated in all nodes, schedules, discoveries, and so on.

> **i** | **NOTE:** You can modify credentials from anywhere you can access the Credential Manager. Be aware when you make a change to a credential, it is applied throughout your deployment, not just in your current context.

If you modify credentials that are used by a discovery node, the node service must be restarted before the changes take effect. Enterprise Reporter will attempt to restart the node; however, if the restart fails, you may need to manually start the service on the computer hosting the node. If there are jobs currently running on the node, they will be canceled. To prevent this, either change the credentials during a down time, or cancel the discoveries yourself and restart them once the change takes effect.

Credential passwords are encrypted with FIPS 140-2 compliant algorithms and an encryption key that has a user-supplied password. For further information on creation and management of the Enterprise Reporter encryption key, see Appendix: Encryption Key Manager on page 154.

The Credential Manager will display a red key icon next to each account that requires a password.

To change the password on an account used by Enterprise Reporter

1   On the System | Configuration page, click **Manage credentials used by Enterprise Reporter**.

    It is recommended that you make these changes while no discoveries are running (or waiting to be run) before changing the password.

2   Select the account, and click **Edit**.

3   Select **Edit Password**.

4   Modify the password and click **OK**.

    A progress dialog box appears.

5   Verify that your changes were processed. If any errors occur, you will need to troubleshoot the issue and manually make any changes.

6   If changing the password on an Office 365 Azure account, any Enterprise Reporter tenant applications using that account must also be reconfigured. For more information, see To reconfigure application for a tenant on page 48.

Changing Account Names Using the Credential Manager

In general, if you want to change an account name, it is recommended that you create a new credential, and delete the old one. However, in the case where you want to replace the credentials in use in a number of places in Enterprise Reporter, the Credential Manager enables you to make a single change and have it be applied across your deployment. For example, if you are provided a new service credential to replace a credential used for a dozen nodes in your environment, you can change the account name on the credential.

If you modify credentials that are used by a discovery node, the node service must be restarted before the changes take effect. Enterprise Reporter will attempt to restart the node; however, if the restart fails, you may need to manually start the service on the computer hosting the node. If there are jobs currently running on the node, they will be canceled. To prevent this, either change the credentials during a down time, or cancel the discoveries yourself and restart them once the change takes effect.

To change the account used by Enterprise Reporter

1   On the System | Configuration page, click **Manage credentials used by Enterprise Reporter**.

    It is recommended that you make these changes while no discoveries are running or waiting to be run before changing the account

2   Select the account, and click **Edit**.

3   Modify the account and click **OK**.

    A progress dialog box appears.

4   Verify that your changes were processed. If any errors occur, you will need to troubleshoot the issue and manually make any changes.

5   If changing the password on an Office 365 Azure account, any Enterprise Reporter tenant applications using that account must also be reconfigured. For more information, see To reconfigure application for a tenant on page 48.

Changing the Credentials used by the Enterprise Reporter Server

The credentials used to run the Quest Enterprise Reporter server service on the host computer are not stored in the Credential Manager. To change these credentials—either to a different account, or to update the password, take the following steps:

1   Stop any currently running discoveries. It is important to wait until all discoveries have finished canceling to ensure the integrity of the data.

    For more information, see Canceling a Task or Discovery on page 118.

2   Disable all nodes in each cluster.

    For more information, see Why Disable a Node? on page 32.

    > **!** | **CAUTION: At this point, you may want to inform your reporting users that they will lose connection to the Enterprise Reporter server, and they may lose work if they have not saved their latest changes.**

3   On the computer hosting the server, use the Services console to stop the service, then modify the properties of the Quest Enterprise Reporter Server service to use the new credentials.
    See your Microsoft® documentation for details.

4   Restart the service.

5   Open the Enterprise Reporter Encryption Key Manager and import the encryption key from the backup file. You must provide the user-supplied password that was used to create the backup file.

    For more information, see Importing a Key File on page 155.

6   In the Configuration Manager, enable the nodes.

    For more information, see Enabling a Node on page 33.

7   Restart any discoveries you canceled.

    For more information, see Manually Running a Discovery on page 109.

Configuring Global Settings

There are several global settings on the Configuration page that you can manage for Enterprise Reporter.

See also:

*   System | Configuration
*   Tips for Customizing the Configuration Manager Views

System | Configuration

The following section outlines the system configuration options.

*   Change History

    You can enable or disable change history for each discovery type. For more information, see Best Practices for Creating Discoveries on page 51.

*   Configure Logging

    You can set the amount of logging information that is collected from the Enterprise Reporter Server.

Each user of Configuration Manager can also configure the amount of logging information that is collected from each Configuration Manager computer. We recommend the default setting of Debug to collect as much information as possible. The following logging levels are available:

- **Debug**

  This is the most verbose logging level (maximum volume setting).

- **Information**

  The Information level is typically used to output information that is useful to the running of ER.

- **Warning**

  Warning is often used for showing when exceptional behavior has occurred. These can usually be handled without issue.

- **Error**

  Error is used to log all unexpected errors.

- **Fatal**

  Fatal is reserved for special exceptions/conditions causing a failure of a component (console or server) with the ER product.

- Credential Manager

  - Manage credentials

    You can manage credentials for use throughout the system. For more information, see Using the Credential Manager on page 36.

- Database Settings

  - Database Timeouts

    You can increase the timeout for the Enterprise Reporter server. Connection timeouts control how long the server has to establish a connection to the database, while the command timeout controls the amount of time available for processing a command on the database.
    For more information, see Troubleshooting Connection Timeouts on page 121.

  - Defragment Indices

    This option is available for systems running SQL Server™ Enterprise edition and is enabled by default. When this option is enabled, database indices are defragmented in the background as a regular part of collections to enhance the performance of both collection and reporting tasks.

    - Frequency Threshold

      Database indices will become more fragmented from frequent collections or from collections with large amounts of data. If Defragment Indices is enabled, you may optionally set the Frequency threshold to control the minimum amount of time that must elapse between tasks to defragment the database indices.

      For example, if the frequency threshold is set to 3 hours, and a defragmentation task finishes at 9am, no other defragmentation task will be allowed to begin until 12pm at the earliest. For any threshold, only one defragmentation task will run at a time.

    Enterprise Reporter running on systems that are not equipped with SQL Server™ Enterprise edition requires database indices to be defragmented manually during a shutdown using the Database Wizard to Perform Database Maintenance.

- Email Notifications

  Enterprise Reporter can be configured to send email notifications that indicate when nodes and discoveries change state and may need administrator intervention. You can configure host, sender, and recipient information, and select whether to send notifications for node state changes, discovery state changes, or both. You can also select whether to include all discovery state changes or just the failures.For more information, see Configuring Email Notifications on page 48.

- Extend Reporter Attributes

You can collect more than the default attributes for Active Directory® and Active Roles objects. For more information, see Managing the Collection of Additional Attributes on page 42.

- IT Security Search

  You can manage the configuration information and the credentials to access IT Security Search Warehouse Rest API. This option can also be used to select the information to be sent to the IT Security Search 11.3 Repository after every discovery. For more information, see Configuring IT Security Search on page 45.

- NAS Host Configuration

  You can manage the configuration information and the credentials for all NAS host devices that contain targets to be collected. For more information, see Configuring a NAS Host Device on page 46, File Storage Analysis Discovery: Configure NAS Host Devices on page 74, or NTFS Discovery: Configure NAS Host Devices on page 84.

- Server Error Notification

  You can manage the suppression of the login notification that the Enterprise Reporter server restarted unexpectedly since the last login. For more information, see Configuring Server Error Notification on page 47.

- Tenant Application Management

  - Manage tenant applications used by Enterprise Reporter

    You can register and manage the configuration information and the credentials for all applications used by Enterprise Reporter in your Azure portal. For more information, see Using the Tenant Application Manager on page 47, Azure AD Discovery: Configuring Enterprise Reporter Applications for Azure on page 89, OneDrive Discovery: Configuring Enterprise Reporter Applications for Azure on page 100, Azure Resource Discovery: Configuring Enterprise Reporter Applications for Azure on page 92, or Microsoft Teams Discovery: Configuring Enterprise Reporter Applications for Azure on page 98.

Tips for Customizing the Configuration Manager Views

To help you view information in the Configuration Manager, you can sort columns and resize panels.

To sort a column

- Click the column header.

To change the size of a panel

- Click and drag a pane divider.

Global Discovery Settings

There are settings which affect all discoveries of a given type:

- Configuring Change History

- Managing the Collection of Additional Attributes

- Configuring IT Security Search

- Configuring a NAS Host Device

- Configuring Server Error Notification

-
-

Configuring Change History

For all discovery types except Azure® Active Directory, Azure® Resource, Exchange, Exchange Online, Microsoft Teams, and OneDrive®, change history allows you to report on changes over time to the objects you collect. For example, if you choose to collect the change history for the NTFS discovery type, and a new file is added to a previously collected folder, you can see this reflected in a change history report.

You configure change history at a global level for each discovery. All discoveries of that type will collect this data. When you create a discovery, the Name page indicates whether change history is enabled for the discovery type.

In addition to the discoveries you can create and run in Enterprise Reporter, there is additional information that is common to more than one type of discovery, such as user accounts, groups or group members. To collect change history information for this data, enable change history for the Common discovery type.

To enable or disable change history for a discovery type

1   Click **System | Configuration**.

2   Click **Configure global change history settings**.

A button shows the current status of the change history configuration for each discovery type.

3   Click the **Enabled** or **Disabled** button to toggle the setting.

4   Click **Close**.

Managing the Collection of Additional Attributes

Enterprise Reporter collects a pre-defined set of attributes for each object in a discovery. The attributes collected vary depending on the type of discovery, the object, and the version of Enterprise Reporter you are using. You can add and remove attributes collected by Active Directory® and computer discoveries. You can extend:

- Active Directory® discovery attributes for users, groups, computers, organizational units, and service account.

- Computer discoveries with attributes from WMI classes added.

- Active Roles attributes for users, groups, computers, contacts, and organizational units.
  The following data types are supported:
    - Boolean
    - Case Ignored String
    - Case Sensitive String
    - DN String
    - Integer
    - Large Integer
    - Numeric String
    - Print Case String
    - Time
    - Unicode

NOTE: Collecting Active Roles attributes for users, groups, computers, and organizational units can significantly increase the collection time. Some attributes that are known to increase collection time are:

- edsaDialinAccessPermissions
- edsaDialinApplyStaticRoutes
- edsaDialinApplyStaticIP
- edsaDialinCallbackNumber
- edsaDialinCallbackOptions
- edsaDialinCallerID
- edsaDialinStaticIP
- edsaDialinStaticRoutes
- edsaDialinVerifyCallerID

Attributes you extended in previous versions of Enterprise Reporter may become default attributes in newer versions. In this case, the extended attribute is preserved to ensure that your reports continue to work, but only the default attribute is available for new reports.

When you add or remove attributes for a discovery type, Enterprise Reporter has to process them. This can take some time, during which any running discoveries of the type you extended may fail. You should perform the extension only after ensuring that no discoveries of that type are running or scheduled to run. Additionally, any attributes that are no longer being collected should be removed from any reports in which they were included.

TIP: In order to minimize collection time, it is recommended that you only collect attributes that you know are required by your reporting users. You cannot remove attributes that are collected by default.

NOTE: Reporting users will need to restart their consoles in order to have the most up to date list of available attributes in their reports. Consider informing users that data is no longer being collected, so they can remove the associated fields from their reports.

To add or remove Active Directory® attributes

1 Ensure that no Active Directory® discoveries are running or scheduled to run while your changes are processed.

2 Click **System | Configuration**.

3 Click **Manage attributes collected by Enterprise Reporter**.

4 Click **Yes** in the warning dialog box.

5 In the Active Directory® section, click **Extend**.

6 If you want to use a different forest to enumerate the schema, click the **Browse** button and select an appropriate domain.

   If your logged in user does not have access to a domain in the forest whose schema you want to enumerate, right-click in the dialog box and choose Connect as user.

7 Click **Get Schema**.

   The Extend Enterprise Reporter Attributes dialog box is displayed. Default attributes have grey check marks. Extended attributes have blue check marks for easy identification.

8 Select the type of attributes to collect from the Type menu.

9 Select to add or deselect to remove attributes.

   To view a list of your currently selected attributes, select Only show selected attributes.

10 Click **Apply**.

11 Click **Close**.

To extend computer attributes (using WMI classes)

1   Ensure that no computer discoveries are running or scheduled to run while your changes are processed.

2   Click **System | Configuration.**

3   Click **Manage attributes collected by Enterprise Reporter**.

4   Click **Yes** in the warning dialog box.

5   In the Computer section, click **Extend**.

6   Click ✚ **Add**.

   If you want to use a different computer to enumerate the WMI classes, click the ellipsis and select the computer.

   The Extend Enterprise Reporter Attributes dialog box is displayed. Any classes that have been extended are shown.

7   Expand the WMI class treeview as necessary to locate the desired classes.

8   Select the classes and click the **Add** button.

   You can only add entire classes, not individual properties.

   To remove classes, select them from the list and click the Remove button.

9   Click **Add**.

   The Extend Enterprise Reporter Attributes dialog box reappears, with the newly added classes in bold letters.

   If you add more than the recommended number of classes, a warning appears. Ensure that you require the selected extended WMI classes, as they will increase your collection time.

10 Verify the list and then click **Apply**.

To add or remove Active Roles attributes

1   Ensure that no Active Directory® discoveries are running or scheduled to run while your changes are processed.

2   Click **System | Configuration**.

3   Click **Manage attributes collected by Enterprise Reporter**.

4   Click **Yes** in the warning dialog box.

5   In the Active Roles section, click **Extend**.

6   If you want to use a different forest to enumerate the schema, click the ellipsis, select an appropriate domain, and click **OK** to continue.

   If your logged in user does not have access to a domain in the forest whose schema you want to enumerate, right-click in the dialog box and choose **Connect as user**. Use the Credential Manager to select or create a user with the required credentials and click **OK**.

7   To select the Active Roles server, click the ellipsis, select a computer, and click **OK**.

8   To set the credentials to access the Active Roles server, click the ellipsis, use the Credential Manager to select or create a user with the required credentials, and click **OK**.

9   Click **Get Schema**.

   The Extend Enterprise Reporter Attributes dialog box is displayed. Default attributes have grey check marks. Extended attributes have blue check marks for easy identification.

10 Select the type of attributes to collect from the Type menu.

11 Select to add or deselect to remove attributes.

   To view a list of your currently selected attributes, select Only show selected attributes.

12 Click **Apply**.

13 Click **Close**.

To remove computer attributes (non-default)

1 Ensure that no computer discoveries are running or scheduled to run while your changes are processed.

2 Click **System | Configuration**.

3 Click **Manage attributes collected by Enterprise Reporter**.

4 Click **Yes** in the warning dialog box.

5 In the Computer section, click **Extend**.

The Extend Enterprise Reporter Attributes dialog box is displayed. Any classes that have been extended are shown.

6 Select the desired classes and click **Remove**.

7 Click **Apply**.

Configuring IT Security Search

You can globally configure Enterprise Reporter to send information collected about Active Directory and Computers, NTFS, Azure Active Directory, Azure Resources, and Microsoft Teams to the IT Security Search Repository after every discovery. To be able to push collected data to the repository, the database account must have db-owner rights to the Enterprise Reporter database. The first discovery transmission includes all of the information collected for IT Security Search (based on your configuration settings). Subsequent discovery transmissions include only information that has been updated, is new, or has been deleted. You can view the statuses of the 100 most recent data transfers sent to the IT Security Search Repository or send all of the information in the database if maintenance or support issues occur.

For Enterprise Reporter to access the IT Security Search Warehouse Rest API, you must provide the appropriate host name, port, and credentials. This information can be obtained from the IT Security Search administrator.

To configure IT Security Search

1 Click **System | Configuration**.

2 Click **Manage configuration of IT Security Search**.

3 Enter the host name.

4 Enter the host port number.

- OR -

Select **Use HTTPS** to set the port number to 443.

Optionally, enter a custom host port number.

5 Click the ellipsis (...) to use the Credential Manager to enter and select the credentials required to access the host. For more information, see Using the Credential Manager on page 36.

6 Optionally, click **Test Connection** to verify access to the host.

> **i** | **NOTE:** If Enterprise Reporter encounters errors connecting to IT Security Search, the Configuration Manager logs will contain those errors. For more information, see Exporting Logs from the Configuration Manager on page 129.

7 Optionally, select the information to send to the IT Security Search Repository after every discovery.

8 If the connection test has been completed successfully, click **OK** to save the configuration settings.

To view the status of IT Security Search data transfers and re-send data if required

1 Click **System | Configuration**.

2   Click **Manage configuration of IT Security Search**.

3   Click **View Status**.

4   Optionally, click **Send All** to send all of the information in the database to the IT Security Search repository, review the confirmation, and click **Yes** to continue.

> **ⓘ** | **NOTE:** Normally, only information that has changed since the last discovery is sent to the IT Security Search Repository. Setting this option to send all available information when the next discovery is complete. You can wait for a scheduled discovery to finish or run a discovery manually.

5   Click **OK**.

Configuring a NAS Host Device

Configuring NAS host devices is only required if a NAS host device target has been added to an NTFS or File Storage Analysis discovery and that discovery is collecting volume information. For more information, see Choosing Your File Storage Analysis Scopes on page 72 and Choosing Your NTFS Scopes on page 77.

If you have NAS host devices that contain targets to be collected, you must enable the collection option and add

the configuration and credentials for each host device. You can optionally omit NAS host device configuration and collect shares as volumes for NAS device targets in the discovery. Use this option when volumes cannot be retrieved or do not need to be collected.

To configure a NAS host device

1   On the System | Configuration page under NAS Host Configuration, click **Manage configuration of NAS host devices**.

- OR -

For the File Storage Analysis Discovery, follow the steps outlined in File Storage Analysis Discovery: Configure NAS Host Devices on page 74.

- OR -

For the NTFS Discovery, follow the steps outlined in NTFS Discovery: Configure NAS Host Devices on page 84.

2   Click **Add**.

- OR -

Select an existing device and click **Edit**.

3   Select the appropriate type of device.

4   Enter the Host IP address.

For FluidFS targets, enter the IP address of the computer where Enterprise Manager is installed.

5   Enter the Host port number.

6   Click the ellipsis (...) to use the Credential Manager to enter and select the credentials required to access this device. For more information, see Using the Credential Manager on page 36.

7   Click **OK** to close the Credential Manager.

8   Click **OK** to close the NAS Configuration details.

9   Click **Close** to close the NAS Host Device Manager.

To delete a NAS host device

1   On the System | Configuration page under NAS Host Configuration, click **Manage configuration of NAS host devices**.

- OR -

For the File Storage Analysis Discovery, follow the steps outlined in File Storage Analysis Discovery: Configure NAS Host Devices on page 74.

- OR -

For the NTFS Discovery, follow the steps outlined in NTFS Discovery: Configure NAS Host Devices on page 84.

2   Select an existing device and click **Delete**.

3   Click **Yes** to confirm the deletion.

4   Click **Close** to close the NAS Host Device Manager.

Configuring Server Error Notification

Enterprise Reporter can be configured to display or suppress a login notification indicating that the Enterprise Reporter server has unexpectedly restarted since the last login.

To configure server error notification

1   On the System | Configuration page, click **Manage server error notification**.

2   Select **Do not show server restarted error** to suppress the notification.

3   Click OK.

Using the Tenant Application Manager

Azure applications are used by Azure Active Directory, Azure Resource, Microsoft Teams, and OneDrive discoveries. The applications used by Enterprise Reporter must be registered in your Azure environment and consent must be provided for the application's permissions. The owner of each Azure application will be the person who first configures the application.

The Tenant Application Manager in Enterprise Reporter will track what tenants are configured and available for discoveries. Basic configuration is available for the Enterprise Reporter applications. Note that any outside changes to Azure applications must be managed separately and can adversely affect the results of your collections.

To open the Tenant Application Manager

- On the System | Configuration page, click **Manage tenant applications**.

  - OR -

  Click the ellipsis in any Office 365 discovery scope page.

To add a tenant

1   Click **Add**.

2   Enter the name of the tenant where Enterprise Reporter applications will be created.

3   Click **OK**.

To delete a tenant

1   Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

    The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

Configuring Email Notifications

Enterprise Reporter can be configured to send email notifications that indicate when nodes and discoveries change state and may need administrator intervention.

To configure email notifications

1   On the System | Configuration page, click **Manage email notifications**.

2   Enter the Host name or IP address.

3   Enter the Host port number.

4   Click the ellipsis (...) to use the Credential Manager to enter and select the credentials required to access the mail server. For more information, see Using the Credential Manager on page 36.

    If you need to remove the credentials, click **Clear**.

If you want to test the Host configuration, press **Click to test your connection**.

5   Enter the email address to be displayed as the sender of the notifications.

6   Enter the email addresses of the people to receive the notifications (separated by commas or semicolons).

7   If there are any email address errors, fix them.

Optionally, click **View email address format errors** to display the errors.

8   Optionally, select **Send notifications on node state changes**.

9   Optionally, select **Send notifications on discovery state changes**.

10 If **Send notifications on discovery state changes** is selected, select **Send all** to send email notifications for every discovery state change.

- OR -

If **Send notifications on discovery state changes** is selected, select **Send only failures** to send email notifications only for only discovery state changes that are failures.

Understanding Discoveries

- Defining the Data Collection (Discoveries)
- Best Practices for Creating Discoveries

Defining the Data Collection (Discoveries)

Once you have configured a cluster, you can begin setting up discoveries. Discoveries define the targets from which you will be collecting data. Enterprise Reporter uses a "collect all" model. After you run a discovery, you can run reports that include the data you have collected. For more information on reporting, see the Quest Enterprise Reporter Report Manager User Guide in the Technical Documentation.

Enterprise Reporter includes the following types of on-premises discoveries:

Table 9. Types of on-premises Discoveries included in Enterprise Reporter

Type	Description
Active Directory	Collects information about your domains and Active Directory® objects within the domains, such as users, groups, sites and trusts.
Computer	Collects information specific to a computer, such as printers, shares and security policies.
Exchange	Collects information about your Exchange organization and permission information about your organization's mailboxes, stores, public folders, contacts, groups, and group members.
File Storage Analysis	Collects information about your network's file storage capacity and usage.
Microsoft SQL	Collects information about your Microsoft® SQL Servers®.
NTFS	Collects information about your NTFS structure— files, folders, and permissions.
Registry	Collects registry keys and values from available registry hives.

> **i** | **NOTE:** The remote registry service needs to be enabled for the collection of some attributes.

Enterprise Reporter includes the following types of Office 355™ discoveries:

Table 10. Types of Office 365 Discoveries included in Enterprise Reporter

Type	Description
Azure Active Directory	Collects information about your Office 365™ tenant information about your Azure™ users and groups.
Azure Resource	Collects information about your Azure subscriptions.
Exchange Online	Collects information about your Office 365™ tenant and permission information about your Exchange Online™ mailboxes, public folders, contacts, groups (static, dynamic, and unified) and group members.

Table 10. Types of Office 365 Discoveries included in Enterprise Reporter

Type	Description
Microsoft Teams	Collects information about your Microsoft Teams.
OneDrive	Collects information about your Office 365™ tenant and permission information about your OneDrive® files and folders for licensed OneDrive users.

Best Practices for Creating Discoveries

To get the best performance, and meet the needs of your reporting users, there are a number of things you should consider.

For On-Premises and Office 365 Discoveries

- The default settings of a discovery are designed to optimize performance. If you are going to change them, ensure that your reporting users require the data you are going to collect.

- Schedule your discoveries to collect data only as frequently as required to satisfy the needs of your reporting users.

For On-Premises Discoveries

- Group together targets in a discovery based on the data you are collecting. A discovery should collect generally the same data from all targets. For example, if you have some targets from which you only want to collect the Windows folder, and other targets from which you want to collect both the Windows and Program Files folders, you should create two different discoveries. This makes it easier to design and maintain discoveries.

- Each discovery must be assigned to a cluster. Assign the discovery to the cluster geographically closest to the targets. If necessary, break a discovery into smaller discoveries to accomplish this. By assigning each discovery to a cluster, you can also control which nodes are used for each discovery. For example, you can create one cluster with two nodes for running your Active Directory discoveries and create a second cluster with 5 nodes for running NTFS discoveries.

- Understand how each discovery type works with tombstoning. For example, you cannot collect SQL information from a single SQL Server using multiple discoveries; however, you can collect NTFS information for a single computer using multiple discoveries. Refer to Table 3. On page 48.

- The following table contains sample collection timings for on-premises discoveries.

Table 11. Sample collection timings for on-premises discoveries

Discovery Type	Number of Objects	Number of Nodes	Collection Time
Active Directory	100,000 users	1	40 minutes
Active Directory	500,000 users, groups, and computers (includes 10 million group members)	2	11 hours
Computer	200 computers	12	25 minutes
File Storage Analysis	1 server with 1 million files and folders	1	2 hours 25 minutes
Exchange	1,000 mailboxes and 15,000 mailbox folders	1	40 minutes
Exchange	10,000 mailboxes with permissions	1	1 hour 20 minutes
NTFS	5 million files and folders with permissions	6	5 hours 50 minutes
NTFS	31 million files and folders with permissions	10	20 hours

Note: All collection times are affected by hardware specifications and network activity.

For Office 365 Discoveries

- Microsoft will throttle traffic when they determine the need for it. This means that when an Azure or Office 365 discovery is running, Microsoft will send throttle warnings and force the collection to slow down or stop and wait for a specific period of time. In testing, we have seen a throttling delay period of 2 minutes. We have implemented methods to handle throttle warnings and keep the collection going. The use of multiple credentials can help minimize throttling.

  For large collections, it is highly recommended to provide multiple properly-provisioned credentials per discovery. For Azure discoveries, the credentials are alternately used to query the data from the tenant, so using multiple credentials will reduce the number of calls per minute. For OneDrive discoveries, throttling is handled differently and the collection will only switch to an alternate credential if a throttling message is received.

  In general, for large Office 365 discoveries, testing has shown that providing 3 or more credentials works very well.If the Enterprise Reporter node log files indicate that the discovery credentials are consistently locked, try adding a credential to the discovery to assist collection.

- Selecting smaller discoveries improves performance.

  For OneDrive discoveries, select only the drives that contain information you need.

  For Azure Active Directory discoveries, only collect Office 365 object attributes if they are needed. For large Azure discoveries, we recommend creating multiple discoveries for the same tenant and spitting up discoveries by object type. It is recommended to collect users and groups in one discovery and remaining Azure objects in different discoveries. For Azure Active Directory discoveries, the reason to collect users and groups together is that users are group members and, therefore, will not have to be collected twice.

  For Azure Resource discoveries, select only the subscriptions that contain information you need. For large subscriptions, we recommend creating multiple discoveries for each subscription.

- Understand how each discovery type works with tombstoning. For example, you can collect Azure AD objects from a single tenant across multiple discoveries and data will not be tombstoned. For more information, see How Scopes Affect Tombstoning on page 56.

- Understanding how your discovery will resolve to collection tasks is very important.

  For Azure discoveries, as show in Table 12, the scoped tenant will resolve to a single collection task for the tenant. This means that if you want to utilize more nodes for the collection of a single tenant, you should create a discovery by object type for optimal performance. For Azure discoveries, there are 5 object types

(users, groups, contacts, roles, and service principals) that can be collected so a maximum of 5 nodes can be used at one time.

For OneDrive discoveries, as shown in Table 12, the scoped tenant will resolve to a single collection task for the tenant, but the discovery will be tombstoned based on the drives, not the object type like Azure AD discoveries. This means that if you want to utilize more nodes for the collection of a single tenant, you must create a discovery for each drive or group of drives. For example, if your tenant has 10 drives, you can create 10 discoveries so a maximum of 10 nodes can be used at the same time.

• The following table contains sample collection timings for Office 365 discoveries.

Table 12. Sample collection timings for Office 365 discoveries

Discovery Type	Number of Objects	Number of Nodes	Number of Discovery Credentials	Number of Threads (OneDrive only)	Collection Time
Azure AD	30,000	1	2	N/A	30 minutes
Azure AD	300,000	1	1	N/A	5.5 hours
Azure Resource	7 subscriptions with 7,000 resources	1	3	N/A	30 minutes
Exchange Online	1,000 mailboxes	1	1	N/A	8 minutes
Exchange Online	1,000 mailboxes with permissions and delegates	1	1	N/A	3 hours 59 minutes
Exchange Online	1,000 mailboxes with folders	1	1	N/A	1 hour 12 minutes
Exchange Online	1,000 mailboxes with folders	1	3	N/A	24 minutes
Exchange Online	5,000 mailboxes	1	1	N/A	15 minutes
Exchange Online	5,000 mailboxes with permissions	1	1	N/A	4 hours 35 minutes
Microsoft Teams	1,000 teams with 10,000 files and folders	1	1	N/A	1 hour 30 minutes
Microsoft Teams	1,000 teams with 200,000 files and folders	1	1	N/A	5 hours
OneDrive	Drives: 1 Files/Folders: 400,000	1	11	10	17 hours
OneDrive	Drives: 5 Files/Folders: 20,000	1	5	10	1
OneDrive	Drives: 100 Files/Folders: 500,000	1	11	10	21 hours

Note: All collection times are affected by hardware specifications and network activity.

Creating Discoveries

• Step 1. Create the Discovery

• Step 2. Choose what to include in your discovery (Scopes)

• Step 2a. Choose scopes for your discoveries

There are several steps for creating a discovery. A wizard guides you through the process, which varies slightly depending on the type of discovery.

Step 1. Create the Discovery

When you are creating a discovery, it is important to consider what cluster will be running the discovery. A discovery can only belong to one cluster. When you run the discovery, the collection is performed by the nodes in the cluster.

By default for on-premises collection, the credentials used to access the targets and read the data are those provided when creating the node. If required, you can specify alternate credentials during the creation of your on-premises discovery. For more information, see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17.

> **i** | **TIP:** It is recommended that you assign the discovery to the cluster closest to the data.

To create a discovery

1   On the Manage Discoveries pane, click **New Discovery.**

    - OR -

    Select an existing discovery on the Manage Discoveries pane and click **Duplicate** to create an exact copy. Click **OK** to confirm that you want to duplicate the selected discoveries and edit the copy to meet your needs. For more information, see Modifying a Discovery on page 117.

2   Select the type of discovery. On the Name page of the Create Discovery wizard, type a unique name for your discovery.

3   Provide an optional description that outlines the data you collect with this discovery.

4   Select the assigned cluster.

    Your change history status is indicated. For more information, see Best Practices for Creating Discoveries on page 51.

    > **i** | **NOTE:** A cluster with a red icon is currently disabled. Your discovery cannot be run until you resolve the issue with your cluster.

5   For all discoveries except Azure, Exchange Online, Microsoft Teams, and OneDrive, select Use default node credentials to target computers that the logged in user can access.

    For access to additional targets, use alternate credentials and choose an account from the Credential Manager. For more information, see Using the Credential Manager on page 36.

    - OR -

    For Azure, Exchange Online, Microsoft Teams, and OneDrive discoveries, click **Add** to use the Credential Manager and select (or Add) an Office 365 administrator account within the target tenant that does not have multi-factor authentication enabled and click **OK**. For more information, see Using the Credential Manager on page 36.

    If any warnings or messages are displayed, review them and click **OK** to continue.

    Optionally, click **Remove** to delete an invalid or unwanted account.

    Optionally, repeat these steps to add additional accounts that can be used to access the same target tenant for this discovery. Specifying additional accounts helps optimize the discovery by collecting tenant information in a more efficient manner.

6   For all discoveries except Active Directory, Azure, Exchange Online, Microsoft Teams, My SQL, and OneDrive, choose whether to ping computers in the discovery, and how long to allow for a response. (NTFS, computer, Exchange, File Storage Analysis, and Registry discoveries do not ping by default).

This sets the amount of time given to confirm a computer's existence prior to attempting to collect data. Leave this option disabled if it is not possible to ping the computers in the discovery.

7   Click **Next** to continue to the Scope page.

Step 2. Choose what to include in your discovery (Scopes)

In this section, we will explore what targets (scopes) you can define for each discovery.

See also:

- Scopes: An Overview
- How Scopes Affect Tombstoning
- Using the Browser to Include and Exclude Scopes
- Including Objects in Your Scope
- Importing Computers to Your Scopes
- Refining Your Scope with Exclusions
- Filtering in the Browser
- Using Queries to Define Your Scopes
- Using Subnets to Define Your Scopes
- Using Subnet Credentials

Scopes: An Overview

Scopes define the targets of the discovery. Scope options vary depending on the type of discovery you are creating. When you are choosing scopes for on-premises collection, the node credentials or alternate credentials you specify on the Name page of the Create Discovery wizard determine the available targets. If you are using default node credentials, only targets that the logged in user can access are shown. If you provided alternate credentials when you created the discovery, those credentials are used to enumerate your scopes. For more information, see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17.

i | **TIP:** It is recommended that you only include a target (computer) in a single discovery for each type.

i | **NOTE:** Use care when changing credentials. Credentials must have read access to all targets of the discovery, or tasks will fail. If this is not possible for all targets in one large discovery, break the discovery into smaller discoveries.

i | **NOTE:** When choosing scopes for Azure Active Directory, Azure Resource, Exchange Online, Microsoft Teams, and OneDrive, the credential provided on the Name page of the Create Discovery wizard provides the default tenant name. If the tenant name is different than the credentials provided, the tenant name can be edited and specified. The tenant name is then used to enumerate scopes.

Some discovery types have additional options to collect related information that adds value to your reports. For example, if an NTFS discovery encounters AD groups in the security settings on an object, you can collect and report on the nested members of the groups. The data is collected for all scopes in the discovery and will add time to your discovery, so take this into consideration when selecting this option.

You may be able to enable this in a subset of your discoveries. For example, if you have six different discoveries with varying schedules that could potentially collect the same group members, you could enable it in only the discovery that is scheduled once a week, assuming that is sufficient to meet your reporting needs. In this way, performance is maximized, and reports have the data they need. It does not matter what discovery type is used to collect the data, as long as you are sure the data will be complete. Results are available for any report that includes the field.

How Scopes Affect Tombstoning

The concept of tombstoning refers to the process of comparing the objects found in previous collections with those found in current collections and subsequently marking objects that are no longer found in the current collection as tombstoned in the database. Tombstoned objects in the database will not be shown in library reports.

Enterprise Reporter makes the best assumptions it can based on what is collected. For some discovery types such as Computer, SQL Server, and Registry, use caution when changing collection options between collection runs on the same target (same computer). If these options are changed, objects that are not found in resulting collections will be marked as tombstoned in the database.

For example:

- For a Computer collection, if you collect shares and services for a particular computer in one collection and then run another collection for the same computer but just select shares, all the services will be marked as tombstoned and only shares will be updated.

Other discovery types such as Active Directory, NTFS, Azure Active Directory, Azure Resource, and Microsoft Teams allow for collection of different object types from the same target without tombstoning the objects.

For example:

- For an Active Directory collection, if you collect users and groups for a particular domain in one collection and then run another collection for the domain computers for the same domain, the users and groups from the previous collection will not be tombstoned.

- For a OneDrive collection, if you collect a drive from a particular tenant in one collection and then run another collection for a different drive from the same tenant, the previous drive will not be tombstoned.

The following table describes how each discovery type works with tombstoning.

Table 13. How each discovery type works with tombstoning

Discovery Type	Root Object	Tombstoning?*
Active Directory	Domain	No
Azure Active Directory	Tenant	No
Azure Resource	Tenant	No
Computer	Computer	Yes
Exchange	Organization	No
Exchange Online	Tenant	No
File Storage Analysis	Computer	Not applicable - collects historical snapshots
Microsoft SQL	Microsoft SQL Server	Yes
Microsoft Teams	Tenant	No
NTFS	Share	Yes
OneDrive	Drive	No
Registry	Computer	Yes

*Yes indicates that objects will be tombstoned when subsequent collections run on the same target with different collection options.

Note: All discovery types have a root object that is not tombstoned by Enterprise Reporter. For example, in Computer, File Storage Analysis, MS SQl, NTFS, and Registry discoveries, the computer is the root object, so computers will never be tombstoned.

Using the Browser to Include and Exclude Scopes

The browser is designed to allow you to drill into the acceptable objects for a given discovery type. Although the browser may vary slightly between discoveries, the basic use of it is consistent.

Your discovery should contain objects for which you want to collect similar data because:

- there are several options that are applied to every object in a discovery, such as global scopes and discovery options.

- it makes it easier to understand what you are collecting.

- you are more likely to meet the needs of your reporting users by providing consistent data.

    i | **NOTE:** The account you are logged in as is used to enumerate the scopes. If you are not seeing the expected objects in your browser, check your permission level.

See also:

- Including Objects in Your Scope

- Importing Computers to Your Scopes

- Refining Your Scope with Exclusions

- Filtering in the Browser

Including Objects in Your Scope

A valid discovery requires that you include at least one object. You can explicitly include high-level objects— domains, OUs and containers. This implicitly adds all computers in the selected object. For some discovery types, you can include objects using a query for more flexibility. For more information, see Using Queries to Define Your Scopes on page 59.

When you run the discovery Enterprise Reporter resolves the high level object to a list of targets, or in the case of an AD discovery, to a list of domains. These can be useful because if the contents of the container change, so do the targets of the discovery. Once Enterprise Reporter resolves this list, the other options in the scope can be applied.

i | **TIP:** It is strongly recommended that each target is included in only one discovery. Including a target in differently-configured discoveries can result in data loss. If you add individual targets, ensure that they have not already been implicitly included in another discovery by way of a domain, OU, or container.

Depending on the discovery type, you may also be able to select:

- Specific computers: You can drill into domains, OUs and containers and select individual computers. All options in the scope are applied to each selected computer, including global scopes.

- Folders and shares: You can drill into a computer and select folders and shares. All relevant options in the scope are applied.

    i | **NOTE:** If you select competing scopes, collection options define what will be collected. If you are collecting public shares and include a computer and a specific share, only the selected share will be collected. If you are collecting public shares and include a computer and a specific volume, both the volume and all public shares will be collected.

- Registry hives and keys: You can drill into a computer and select specific hives and keys. Because the local registry hives are not available for collection, they are not available to select. When you include a specific

hive or key, Enterprise Reporter interprets this to mean that is all you want to collect from the host computer; global scopes will not be applied.

- DFS Shares: A published Windows® Server DFS share can be added like any other share, from the System\Dfs-Configuration container within a domain. You can use the Browse dialog box to manually add all other DFS shares.

For more details on including objects in your scopes for each discovery type, see also:

- Step 2a. Choose scopes for your discoveries
- Step 2b: Choose scopes for your Office 365 discoveries

Importing Computers to Your Scopes

On the scopes page, instead of using the Add button to add one computer at a time, you may choose to import multiple computers to be included or excluded from collection for the following types of discoveries:

- Computer
- File Storage Analysis
- Microsoft® SQL
- NTFS
- Registry

You may import from any text file containing a list of the fully qualified computer name (or IP addresses) of the computers to be targeted with one computer per line, as in the following example:

        \\computer1.domain.com
        \\192.168.10.25
        computer3.domain.com

Only unique computer names free of invalid characters will be imported.

To import computers to your scopes

1. On the scopes page, click the **Import** button.
2. Select **Include Computers**.

    - OR -

    Select **Exclude Computers**.
3. Browse to locate the file containing the computers to be imported.
4. Select the file and click **Open** to start the import.
5. Review the import progress and results messages and click **OK** when complete.

    > **i** | **NOTE:** Any computer that is not successfully imported will be indicated by a red dot next to its name with an error description in the Message column.

6. Successfully imported computers will be visible in the Selected Scopes pane.
7. Optionally, fix any errors in the text file and repeat steps 1 through 5 until all computers are imported.

    > **i** | **NOTE:** There is no need to remove previously imported computers from the text file as only unique computer names will be loaded during subsequent imports.

Refining Your Scope with Exclusions

This step can be done in conjunction with the inclusions. Exclusions refine your scope further. Use the browser, a query, or the **Import** option to exclude scopes. For more information, see Importing Computers to Your Scopes on page 58.

For example, you can add a computer, but exclude a specific folder; or add a domain and exclude a specific OU. The following rules are used to process your exclusions:

- If you add an exclusion without the inclusion of a parent, it has no meaning and will be ignored. As long as the exclusion can be traced up the AD tree to a parent object, the exclusion can be processed.

- If you explicitly exclude a folder or share, that data will not be available for reporting.

- You should not collect the same shares or volumes in multiple discoveries as data loss can occur.

- When selecting a domain, OU or container as a target to collect, a whole computer can be excluded from the discovery.

- For an Active Directory® discovery, excluded objects are processed after included objects. Any OU or container that is included below an excluded one in the AD structure will not be included, as the exclude will take precedence.

Filtering in the Browser

It can be difficult to find the objects you are looking for when you are selecting your scope. To address this, you can use filtering in the browser. Filtering starts from your selected location, and is applied one level down. For example, if you filter at the domain level, the filter is applied to the first level containers. You remove filters one at a time, starting with the last filter you applied.

To apply a filter

1 Expand the tree, and then select your starting point.

   You must expand the first level of the selected node in order for the filter to work.

2 In the **Apply filter to** box, type your search string.

3 Click **Apply Filter**.

   > **NOTE:** You can also right-click any node on the tree and choose Filter.

To remove a filter

1 Click **Undo Filter** to remove the last filter applied.

   The details of the filter you will remove are displayed in the "Last Filter applied to" box.

2 Repeat if needed to remove other filters.

Using Queries to Define Your Scopes

Queries can be used on their own or to complement explicit scopes for the following types of discoveries:

- Computer
- File Storage Analysis
- NTFS
- Registry

Queries allow you to define a set of criteria that will be resolved when the discovery runs. For example, you can create a query that looks for computer names that contain "Finance" across an entire domain. When the discovery

runs, the query will execute and resolve a list of targets. Queries can be run against a particular domain, a container in a domain, or against all child domains of the domain you are currently logged in to. By default, queries create an included scope, but you can modify it to be excluded if you want.

> **i** | **NOTE:** Only queries based on computers, organization units and containers are relevant. Other parameters will be ignored.

To include a scope using a query

1   Click ➕ **Add**.

    If you get a warning message to select an Active Directory domain, click **OK** to close the message, then expand the Available Scopes list of Active Directory domains, select one and then click **Add** again.

2   Click **Add scope using a query**.

3   To choose a specific domain to target with the query, select it from the In list.

    - OR -

    To choose a specific container or OU to target, click **Browse**, locate the target and click **OK**.

    - OR -

    To target the current domain and all child domains, from the In list, select **Entire Directory**.

4   From the Find list, select **Computers**, **Organizational Units** or **Custom Search**, then make your selections on the tabs shown.

    Each selection you add builds a query. Ensure the Custom Search only includes computers, containers and organizational units. Note that on the Advanced tab of the Custom Search option, you can use any appropriate LDAP query. To test your query, click Find Now.

5   Click **OK** to save your query.

    The query is saved, and the domain against which it will run is shown in the scopes list. To view the query, you can hover over the domain name. If the query was run against the Entire Directory, the originating domain is shown, and all its trusted domains will be queried when the discovery is executed.

To exclude objects using a query

1   Follow the steps for adding a query.
    For more information, see .

2   In the Selected Scopes list, right-click the query and select **Exclude**.

Using Subnets to Define Your Scopes

You can use subnets on their own or to complement explicit scopes for the following types of discoveries:

* Computer
* File Storage Analysis
* NTFS
* Registry

Specifying a subnet allows you to specify a group of targets for a discovery based on their IP addresses. By default, a subnet mask is included in the scope, but you can modify it to be excluded.

> **i** | **NOTE:** Network and Broadcast IPs will not be included in the list of targets.

To include a scope using a subnet

1   Click ➕ **Add**.

2   Click **Add Subnet**.

3   Enter the digits that complete the subnet.

The mask, the number of IPs, and number of usable IPs are displayed.
The network and broadcast IPs that will be excluded from the list of targets are also displayed.

4   Optionally, click **Edit Subnet Credentials** to enter credentials that override the node credentials or the alternate credentials for this subnet. For more information, see Using Subnet Credentials on page 61.

5   Click **OK** to add this subnet to the selected scopes.

To edit a subnet scope

1   In the Selected Scopes listing, click on a subnet scope.

2   Click the **Edit** button.

3   Optionally, change the digits that complete the subnet.

4   Optionally, click **Edit Subnet Credentials** to change the credentials for this subnet.

To exclude objects using a subnet

1   Follow the steps **To include a scope using a subnet**.

2   In the Selected Scopes list, right-click the subnet and select **Exclude**.

Using Subnet Credentials

For each subnet scope in Computer, File Storage Analysis, NTFS, and Registry discoveries, you can define subnet credentials for a computer or group of computers using wildcard expressions. For example, you can enter one credential for all computers in a domain using an expression such as *.domain1.corp. During collection, the subnet credentials will be used to access target computers instead of the node credentials or the alternate credentials.

> **i** | **NOTE:** During collection, the subnet credentials will be processed in the order in which they are listed. If a computer happens to match two expressions within the subnet credentials, it will only be processed once using the topmost entry in the table. Place more specific entries toward the top of the list and entries using wildcard characters toward the bottom of the list.

To add subnet credentials

1   When adding or editing a subnet on the scopes page, click **Edit Subnet Credentials**.

2   Click **Add** to enter a new credential.

3   Enter the fully qualified name of a computer or an expression that represents a group of computers.

4   Click the **Account** ellipsis (...) to open the Credential Manager and add or select the credentials for the specified computers.

5   Click **OK** to return to the credential mappings.

6   Optionally, click **Move Up** or **Move Down** to position the entry appropriately.

7   Click **OK** to return to the Scopes page.

To edit subnet credentials

1   On the scopes page of a discovery, select a subnet and click **Edit**.

2   Click **Edit Subnet Credentials**.

3   Select the credential to edit and click **Edit**.

4   Change the **Computer name or expression** or the credential **Account.**

5   Click **OK** to return to the credential mappings.

6   Optionally, click **Move Up** or **Move Down** to re-position the entry.

7   Click **OK** to return to the Scopes page.

Step 2a. Choose scopes for your discoveries

In this section, we will explore choosing scopes for each discovery type.

- Choosing your Active Directory Scopes
- Choosing your Computer Scopes
- Choosing Your Exchange Scopes
- Choosing Your File Storage Analysis Scopes
- Choosing Your Microsoft SQL Scopes
- Choosing Your NTFS Scopes
- Choosing Your Registry Scopes

See also Step 2b: Choose scopes for your Office 365 discoveries.

Choosing your Active Directory Scopes

Active Directory® scopes determine what information will be collected when you run the discovery. There are several steps you should take to properly design your discovery.

See also:

- AD Discovery: Include scopes
- AD Discovery: Optionally refine your scope list with exclusions
- AD Discovery: Optionally select a domain controller
- AD Discovery: Decide what to collect from any domain in the discovery

AD Discovery: Include scopes

You can include domains, OUs and containers in your scope. You can collect the same domain information in more than one Active Directory discovery with an additional option to create multiple tasks for each domain. For full information on using the browser to add scopes, see Step 2b: Choose scopes for your Office 365 discoveries on page 88.

To explicitly include objects in your Active Directory scope

1   Click ➕ **Add**.

2   In the treeview, locate the desired object and select it.
    To select multiple objects, press Ctrl during selection.

    For more information, see Filtering in the Browser on page 59.

3   Click **Include** to add to your selected scopes list.

4   Click **OK**.

> ℹ️ | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly added. When you add an OU or container, all children are included.

AD Discovery: Optionally refine your scope list with exclusions

Active Directory® excluded objects are resolved after included objects. For example, if you include an object that exists below an excluded object in the AD structure, it will not be collected, as the exclusion will take precedence. For more information, see Using the Browser to Include and Exclude Scopes on page 57.

To explicitly exclude objects from your selected scope

1   Click ➕ **Add**.

2   Expand the treeview to locate the object you want to exclude.

    You can press Ctrl to select multiple objects.

3   Click **Exclude**.

AD Discovery: Optionally select a domain controller

In order to enumerate your domain, Enterprise Reporter looks to a domain controller. Depending on the domain controller chosen, the time it takes to perform the collection may vary. You can either allow Enterprise Reporter to choose the first available domain controller returned by Active Directory®, or you can select one that will optimize collection time. In this case, select a domain controller that is located physically close to the cluster you assigned to the discovery, or one that you know is a fast computer.

To set the domain controller used to enumerate a domain

1   Add your scopes.

    For more information, see AD Discovery: Include scopes on page 62 or AD Discovery: Optionally refine your scope list with exclusions on page 63.

2   In the scopes list for the included domain you want to configure, click "**...**"

3   To use the first available domain controller returned by Active Directory®, select **Automatic selection**.

    - OR -

    To assign a specific domain controller, choose **Select an available domain controller**, and then select a domain controller from the list.

4   Click **OK**.

AD Discovery: Decide what to collect from any domain in the discovery

When you run an Active Directory® discovery, it is resolved to a list of domains. You can collect a variety of information from these domains. The basic domain information, including a list of all OUs in the domain, is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 14. Active Directory Discovery: Optionall collect information for the following objects

Option	Notes	Performance Cost
Users		Medium
• Token groups count	Collect the count of items in the TokenGroups attribute to monitor access token sizes and to assist with group cleanup and consolidation projects.	High

• Query all domain controllers for last logon	The Active Directory® LastLogonTimestamp attribute is always collected for users and computers and can be used to determine if a user or computer account has recently logged onto the domain. To determine the actual time of the most recent logon for a user or computer account, enable this option to query all domain controllers for the LastLogon attribute. The LastLogon attribute is subsequently used to calculate the total number of times each domain user has successfully logged on (Number of Logons field).
	
	There are two known cases where the Number of Logons total may be inaccurate.
	
	  • Since the LastLogon attribute is not replicated, domain controllers that are retired, and are therefore not collected, may have contained logons for users that would be missing from the count.
	  • The attribute can track up to 65535 logons. After this limit has been reached, it becomes an inaccurate indicator of user activity on this domain controller.
		High

• In addition, you may choose to collect remote desktop information for the accounts collected.		Medium
• You may also choose to collect thumbnail photographs of domain users.		Medium
Groups and members	When this option is selected, groups and members are collected from the target domain.	High

Table 14. Active Directory Discovery: Optionall collect information for the following objects

Option	Notes	Performance Cost
• Include members from foreign domains	When this option is selected, all members of the group are collected, even if those members are from other domains or other forests.	Medium
▪ Collect nested groups and members from foreign domains	If Include members from foreign members is enabled, you can select this option to collect all nested groups and their members from foreign domains.	High
Computers		Medium
Domain Controllers		Low
Permissions	In addition to collecting Active Directory® objects, you may also choose to collect Active Directory® object permissions.	High
Contacts		Low
Trusts		Low
Sites		Low
Deleted Objects		Medium

Table 14. Active Directory Discovery: Optionall collect information for the following objects

Option	Notes	Performance Cost
Service Accounts		Low
Active Roles Virtual Attributes and Active Roles Server	When this option is selected, Active Roles virtual attributes are collected from all targets managed by the Active Roles Server. **NOTE:** Collecting Active Roles attributes for users, groups, computers, and organizational units can significantly increase the collection time. Some attributes that are known to increase collection time are: • edsaDialinAccessPermissions • edsaDialinApplyStaticRoutes • edsaDialinApplyStaticIP • edsaDialinCallbackNumber • edsaDialinCallbackOptions • edsaDialinCallerID • edsaDialinStaticIP • edsaDialinStaticRoutes • edsaDialinVerifyCallerID If Active Roles Virtual Attributes is enabled, you must specify the Active Roles Server that manages the target domains. Optionally, you may supply credentials to be used for contacting the Active Roles Server during collection. If these credentials are left empty, the node credentials or alternate credentials specified in on the Name page of the discovery in Step 1. Create the Discovery on page 54 will be used.	High

The following options further refine how collection tasks are handled.When collecting additional attributes, especially ones with a high performance cost, consider enabling these performance options to help improve collection performance.

Table 15. Active Directory Discovery: Additonal Performance Options

Options	Notes
Performance Options	

Table 15. Active Directory Discovery: Additonal Performance Options

Options	Notes
• Do not collect object counts	When this option is cleared, the discovery collects the total number of users, groups, contacts, computers, and organizational units. To collect summary counts of these objects and to be able to display them, reports such as Domain Summary, require this option to be cleared.
• Create multiple tasks for each domain	A collection task is always created for each domain in an Active Directory® discovery. Selecting this option will create individual tasks for domains, users, groups, computers, domain controllers, permissions, and deleted objects. This allows tasks to be load balanced using multiple nodes.

Choosing your Computer Scopes

Computer scopes determine what information will be collected when you run the discovery. There are several steps you should take to properly design your discovery.

See also:

- Computer Discovery: Include scopes
- Computer Discovery: Optionally refine your scope list with exclusions
- Computer Discovery: Decide what to collect from any computer in the discovery

Computer Discovery: Include scopes

You can specifically add domains, OUs, containers or computers. Domains, OUs and containers can contain many computers, which can significantly increase the time it takes to run the discovery.

For full information on using the browser to add scopes, see Step 2b: Choose scopes for your Office 365 discoveries on page 88. You can create scopes using a dynamic query, which is resolved when the discovery is run. This gives you the flexibility to describe the computers you want to target. For more information, see Using Queries to Define Your Scopes on page 59.

> **i** | **NOTE:** Use caution when creating your queries. Ensure that the resulting set of targets is not too large for a single discovery. As well, query results should not include computers included in another discovery of the same type. If a target is in more than one discovery of a particular type, rejected tasks will appear. For more information, see Data Collection Issues on page 127.

To explicitly include objects in your Computer scope

1 Click ✚ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2 In the treeview, locate the desired object and select it.
To select multiple objects, press Ctrl during selection.

For more information, see Filtering in the Browser on page 59.

- OR -

To add a scope using a query, expand the Add to scope menu, select **Add scope using a query**, and click Add. For more information, see Using Queries to Define Your Scopes on page 59.

- OR -

To include a scope using a subnet, expand the Add to scope menu, select **Add Subnet**, and click Add. For more information, see Using Subnets to Define Your Scopes on page 60.

3   Click **Include** to add to your selected scopes list.

4   Click **OK**.

> ℹ | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly added. When you add an OU or container, all children are included.

Computer Discovery: Optionally refine your scope list with exclusions

Exclusions refine the inclusions you have defined. You can do this optional step in conjunction with inclusions. For full details, see Refining Your Scope with Exclusions on page 59.

To explicitly exclude objects from your selected scope

1   Click ➕ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2   Expand the treeview to locate the object you want to exclude.

You can press Ctrl to select multiple objects.

3   Click **Exclude**.

Computer Discovery: Decide what to collect from any computer in the discovery

You can collect a variety of information from the computers in your discovery. The computer attributes (such as operating system details) are always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 16. Computer Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Printers		Low
Shares	**NOTE:** If you want to collect printer shares, you must enable both the Printers and the Shares check box.	High
Volumes		Medium
Accounts (Users and Groups)	**NOTE:** Accounts are not collected from domain controllers, as they are domain accounts associated with the computer, not computer accounts.	Medium
Installed Software		Medium
Microsoft Store Applications		Medium
Hotfixes		High

Table 16. Computer Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Policies		Medium
Services		High
Event Log Configuration	**NOTE:** Permissions are automatically collected where applicable. You can collect the members of any groups found during this collection by enabling the Collect group members check box.	
Extended WMI Entities		High

Choosing Your Exchange Scopes

Exchange® scopes determine what information will be collected when you run the discovery. There are several steps you should take to properly design your discovery.

See also:

- Exchange Discovery: Include Scopes
- Exchange Discovery: Optionally refine your scope list with exclusions
- Exchange Discovery: Decide what to collect from any server in the discovery

Exchange Discovery: Include Scopes

You can include Domains, Organizations, Exchange® Servers, in your scope. For full information on using the browser to add scopes, see Using the Browser to Include and Exclude Scopes on page 57.

To explicitly include objects in your Exchange scope

1 Click ➕ **Add**.

2 In the treeview, locate the desired object and select it.
   To select multiple objects, press Ctrl during selection.

   For more information, see Filtering in the Browser on page 59.

3 Click **Include** to add to your selected scopes list.

4 Click **OK**.

   ℹ | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly
   | added. When you add an OU or container, children are not included.

Exchange Discovery: Optionally refine your scope list with exclusions

Exclusions refine the inclusions you have defined. You can do this optional step in conjunction with inclusions. For full details, see Refining Your Scope with Exclusions on page 59.

To explicitly exclude objects from your selected scope

1 Click ➕ **Add**.

2 Expand the treeview to locate the object you want to exclude.

You can press Ctrl to select multiple objects.

3   Click **Exclude**.

Exchange Discovery: Decide what to collect from any server in the discovery

You can collect a variety of information from the computers in your discovery. Basic information from the root organization (the organization name and full LDAP path) is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 17. Exchange Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Mailboxes	Collects basic information for Mailboxes, Public Folder Mailboxes, and Mailbox Stores.	Medium
• Mailbox subfolders	If collecting mailboxes, you can enable this option to collect subfolders of Exchange mailboxes.	High
• Mailbox delegates	If collecting mailboxes, you can enable this option to collect mailbox delegates on Exchange mailboxes. This option excludes Mail-Enabled Users, Mail Contacts, Administrators, and Distribution Groups.	High
Mail-Enabled Users	Collects basic account information for Mail-Enabled Users.	Low
Mail Contacts	Collects basic account information for Mail Contacts.	Low
Public Folders	Collects basic information for Public Folders and, where applicable, Public Folder Stores.	Medium
• System Public Folders	If collecting public folders, you can enable this option to collect system public folders.	High
Distribution Groups	Collects basic account information for Distribution Groups.	Medium
Permissions	Collects permissions on Exchange objects such as mailboxes, public folders, organizations, administration groups, and servers. This option excludes permissions for Mail-Enabled Users, Mail Contacts, Administrators, and Distribution Groups.	Medium

Table 17. Exchange Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
• Mailbox AD permissions	If collecting permissions, you can enable this option to collect mailbox Active Directory permission information, including mailbox subfolder permissions.	High
• Mailbox Exchange permissions	If collecting permissions, you can enable this option to collect mailbox Exchange permission information.	High
• Mailbox folder permissions	If collecting permissions, you can enable this option to collect mailbox folder permission information.	High
• Public folder permissions	If collecting permissions, you can enable this option to collect public folder permission information.	High
Nested group members	Recursively collects the members of any groups found in the Exchange discovery.	High

The following options further refine how collection tasks are handled.When collecting additional attributes, especially ones with a high performance cost, consider enabling these performance options to help improve collection performance.

Table 18. Exchange Discovery: Performance Options

Option	Notes
Performance Options	When both sub-options are selected, one task per mailbox server is created. One additional task per object type (except Mailboxes) is also created.
	When neither option is selected, one task is created for the entire collection.
• Create multiple tasks per server	One task per mailbox server is created. One additional task that contains all other object types is also created.
• Create multiple tasks for each object type being created	One task per object type is created. For example, if there are Mailboxes, Distribution Groups, Mail-Enabled Users, Mail Contacts, and Public Folders, 5 tasks are created.
Mailboxes	
• Include system delegates	If collecting mailboxes, you can enable this option to collect all system delegates. Disabling this option will decrease collection time by excluding system delegates with the following account names from the collection: • Domain Admins • Enterprise Admins • Organization Management • Exchange Trusted Subsystem • Exchange Servers • System • Administrator
Public Folders	

Table 18. Exchange Discovery: Performance Options

Option	Notes
Permissions	
• Only collect explicit permissions	If collecting permissions, you can enable this option to collect permission information only for Exchange objects with explicit permissions.

The following options determine how target computers are resolved during discoveries.

Table 19. Discovery Resolution Options

Option	Notes
Resolution Options	
• Allow a configurable number of seconds for a target computer to respond to a ping request	Pings target computers at collection time to determine the existence of target computers on the network before processing tasks.
• Resolve target computers with disjoint namespaces	Verifies the existence of target computers with disjoint namespaces by using their DNS host names and, if that fails, by using their NetBIOS names.

Choosing Your File Storage Analysis Scopes

File Storage Analysis scopes determine what information will be collected when you run the discovery. There are several steps you should take to properly design your discovery.

See also:

- File Storage Analysis Discovery: Include scopes
- File Storage Analysis Discovery: Optionally refine your scope list with exclusions
- File Storage Analysis Discovery: Decide what to collect from any server in the discovery
- File Storage Analysis Discovery: Configure NAS Host Devices

File Storage Analysis Discovery: Include scopes

You can include domains, OUs, containers, individual computers, and volumes in your scope. For full information on using the browser to add scopes, see Using the Browser to Include and Exclude Scopes on page 57. You can create scopes using a dynamic query, which is resolved when the discovery is run. This gives you the flexibility to describe the computers you want to target. For more information, see Using Queries to Define Your Scopes on page 59.

> **i** | **NOTE:** Use caution when creating your queries. Ensure that the resulting set of targets is not too large for a single discovery. As well, query results should not include computers included in another discovery of the same type. If a target is in more than one discovery of a particular type, rejected tasks will appear. For more information, see Data Collection Issues on page 127.

To explicitly include objects in your File Storage Analysis scope

1 Click ➕ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2 In the treeview, locate the desired object and select it.
To select multiple objects, press Ctrl during selection.

For more information, see Filtering in the Browser on page 59.

- OR -

To add a scope using a query, expand the Add to scope menu, select **Add scope using a query**, and click Add. For more information, see Using Queries to Define Your Scopes on page 59.

- OR -

To include a scope using a subnet, expand the Add to scope menu, select **Add Subnet**, and click Add. For more information, see Using Subnets to Define Your Scopes on page 60.

3   Click **Include** to add to your selected scopes list.

4   Click **OK**.

> **i** | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly added. When you add an OU or container, all children are included.

File Storage Analysis Discovery: Optionally refine your scope list with exclusions

Exclusions refine the inclusions you have defined. You can do this optional step in conjunction with inclusions. For full details, see Refining Your Scope with Exclusions on page 59.

To explicitly exclude objects from your selected scope

1   Click ➕ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2   Expand the treeview to locate the object you want to exclude.

You can press Ctrl to select multiple objects.

3   Click **Exclude**.

File Storage Analysis Discovery: Decide what to collect from any server in the discovery

You can collect a variety of information from the computers in your discovery. The storage attributes and volume information are always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 20. File Storage Analysis Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Shares		
• Hidden Shares	If collecting shares, you can enable this option to collect hidden shares.	High

Table 20. File Storage Analysis Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
• Home Directories	If collecting shares, you can enable this option to collect home directories. **NOTE:** Selecting the Home Directory option automatically collects shares (including hidden shares), regardless of the other options selected.	High
Collect NAS using Configuration	If collecting targets on NAS host devices, enable this option and configure the NAS host device. For more information, see File Storage Analysis Discovery: Configure NAS Host Devices on page 74.	Medium
• If unable to collect volumes use shares as volumes		
Files		Medium
Folders		Medium
• Follow Directory Symbolic Links	If collecting folders, you can enable this option to collect and follow directory symbolic links.	Medium
Owners		Medium

File Storage Analysis Discovery: Configure NAS Host Devices

If you have NAS host devices that contain targets to be collected, you must enable the collection option and add the configuration and credentials for each host device. You can optionally omit NAS host device configuration and collect shares as volumes for NAS device targets in the discovery. Use this option when volumes cannot be retrieved or do not need to be collected.

See also:

- Configuring a NAS Host Device
- File Storage Analysis Discovery: Configure NAS Host Devices

To configure a NAS host device

1   On the Scopes page of the File Storage Analysis Discovery, select **Collect NAS using Configuration** and click **Configuration**.

2   Click **Add**.

- OR -

Select an existing device and click **Edit**.

3   Select the appropriate type of device.

4   Enter the Host IP address.

For FluidFS targets, enter the IP address of the computer where Enterprise Manager is installed.

5   Enter the Host port number.

6   Click the ellipsis (...) to use the Credential Manager to enter and select the credentials required to access this device. For more information, see Using the Credential Manager on page 36.

7   Click **OK** to close the Credential Manager.

8   Click **OK** to close the NAS Configuration details.

9   Click **Close** to close the NAS Host Device Manager.

To delete a NAS host device

1   On the Scopes page of the File Storage Analysis Discovery, select **Collect NAS using Configuration** and click **Configuration**.

2   Select an existing device and click **Delete**.

3   Click **Yes** to confirm the deletion.

4   Click **Close** to close the NAS Host Device Manager.

Choosing Your Microsoft SQL Scopes

The scopes of a SQL discovery can consist of database servers, instances or databases, in any combination. There are two steps for creating your scope:

- Select what to include.

- Select what to exclude. You can exclude targets explicitly for a server or instance, or globally for all servers.

    i | **TIP:** It is strongly recommended that each SQL Server® is included in only one discovery. Including a server in differently configured discoveries can result in data loss.

See also:

- MS SQL Discovery: Include scopes

- MS SQL Discovery: Optionally refine your scope list with exclusions (Scope Exclusions and Global Exclusions)

MS SQL Discovery: Include scopes

To explicitly include objects in your Microsoft SQL scope

1   On the Scopes page of the Create Discovery wizard, click ➕ **Add**.

    - OR -

    If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2   Click the drop down arrow in the Server selection field to display a list of the SQL Servers®, instances and databases known to Configuration Manager.

3   To add servers, select the servers and click **Add**.

    If your server does not appear on the list, you can type it in the form "SQLServerName" and click Add.

    Ensure this server is not included in any other discovery.

    - OR -

    To add instances, select the instances and click **Add**.

    If your instance does not appear on the list, you can type it in the form "SQLServerName\InstanceName" and click Add.

    Ensure that this is the only discovery where the contents of the host SQL Server® are included.

    - OR -

To add databases, select the databases and click **Add**.

If your instance does not appear on the list, you can type it in the form "SQLServerName\InstanceName\Database Name" and click Add.

Ensure that this is the only discovery where the contents of the host SQL Server® are included.

> **i** | **NOTE:** You can remove targets in both the Add SQL Scopes dialog box and the Scopes page of the Create Discovery wizard. Select the scopes from the list, and click Remove.

4   Click **OK** to close the scopes selection.

5   If desired, select **Collect Group Members**.

If you want to report on the members of any group included in SQL permissions, select this option. Collecting group members increases the discovery time.

6   Click **Next** to continue to the Exclusions page.

MS SQL Discovery: Optionally refine your scope list with exclusions (Scope Exclusions and Global Exclusions)

There are two ways to refine the scope of your discovery. Both are accessed on the Exclusions page of the Create Discovery wizard.

1   SQL scope exclusions allow you to specifically exclude:

- an instance or database within a server

- a database within an instance.

> **i** | **NOTE:** If you exclude a specific database or instance, that data will not be available for reporting, as you should not add the database or instance to another discovery.

2   Global exclusions allow you to exclude databases from all servers or instances included in the scope based on their names. By default, all databases that ship with Microsoft® SQL Server® are excluded.

> **i** | **NOTE:** To exclude an instance on a server, the server must be included in the scopes. To exclude a database, the server and instance must be included in the scopes.

To exclude specific instances or databases from your discovery (Scope exclusions)

1   On the Exclusions page of the Create Discovery wizard, click the **Scope** button.

2   Select the scope to modify, and click the ✚ **Modify** button.

3   To exclude instances, ensure the server is expanded, then select the instances and click **Add**.

The SQL instance "MSSQLServer" is the name Enterprise Reporter gives to an unnamed instance on the server.

- OR -

To exclude databases, expand the instance, select the databases and click **Add**.

> **i** | **NOTE:** You can restore removed targets to the scope by selecting them from the Excluded Scopes list, and clicking Remove.

4   Click **OK**.

5   Repeat steps 1 through 4 for any other scopes from which you want to set scope exclusions.

6   Click the **Global** button to set global exclusions.

- OR -

Click **Next** to continue to the Schedule page.

To exclude all databases with a specific name from your discovery (Global exclusions)

1   On the Exclusions page of the Create Discovery wizard, click the **Global** button.

2   In the text box, type the name of the database to exclude.

> **i** | **NOTE:** You can use the * and ? wildcards when listing global exclusions.

3   Click ✚ **Add**.

4   Click the **Scope** button to set scope exclusions.

- OR -

Click **Next** to continue to the Schedule page.

Choosing Your NTFS Scopes

NTFS scopes determine what information will be collected when you run the discovery. There are a number of steps you should take to properly design your discovery.

See also:

- NTFS Discovery: Best Practices
- NTFS Discovery: Include scopes
- NTFS Discovery: Optionally refine your scope list with exclusions
- NTFS Discovery: Select your global scopes
- NTFS Discovery: Decide what to collect from any computer in the discovery
- NTFS Discovery: Configure NAS Host Devices
- NTFS Discovery: Configure your file collection

NTFS Discovery: Best Practices

This section outlines some best practices to ensure the NTFS discovery and reports are optimized.

Hardware and software

- It is recommended to have Enterprise Reporter on its own SQL Server or at least a SQL Server that is not already used heavily by other products such as Change Auditor.

- When collecting Microsoft Office file attributes, Microsoft Office must be installed on the Enterprise Reporter node computer.

- It is recommended that each node have at least 100GB of hard drive space and 32GB of RAM.

- Ensure Enterprise Reporter has been updated with the latest hotfix. The product team regularly adds improvements to optimize discovery times.

Discovery setup and scheduling

- Create a separate discovery for each file server collected. This will allow the schedules to be staggered so that not all file servers are being collected and stored in the database at the same time.

  - To properly stagger the schedules, the discoveries may have to be run a few times to get a baseline of how long each discovery takes to run.

Discovery Scopes page: minimize information collected

- Only enable **Include hidden shares** if this data is needed. Selecting this option will cause a great increase in discovery time because all files on a server will be accessed through system shares such as C$ (C:\ drive) which are all the files and folders on the server.

- Folder Options
  - Only enable the option to collect **All folder levels** if it is required.
  - Set the **Folder depth** to limit the collection to only the levels of folder structure required.
- File Options
  - Only enable the option to **Collect files and their basic details like size and attributes** if file information is required.
- Permission Options
  - Only enable the option to **Calculate differences between folders and subfolders** if the File and Folder Permissions with Differences report is required.
- Advanced Options
  - Enterprise Reporter 3.0 (and later) has a performance option to break up tasks by share. Enable **Create a task per share for each computer** to allow multiple nodes/threads to collect one large file server.
    - Prior to Enterprise Reporter 3.0, one computer was considered to be one task. On these versions, if you are collecting 10 file servers, it is recommended to have one node per server that is manually configured to accept one task. This will force each node to collect one file server.

Running Reports

- When generating large NTFS reports, use the Export to CSV option. This will be much faster than generating the report in other formats.

NTFS Discovery: Include scopes

The scopes you add will determine the computers and specific folders that will be the targets of the discovery when it is run. You can explicitly add domains, OUs, containers, computers, folders, shares, DFS shares, or NAS devices. Domains, OUs and containers can contain many computers, which can significantly increase the time it takes to run the discovery. For more information, see Including Objects in Your Scope on page 57.

You can create scopes using a dynamic query, which is resolved when the discovery is run. This gives you the flexibility to describe the computers you want to target. For more information, see Using Queries to Define Your Scopes on page 59.

i | **NOTE:** Use caution when creating your queries. Ensure that the resulting set of targets is not too large for a single discovery. As well, query results should not include computers included in another discovery of the same type. If a target is in more than one discovery of a particular type, rejected tasks will appear. For more information, see Data Collection Issues on page 127.

To explicitly include objects in your NTFS scope

1 Click ✚ **Add**.

   - OR-

   If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2 In the treeview, locate the desired object and select it.
   To select multiple objects, press Ctrl during selection.

   For more information, see Filtering in the Browser on page 59.

   - OR -

   To add a scope using a query, expand the Add to scope menu, select **Add scope using a query**, and click Add. For more information, see Using Queries to Define Your Scopes on page 59.

- OR -

To include a scope using a subnet, expand the Add to scope menu, select **Add Subnet**, and click Add. For more information, see .

3   Click **Include** to add to your selected scopes list.

4   Click **OK**.

> **ℹ** | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly added. When you add an OU or container, all children are included.

To explicitly include a DFS share (NTFS discoveries only)

1   Click ➕ **Add**.

2   Click **Add to Scope | Add DFS share**.

3   Type the DFS root path using the fully qualified domain name.

4   Click **Include** to include the DFS share.

NTFS Discovery: Optionally refine your scope list with exclusions

Exclusions refine the inclusions you have defined. You can do this optional step in conjunction with inclusions. For full details, see .

To explicitly exclude objects from your selected scope

1   Click ➕ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in .

2   Expand the treeview to locate the object you want to exclude.

You can press Ctrl to select multiple objects.

3   Click **Exclude**.

To explicitly exclude a DFS share (NTFS discoveries only)

1   Click ➕ **Add**.

2   For Add to Scope option, select **Add DFS share**.

3   Click **Add**.

4   Type the DFS root path using the fully qualified domain name.

5   Click **Exclude** to exclude the DFS share.

NTFS Discovery: Select your global scopes

Global scopes are folders that are collected from every computer in the scope. When you run the discovery, Enterprise Reporter uses the scopes you selected in the first two steps to resolve a list of computers that will be targeted. Anything you include or exclude globally is applied to these computers. Use global scopes to:

• Collect common paths on all computers. For example, if you have included computers from different Windows® Operating Systems, the path to the Windows folder may be different. Instead of having to specifically add the Windows folder, you can choose the global [WINDOWS] scope, and Enterprise Reporter can determine where that folder is located on each computer.

The following table outlines the pre-defined filters available in the global scope selection.

Table 21. Filters in the global scope selection

Filter	Path selected
COMMONFILES	The location of the Common Files folder. Usually C:\Program Files\Common Files
PROGRAMFILES	The location of the Program Files folder. Usually C:\Program Files
SYSTEM	The location of the Windows System32 folder. Usually C:\Windows\System32
SYSTEMDRIVE	The drive where the operating system is installed. Usually C:\
WINDOWS	The location where Windows is installed. Usually C:\Windows

- Collect specific folders or the contents of specific shares on all computers. If there is a folder that you want to collect on many of the computers, it is more efficient to include it globally. If the folder or share is not on all targets, the inclusion will be ignored for those targets.

- Exclude common paths, specific folders or the contents of specific shares on all computers. You can make the same selections outlined above, but choose to exclude them instead of collecting them.

  i | **NOTE:** A global scope is not valid alone. You must first include a domain, OU, container, computer, NetApp filer, folder or share. See NTFS Discovery: Best Practices on page 77 for more information.

To add a global scope

1 From the Add Global list, choose the common path.

- OR -

From the Add Global list, choose the common path, then type "\" and a share or folder.

For example, select [WINDOWS] then type \Temp to include or exclude the temp folder on all computers.

- OR -

Type a share or folder.

Because these are going to be collected for all computers in the resolved discovery, you cannot include a computer name in the path. Invalid characters include "|", "[", and "]".

2 Click **Include** or **Exclude**.

You can press Enter to add to your list. Enter functions as the last button you clicked - for example, if you have just clicked Include, Enter includes your global scope.

3 Optionally select **Requirements | Collect only selected shares, folders and DFS shares** to suppress the collection of all public shares or volumes.

NTFS Discovery: Decide what to collect from any computer in the discovery

When you run an NTFS discovery, it is resolved to a list of computers and folders. Folders will be accessed using the share path and you must have read access to the share. Accessing folders using the share path is useful if shares are distributed, as in the case for DFS shares or Net App filers where administrative shares are disabled or not available.

The NTFS object that is being shared (for example a folder) is collected, and will be displayed in the report. For example, folders and files on a share called \\NYC_SVR\TrainingMaterials that is physically located on the computer named NYC_SVR in the path C:\HR\NYC\NewHires\TrainingMaterials would be accessed and displayed as \\NYC_SVR\TrainingMaterials.

For all computers in the discovery, you need to decide the starting point for your collection. This combines with your global scopes and recursion level to determine what folders are collected from each computer in the scope.

You must specify how to collect the information from each computer in the discovery. You can collect all available public shares, collect all available volumes, or collect only the selected shares, folders, and DFS shares. The following table outlines these options.

Table 22. NTFS Discovery: How to collect the information from each computer in the discovery

Option	Notes	Performance Cost
Collect all available public shares	The following data is collected: • All folders available through public shares on any computer included in the scope, directly or indirectly • Explicitly added folders • Folders in explicitly added shares, DFS shares, and NetApp filers • All globally included folders • All folders in globally included shares	Medium
• Include hidden shares		High
Collect all available volumes	All folders on any computer resolved from the scope are collected, unless: • they are specifically excluded from the discovery. • specific folders or shares are the only objects included on the computer. In this case, the specified folders or the folders in the specified shares are the only data that will be collected. NOTE: Based on your selected scope, if you are collecting against NAS devices, you have the option to use specified NAS Configurations.	High
Collect only selected shares, folders and DFS shares	This option is displayed when either: • both computers and global targets are included in the scope or • explicit shares or folders are included in the scope. When this option is selected, the NTFS discovery will only collect applicable folders and shares without performing the collection of all public shares and volumes.	Medium

The following table outlines the additional information that can be collected during the discovery. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 23. NTFS Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Folder Options	You have options to choose how deep into the tree the discovery will collect data.	
• All folder levels	The default is to collect all folder levels, starting from the included scope.	High
• Folder depth	You can collect just the root level by setting the folder depth to 0. You can choose the number of levels to collect, starting from the included scope, by setting the folder depth as desired. The root is not counted as a recursion level. If you have an excluded scope within an included scope, no folders below the exclusion are collected. In a very complicated nested set of includes and excludes along the same branch, the recursion level is reset with each includes scope.	Medium
File Options	You have options to choose what file information is collected.	
• Collect files and their basic details like size and attributes		Medium
• Collect advanced file metadata such as author and title	Option available if collecting files and their basic details. **Requirement:** Microsoft Office must be installed on the Enterprise Reporter node computer.	High
• Calculate duplicate files within the same computer	Option available if collecting files and their basic details.	High
Permission Options		
• Collect folder permissions		High
• Collect file permissions	Option available if collecting files and their basic details.	High
▪ Only collect and store files which have explicitly granted permissions	Selecting this option reduces the high performance cost of collecting all file permissions	
• Calculate permission differences between folders and subfolders (and files)	Collects the permissions that are Added, Modified, and Removed for a file or folder as compared to the parent. **NOTE:** Only the DACL and Owner permissions are collected.	High
• If a group account is found when collecting permissions, recursively collect group members.	You can choose to recursively collect the members of any groups found in the collection of file or folder permissions by selecting this option.	High

Table 23. NTFS Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Advanced Options		
• Create a task per share for each computer	A collection task is always created for each computer in an NTFS discovery. Selecting this option will create a task per share for each computer to allow tasks to be load balanced using multiple nodes.	

The following options further refines how collection tasks are handled. When collecting additional attributes, especially ones with a high performance cost, consider enabling these performance options to help improve collection performance.

Table 24. NTFS Discovery: Performance Options

Performance Option	Notes
Folder Options	You have options to choose how deep into the tree the discovery will collect data.
• All folder levels	The default is to collect all folder levels, starting from the included scope.
• Folder depth	You can collect just the root level by setting the folder depth to 0.
	You can choose the number of levels to collect, starting from the included scope, by setting the folder depth as desired. The root is not counted as a recursion level. If you have an excluded scope within an included scope, no folders below the exclusion are collected. In a very complicated nested set of includes and excludes along the same branch, the recursion level is reset with each includes scope.
Permission Options	
• Only collect and store files which have explicitly granted permissions	Selecting this option reduces the high performance cost of collecting all file permissions.
Advanced Options	
• Create a task per share for each computer	A collection task is always created for each computer in an NTFS discovery. Selecting this option will create a task per share for each computer to allow tasks to be load balanced using multiple nodes.

The following options determine how target computers are resolved during discoveries.

Table 25. Discovery Resolution Options

Option	Notes
Resolution Options	
• Allow a configurable number of seconds for a target computer to respond to a ping request	Pings target computers at collection time to determine the existence of target computers on the network before processing tasks.
• Resolve target computers with disjoint namespaces	Verifies the existence of target computers with disjoint namespaces by using their DNS host names and, if that fails, by using their NetBIOS names.

The scopes are now selected and configured. It is a good idea to review your scopes before continuing.

NTFS Discovery: Configure NAS Host Devices

If you have selected to Collect all available volumes for an NTFS discovery and the targets to be collected contain NAS host devices, you must add the configuration and credentials for each NAS host device.

See also:

- Configuring a NAS Host Device
- File Storage Analysis Discovery: Configure NAS Host Devices

To configure a NAS host device

1 Select **Collect all available volumes**.

   An additional question is displayed, "Based on your selected scope, if you are collecting against NAS devices, do you want to use specified NAS Configurations?"

2 Select **Yes**.

3 Click **Configure your global NAS Configuration**.

4 Click **Add**.

   - OR -

   Select an existing device and click **Edit**.

5 Select the appropriate type of device.

6 Enter the Host IP address.

   For FluidFS targets, enter the IP address of the computer where Enterprise Manager is installed.

7 Enter the Host port number.

8 Click the ellipsis (...) to use the Credential Manager to enter and select the credentials required to access this device. For more information, see Using the Credential Manager on page 36.

9 Click **OK** to close the Credential Manager.

10 Click **OK** to close the NAS Configuration details.

11 Click **Close** to close the NAS Host Device Manager.

To delete a NAS host device

1 Select **Collect all available volumes**.

   An additional question is displayed, "Based on your selected scope, if you are collecting against NAS devices, do you want to use specified NAS Configurations?"

2 Select **Yes**.

3 Click **Configure your global NAS Configuration**.

4 Select an existing device and click **Delete**.

5 Click **Yes** to confirm the deletion.

6 Click **Close** to close the NAS Host Device Manager.

7

NTFS Discovery: Configure your file collection

You can decide what files, if any, to collect from the folders you targeted. You can use both wildcards and regular expressions to include or exclude groups of files. These are applied to every scope in the discovery. Both explicit and inherited permissions are collected for all files included in the discovery.

To collect all files

- Click **Collect files and their basic details like size and attributes**.

  Include *.* appears in the list.

To use wildcards to include or exclude files

1 Ensure that **Using wildcards** is selected. If necessary, click to change.

2 Type in the desired pattern using the acceptable wildcards.

  Use * to replace any number of characters, and ? to replace a single character.

3 Click **Include** or **Exclude**.

  Your file pattern appears on the list, replacing *.*.

  You can press Enter to add to your list. Enter functions as the last button you clicked—for example, if you have just clicked Include, Enter includes your file pattern.

To use regular expressions to include or exclude files

1 Ensure you **Using Regex** is selected. If necessary, click to change.

2 Type in the desired expression, using Microsoft$^®$ .NET Framework or Perl$^®$ 5 syntax.

  For more information, search Microsoft.com for Microsoft$^®$ .NET Framework regular expressions.

3 Click **Include** or **Exclude**.

To remove a file pattern

1 Select the file pattern on the list.

2 Click **Remove**.

To collect file permission data

1 Select **File Options | Collect files and their basic details like size and attributes**.

2 Select **Permission Options | Collect file permissions**.

3 To refine your collection of files, select **Only collect and store files which have explicitly granted permissions**.

  This option may shorten the time your discovery takes to run.

  > **NOTE:** If you select **Only collect and store files which have explicitly granted permissions**, files that only have inherited permissions are ignored when you run the discovery. These files will not be available to reporting users.

  > **NOTE:** To collect the members of all groups included in the file permissions, make sure **If a group is found when collecging permissions, recursively collect group members is selected** on the Scopes page.

  > > **NOTE:** Only the DACL and Owner permissions are collected.

To collect duplicate file information

1 Select **File Options | Collect files and their basic details like size and attributes**.

2 Select **File Options | Calculate duplicate files within the same computer.**

  > **NOTE:** Enabling this option locates duplicate files per computer within a single discovery by performing CRC checks to compare files that have both the same name and the same size. Collecting duplicate file information increases the discovery time.

To collect file details

    1   Select **File Options | Collect files and their basic details like size and attributes**.

    1   Select **File Options | Collect advanced file metadata such as author and title**.

Choosing Your Registry Scopes

Registry scopes determine what information will be collected when you run the discovery. There are several steps you should take to properly design your discovery.

See also:

- Registry Discovery: Include scopes
- Registry Discovery: Optionally refine your scope list with exclusions
- Registry Discovery: Select your global scopes
- Registry Discovery: Decide what to collect from any computer

Registry Discovery: Include scopes

You can specifically add domains, OUs, containers, computers, hives and keys. For full information on using the browser to add scopes, see Using the Browser to Include and Exclude Scopes on page 57.

You can create scopes using a dynamic query, which is resolved when the discovery is run. This gives you the flexibility to describe the computers you want to target. For more information, see Using Queries to Define Your Scopes on page 59.

> **i** | **NOTE:** Use caution when creating your queries. Ensure that the resulting set of targets is not too large for a single discovery. As well, query results should not include computers included in another discovery of the same type. If a target is in more than one discovery of a particular type, rejected tasks will appear. For more information, see Data Collection Issues on page 127.

To explicitly include objects in your Registry scope

    1   Click ➕ **Add**.

        - OR-

        If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

    2   In the treeview, locate the desired object and select it.
        To select multiple objects, press Ctrl during selection.

        For more information, see Filtering in the Browser on page 59.

        - OR -

        To include a scope using a subnet, expand the Add to scope menu, select **Add Subnet**, and click Add. For more information, see Using Subnets to Define Your Scopes on page 60.

    3   Click **Include** to add to your selected scopes list.

    4   Click **OK**.

>   **i** | **NOTE:** When you add a domain, its child domains are not included - each domain must be explicitly added. When you add an OU or container, all children are included.

Registry Discovery: Optionally refine your scope list with exclusions

Exclusions refine the inclusions you have defined. You can do this optional step in conjunction with inclusions. For more information, see Refining Your Scope with Exclusions on page 59.

To explicitly exclude objects from your selected scope

1   Click ➕ **Add**.

- OR-

If you have a text file containing the computers to be selected, click **Import** and follow the steps as outlined in Importing Computers to Your Scopes on page 58.

2   Expand the treeview to locate the object you want to exclude.

You can press Ctrl to select multiple objects.

3   Click **Exclude**.

Registry Discovery: Select your global scopes

Global scopes are hives or keys that are collected from every computer in the scope. When you run the discovery, Enterprise Reporter uses the scopes you selected in the first two steps to resolve a list of computers that will be targeted. Anything you include or exclude globally is applied to these computers. Use global scopes to:

•   Collect a registry hive on all computers. User specific hives cannot be collected; only available hives are listed.

•   Collect specific registry keys on all computers. If there is a key that you want to collect on many of the computers, it is more efficient to include it globally. If the key is not on all targets, the inclusion will be ignored for those targets.

•   Exclude hives or keys on all computers. You make the same selections outlined above, but choose to exclude them instead of collecting them.

> **i** | **NOTE:** A global scope is not valid alone. You must first include a domain, OU, container, computer, hive or key. See Registry Discovery: Include scopes on page 86 for more information.

> **i** | **NOTE:** If you add a global exclusion that conflicts with an explicit inclusion from Step 1, the explicit inclusion will be processed, and the global exclusion will be ignored.

To add a global scope

1   From the Add Global list, choose a hive.

- OR -

From the Add Global list, choose the hive, then type "\" and a key.

For example, select HKEY_LOCAL_MACHINE then type \Software to include all software registry keys on all computers.

- OR -

Type a hive and key.

Because these are going to be collected for all computers in the resolved discovery, you cannot include a computer name in the path. Invalid characters include "|", "[", and "]".

2   Click **Include** or **Exclude**.

You can press Enter to add to your list. Enter functions as the last button you clicked - for example, if you have just clicked Include, Enter includes your global scope.

Registry Discovery: Decide what to collect from any computer

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 26. Registry Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Collect values	This option is selected to collect key values. You can disable this option to collect only registry keys and enhance discovery performance.	Medium
Key permissions		High
• Group members	If collecting key permissions, you can enable this option to recursively collect the members of any groups found in your discovery.	High

The following options further refine how collection tasks are handled.When collecting additional attributes, especially ones with a high performance cost, consider enabling this performance option to help improve collection performance.

Table 27. Registry Discovery: Performance Options

Performance Option	Notes
Recursion options	
• All branches	
• Branch depth	Large registry hives can take a while to collect, so you can improve performance by restricting the number of branches using the Branch depth option.

The following options determine how target computers are resolved during discoveries.

Table 28. Discovery Resolution Options

Option	Notes
Resolution Options	
• Allow a configurable number of seconds for a target computer to respond to a ping request	Pings target computers at collection time to determine the existence of target computers on the network before processing tasks.
• Resolve target computers with disjoint namespaces	Verifies the existence of target computers with disjoint namespaces by using their DNS host names and, if that fails, by using their NetBIOS names.

Step 2b: Choose scopes for your Office 365 discoveries

In this section, we will explore choosing scopes for each discovery type.

- Choosing Your Azure Active Directory Scopes
- Choosing Your Azure Resource Scopes
- Choosing Your Exchange Online Scopes
- Choosing Your Microsoft Teams Scopes
- Choosing Your OneDrive Scopes

See also Step 2a. Choose scopes for your discoveries

> **i** | **NOTE:** Azure Active Directory, Azure Resource, and OneDrive discoveries are multi-threaded to optimize performance and require an Intel® or AMD 2 GHz multiprocessor with at least 2 cores.

Choosing Your Azure Active Directory Scopes

Azure Active Directory discoveries, by default, target the tenant that you specified with the credentials on the Name page in Step 1. Create the Discovery on page 54. You can edit the name of the Azure Active Directory tenant if it is different from the credentials specified. Information about your contacts, users, groups, group members, roles, and application service principals is collected based on the discovery options.

See also:

- Azure AD Discovery: Configuring Enterprise Reporter Applications for Azure
- Azure AD Discovery: Decide what to collect from the tenant in the discovery

Azure AD Discovery: Configuring Enterprise Reporter Applications for Azure

You must register the applications used by Enterprise Reporter in your Azure environment. You can complete this process using the System | Configuration | Tenant Application Management option in the main menu before you create discoveries.

If the Enterprise Reporter Azure application has not yet been registered in your Azure environment, the Azure Active Directory Scopes page of the discovery will display a warning message above the Azure Tenant name at the top of the screen. If you have already configured the required tenant application and it indicates the option to **Reconfigure**, nothing further is necessary.

> **i** | **IMPORTANT:** If changes are made to Azure applications outside of Enterprise Reporter (for example, the removal of permissions or the deletion of applications), Enterprise Reporter will not receive these updates. Please use the Reconfigure link next to an application to re-create the application and consent to application permissions again. For more information, see To reconfigure application for a tenant on page 90.

To open the Tenant Application Manager from an Azure discovery

1 On the Azure discovery Scopes page, click **Configure** next to the tenant name to open the Tenant Application Manager.

To add a tenant

1 Click **Add**.

2 Enter the name of the tenant where Enterprise Reporter applications will be created.

3 Click **OK**.

To delete a tenant

1 Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

Azure AD Discovery: Decide what to collect from the tenant in the discovery

You can collect a variety of information from the Office 365 tenant in your discovery. Basic information from the tenant (the tenant name and full LDAP path), subscriptions, licenses, and service plans is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 29. Azure Active Directory Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Users	Collects basic information for all user accounts.	Medium
• Additional attributes from Office 365	If collecting users, you can enable this option to collect additional attributes such as AboutMe, Birthday, HireDate, Interests, MySite, PastProjects, PreferredName, Responsibilities, Schools, and Skills.	High
• Multi-factor authentication attributes	If collecting users, you can enable this option to collect multi-factor authentication information. **NOTE:** Microsoft Azure Active Directory Module for Windows PowerShell is required on the node machines to collect multi-factor authentication attributes for Azure Users.	Medium
• Users flagged as risky	If collecting users, you can enable this option to collect users flagged for risk.	Medium
Groups	Collects basic information for all groups including direct members.	Medium
• Additional attributes from Office 365	If collecting groups, you can enable this option to collect additional attributes such as AutoSubscribeNewMembers, and IsSubscribedByMail.	High
Contacts	Collects basic information for all contacts.	Low
Devices	Collects basic information for all devices and their users.	Medium
Roles	Collects basic information for all standard roles including role direct members.	Low
Applications and Service Principals	Collects information for all applications and service principals including the user and group assignments.	High

Choosing Your Azure Resource Scopes

Azure Resource discoveries, by default, target the tenant that you specified with the credentials on the Name page in Step 1. Create the Discovery on page 54. You can edit the name of the tenant if it is different from the credentials specified. Information about your resources is collected based on the discovery options.

See also:

- Azure Resource Discovery: Configuring Enterprise Reporter Applications for Azure
- Azure Resource Discovery: Include Scopes

- Azure Resource Discovery: Optionally refine your scope list with exclusions
- Azure Resource Discovery: Decide what to collect from the subscriptions in the discovery

Azure Resource Discovery: Configuring Enterprise Reporter Applications for Azure

You must register the applications used by Enterprise Reporter in your Azure environment. You can complete this process using the System | Configuration | Tenant Application Management option in the main menu before you create discoveries.

If the Enterprise Reporter Azure and Resource applications have not yet been registered in your Azure environment, the Azure Resource Scopes page of the discovery will display a warning message above the Azure Tenant name at the top of the screen. If you have already configured the required tenant application and it indicates the option to **Reconfigure**, nothing further is necessary.

To open the Tenant Application Manager from an Azure Resource discovery

1   On the Azure Resource discovery Scopes page, click **Configure** next to the tenant name to open the Tenant Application Manager.

To add a tenant

1   Click **Add**.

2   Enter the name of the tenant where Enterprise Reporter applications will be created.

3   Click **OK**.

To delete a tenant

1   Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

    The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1. Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2. Click **Reconfigure** next to the application.

3. Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

   If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

   Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

Azure Resource Discovery: Include Scopes

By default, the discovery collects information from the subscriptions you specify. You can change the scope to collect all subscriptions. For full information on using the browser to add scopes, see Step 2b: Choose scopes for your Office 365 discoveries on page 88.

To specify whether to collect all subscriptions or only specific subscriptions

- Select **Collect only selected subscriptions** to explicitly include individual subscriptions in the scope.

  > **NOTE:** This option is recommended when the majority of the subscriptions on the target tenant do not need to be collected. For more information, see To explicitly include objects in your Azure Resource scope on page 93.

  - OR-

  Select **Collect all subscriptions with the option to exclude selected subscriptions** to collect all available information from most or all of the subscriptions on the target tenant.

  > **NOTE:** This option is recommended when the majority of the subscriptions on the target tenant need to be collected. For more information, see Azure Resource Discovery: Optionally refine your scope list with exclusions on page 93.

To explicitly include objects in your Azure Resource scope

1. Click ➕ **Add** to populate the subscription list with all subscriptions for all credentials provided on the Name page.

2. In the treeview, locate the object to include and select it.

   To select multiple objects, press **Ctrl** during selection.

3. Click **Include** to add to your selected scopes list.

4. Click **OK**.

Azure Resource Discovery: Optionally refine your scope list with exclusions

Exclusions refine the scope by limiting the collection to only the necessary information.

To explicitly exclude objects from your selected Azure Resource scope

1   Click ➕ **Add** to populate the subscription list with all subscriptions for all credentials provided on the Name page.

2   In the treeview, locate the object you want to exclude.

To select multiple objects, press **Ctrl** during selection.

3   Click **Exclude**.

Azure Resource Discovery: Decide what to collect from the subscriptions in the discovery

Basic information from the tenant (the tenant name and full LDAP path) is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 30. Azure Resource Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Virtual machines		Medium
Disks		Medium
Networking		Medium
Storage Accounts		Medium
Access control	Collects basic information about access control.	Medium
• Nested group members	If collecting access control information, you can enable this option to collect group member information.	Medium

Choosing Your Exchange Online Scopes

Exchange Online discoveries, by default, target the Office 365 tenant that you specified with the credentials on the Name page in . You can edit the name of the Office 365 tenant if it is different from the credentials specified. Information about your online mailboxes, public folders, contacts, distribution groups, and distribution group members is collected based on the discovery options.

See also:

* Exchange Online Discovery: Configuring Enterprise Reporter Applications for Azure

* Exchange Online Discovery: Decide what to collect from the tenant in the discovery

Exchange Online Discovery: Configuring Enterprise Reporter Applications for Azure

You must register the applications used by Enterprise Reporter in your Azure environment. You can complete this process using the System | Configuration | Tenant Application Management option in the main menu before you create discoveries.

If the Enterprise Reporter Azure application has not yet been registered in your Azure environment, the Exchange Online Scopes page of the discovery will display a warning message above the Azure Tenant name at the top of

the screen. If you have already configured the required tenant application and it indicates the option to **Reconfigure**, nothing further is necessary.

> **i** | **IMPORTANT:** If changes are made to Azure applications outside of Enterprise Reporter (for example, the removal of permissions or the deletion of applications), Enterprise Reporter will not receive these updates. Please use the Reconfigure link next to an application to re-create the application and consent to application permissions again. For more information, see To reconfigure application for a tenant on page 90.

To open the Tenant Application Manager from an Azure discovery

1   On the Azure discovery Scopes page, click **Configure** next to the tenant name to open the Tenant Application Manager.

To add a tenant

1   Click **Add**.

2   Enter the name of the tenant where Enterprise Reporter applications will be created.

3   Click **OK**.

To delete a tenant

1   Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

    The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1 Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2 Click **Reconfigure** next to the application.

3 Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

   If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

   Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

Exchange Online Discovery: Decide what to collect from the tenant in the discovery

You can collect a variety of information from the Office 365 tenant in your discovery. Basic information from the tenant (the tenant name and full LDAP path) is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 31. Exchange Online Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Mailboxes	Collects basic information for Mailboxes and Public Folder Mailboxes (except FirstName and LastName fields).	Medium
• Mailbox delegates	If collecting mailboxes, you can enable this option to collect mailbox delegates on Exchange mailboxes. This option excludes Mail-Enabled Users, Mail Contacts, Administrators, and Distribution Groups.	High
• Mailbox statistics	If collecting mailboxes, you can enable this option to collect mailbox statistics on Exchange mailboxes.	High
• Mailbox folders	If collecting mailboxes, you can enable this option to collect folders in Exchange mailboxes.	High
Public Folders	Collects basic information for Public Folders.	Medium
Mail-Enabled Users	Collects basic account information for Mail-Enabled Users (except FirstName, FullName, LastName, and Initials fields).	Low
Mail Contacts	Collects basic account information for Mail Contacts (except FirstName, FullName, LastName, and initials fields).	Low

Table 31. Exchange Online Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Distribiution Groups	Collects basic account information for Distribution Groups.	Medium
• Group members	If collecting distribution groups, you can enable this option to recursively collect the members of any groups found in the Exchange Online discovery.	High
• Dynamic group members	If collecting distribution groups, you can enable this option to recursively collect the members of any dynamic groups found in the Exchange Online discovery.	High
Permissions	Collects permissions on Exchange objects such as mailboxes and public folders. This option excludes permissions for Mail-Enabled Users, Mail Contacts, and Distribution Groups.	High
• Mailbox folder permissions	If collecting permissions, you can enable this option to collect mailbox folder permissions.	High
Nested group members	Recursively collects the members of any groups found in the Exchange Online discovery.	High

The following options further refine how collection tasks are handled.When collecting additional attributes, especially ones with a high performance cost, consider enabling this performance option to help improve collection performance.

Table 32. Exchange Discovery: Performance Options

Performance Option	Notes
Collect only files with explicit permissions	If collecting permissions, you can enable this option to collect permission information only for Exchange Online objects with explicit permissions.
	Selecting this option reduces the high performance cost of collecting all file permissions.

Choosing Your Microsoft Teams Scopes

Azure Resource discoveries, by default, target the tenant that you specified with the credentials on the Name page in Step 1. Create the Discovery on page 54. You can edit the name of the tenant if it is different from the credentials specified. Information about your resources is collected based on the discovery options.

See also:

- Microsoft Teams Discovery: Configuring Enterprise Reporter Applications for Azure

- Microsoft Teams Discovery: Include Scopes

- Microsoft Teams Discovery: Decide what to collect from the subscriptions in the discovery

Microsoft Teams Discovery: Configuring Enterprise Reporter Applications for Azure

You must register the applications used by Enterprise Reporter in your Azure environment. You can complete this process using the System | Configuration | Tenant Application Management option in the main menu before you create discoveries.

If the Enterprise Reporter Azure and Teams applications have not yet been registered in your Azure environment, the Microsoft Teams Scopes page of the discovery will display a warning message above the Azure Tenant name at the top of the screen. If you have already configured the required tenant application and it indicates the option to **Reconfigure**, nothing further is necessary.

To open the Tenant Application Manager from an Microsoft Teams discovery

1   On the Microsoft Teams discovery Scopes page, click **Configure** next to the tenant name to open the Tenant Application Manager.

To add a tenant

1   Click **Add**.

2   Enter the name of the tenant where Enterprise Reporter applications will be created.

3   Click **OK**.

To delete a tenant

1   Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

    The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1. Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2. Click **Reconfigure** next to the application.

3. Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

   If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

   Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

Microsoft Teams Discovery: Include Scopes

Microsoft Teams discoveries, by default, target the Azure tenant that you specified with the credentials on the Name page in Step 1. Create the Discovery on page 54. You can edit the name of the Azure tenant if it is different from the credentials specified. Information about your teams, channels, owners, applications, files, and folders is collected based on the discovery options.

Microsoft Teams Discovery: Decide what to collect from the subscriptions in the discovery

Basic information from the tenant (the tenant name and full LDAP path) is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 33. Microsoft Teams Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Channels		Medium
Applications		Medium
Files and Folders		Medium

Choosing Your OneDrive Scopes

OneDrive discoveries, by default, target the Office 365 tenant that you specified with the credentials on the Name page in Step 1. Create the Discovery on page 54. You can edit the name of the Office 365 tenant if it is different from the credentials specified. Information about your files, folders, file and folder permissions, and configuration settings is collected based on the discovery options.

See also:

- OneDrive Discovery: Configuring Enterprise Reporter Applications for Azure

- OneDrive Discovery: Include Scopes

- OneDrive Discovery: Optionally refine your scope list with exclusions

- OneDrive Discovery: Decide what to collect from the drives in the discovery

OneDrive Discovery: Configuring Enterprise Reporter Applications for Azure

You must register the applications used by Enterprise Reporter in your Azure environment. You can complete this process using the System | Configuration | Tenant Application Management option in the main menu before you create discoveries.

If the Enterprise Reporter Azure and OneDrive applications have not yet been registered in your Azure environment, the OneDrive Scopes page of the discovery will display a warning message above the Azure Tenant name at the top of the screen. If you have already configured the required tenant application and it indicates the option to **Reconfigure**, nothing further is necessary.

To open the Tenant Application Manager from a OneDrive discovery

1   On the OneDrive discovery Scopes page, click **Configure** next to the tenant name to open the Tenant Application Manager.

To add a tenant

1   Click **Add**.

2   Enter the name of the tenant where Enterprise Reporter applications will be created.

3   Click **OK**.

To delete a tenant

1   Select the name of the tenant to be deleted.

2   Click **Delete**.

3   Click **Yes** to confirm that discoveries for any tenant that you delete will be unable to collect data

    The tenant and the Enterprise Reporter applications for this tenant are removed.

To edit administrator credentials for the tenant application

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** or **Reconfigure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To configure applications for a tenant

1   Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2   Click **Configure** next to the application.

3   Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

    If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

    Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

To reconfigure application for a tenant

1 Open the Tenant Application Manager. For more information, see To open the Tenant Application Manager on page 47.

2 Click **Reconfigure** next to the application.

3 Complete the Microsoft wizard that guides you through registering the application and consenting to application permissions using credentials with administrative access to create Enterprise Reporter applications on the tenant.

   If multi-factor authentication has been configured for this account, you will be prompted to complete the verification steps to complete the authentication process.

   Once the consent is complete, the Tenant Application Manager will display a **Reconfigure** link next to the application.

OneDrive Discovery: Include Scopes

By default, the discovery collects information from the drives you specify. You can change the scope to collect all drives. For full information on using the browser to add scopes, see Step 2b: Choose scopes for your Office 365 discoveries on page 88.

To specify whether to collect all drives or only specific drives

- Select **Collect only selected drives** to explicitly include individual drives in the scope.

  > **i** | **NOTE:** This option is recommended when the majority of the drives on the target tenant do not need to be collected. For more information, see To explicitly include objects in your OneDrive scope on page 101.

  - OR-

  Select **Collect all drives with the option to exclude selected drives** to collect all available information from most or all of the drives on the target tenant.

  > **i** | **NOTE:** This option is recommended when the majority of the drives on the target tenant need to be collected. For more information, see OneDrive Discovery: Optionally refine your scope list with exclusions on page 101.

To explicitly include objects in your OneDrive scope

1 Click ➕ **Add**.

2 Type a letter or combination of letters and click **Search** to populate the treeview. For example, type the letter **p** to return all drives that start with the letter p, or type **po** to return all drives that start with the letters po. Entering additional letters will narrow the search results.

3 In the treeview, locate the desired object and select it.
   To select multiple objects, press Ctrl during selection.

4 Click **Include** to add to your selected scopes list.

5 Click **OK**.

OneDrive Discovery: Optionally refine your scope list with exclusions

Exclusions refine the scope by limiting the collection to only the necessary information.

To explicitly exclude objects from your selected OneDrive scope

1 Click ➕ **Add**.

2   Enter search criteria and click **Search** to populate the treeview.

> **NOTE:** The drives of users whose UserPrincipalName, DisplayName, GivenName, or Surname starts with the search criteria you provide will be displayed in the treeview. For example, if the DisplayName is John Sitraka, entering "sit" will display this drive; however, entering "rak" will not.

3   Expand the treeview to locate the object you want to exclude.

You can press Ctrl while you click to select multiple objects.

4   Click **Exclude**.

OneDrive Discovery: Decide what to collect from the drives in the discovery

Basic information from the tenant (the tenant name and full LDAP path) is always collected.

The following table outlines the additional information that can be collected. Collecting additional information impacts discovery performance. Options with a high performance cost will slow discovery performance more than options with a medium or low performance cost.

Table 34. OneDrive Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
Folders	Collects basic information for OneDrive folders.	Medium
• Files	If collecting folders, you can enable this option to collect basic information for One Drive files.	Medium
Permissions	Collects permissions on OneDrive roots, folders, and files.	High

Table 34. OneDrive Discovery: Optionally collect information for the following objects

Option	Notes	Performance Cost
• Nested group members	Recursively collects the members of any groups found in the OneDrive discovery.	High
Configuration settings	Collects OneDrive configuration information for settings such as sharing, sync, storage, device access, and notifications. **NOTE:** PowerShell 3.0 and Microsoft SharePoint Online Management Shell must be installed on the node that will be used to perform the collection of this information. **NOTE:** In order for OneDrive configuration settings to be collected successfully, an authorized connection must be established to the SharePoint Online service. To allow for credentials to be specified for your tenant, the **LegacyAuthProtocols** setting must be enabled. To set this on your tenant, run the following commands using the Microsoft SharePoint Online Management Shell. This action must be performed on the node machine with Microsoft SharePoint Online Management Shell installed. • Set-ExecutionPolicy -Scope Process -ExecutionPolicy RemoteSigned • Import-Module -Name Microsoft.Online.SharePoint.PowerShell • Connect-SPOService -Url "<full tenant name>" • Set-SPOTenant -LegacyAuthProtocolsEnabled $True • Disconnect-SPOService	Medium

The following options further refine how collection tasks are handled.When collecting additional attributes, especially ones with a high performance cost, consider enabling this performance option to help improve collection performance.

Table 35. OneDrive Discovery: Performance Options

Performance Option	Notes
Active Thread Maximum	You may adjust the maximum number of threads that can be running at any one time when extracting OneDrive data for folders and files. Increasing the number of threads may improve performance but could increase the chance of receiving a throttling delay from Microsoft. We recommend using the default threading value.

Step 3. Schedule your Discovery

There are four types of schedules you can create (once, daily, weekly, and monthly). These are the same as the schedule types you can create to run reports in the Report Manager. You can create up to five schedules per discovery. Schedules will be grouped alphabetically by type (daily, monthly, once, and weekly) and sorted within the groups chronologically. A schedule can be enabled or disabled at any time. When creating or editing the schedule for a discovery, you can view a calendar showing when the discovery will run based on the schedule. A discovery can also be run manually at any time — unless it is already running. For more information, see How is a Discovery Processed? on page 108 and Manually Running a Discovery on page 109.

> **i** | **NOTE:** All times are stored in local time. They will be adjusted when your local time changes, such as for Daylight Savings Time.

> **i** | **NOTE:** Once you create a schedule for the discovery, you can see the next scheduled run in the main Manage Discoveries pane. For more information, see Viewing your Discoveries on page 109.

See also:

- Run your discovery once
- Run your discovery on a daily interval
- Run your discovery on specified days of the week
- Run your discovery on a specified day of the month
- Run Your Schedule Yearly
- Enabling and Disabling Discovery Schedules

Run your discovery once

This allows you to run your schedule a single time, at a date and time you provide.

To schedule a single run of your discovery

1   On the Schedule page of the Create Discovery wizard, click **Run Once**.

2   Select the start date and time.

3   Click ➕ **Add**.

    Your schedule appears in the bottom pane of the Wizard. If it is not the schedule you want, click Remove.

4   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

5   Click **Finish** to complete the creation of your discovery.

Run your discovery on a daily interval

You can run your schedule every day, or at an interval you choose. For example, if you set the interval to two days, starting on the 22nd day of the month, it will run on the 24th, 26th, 28th and 30th.

The schedule resets at the beginning of the month, so if you have a daily schedule it runs on the first of the month, and then at the set interval. In the above example, after the run on the 30th, the next scheduled runs are the 1st and 3rd of the following month.

To schedule a discovery to run daily

1   On the Schedule page of the Create Discovery wizard, click **Daily**.

2   Select the start date and time.

3   Set the interval for your daily run.

4   Click ✚ **Add**.

    Your schedule appears in the bottom pane of the Wizard. If it is not the schedule you want, click Remove.

5   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

6   Click **Finish** to complete the creation of your discovery.

Run your discovery on specified days of the week

You can set any number of days of the week, and your schedule will run at the set time on those days.

To schedule a discovery to run on selected days

1   On the Schedule page of the Create Discovery wizard, click **Weekly**.

2   Select the start date and time.

3   Select the days on which you want your discovery to run.

4   Click ✚ **Add**.

    Your schedule appears in the bottom pane of the Wizard. If it is not the schedule you want, click Remove.

5   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

6   Click **Finish** to complete the creation of your discovery.

Run your discovery on a specified day of the month

You can select a certain day of the month to run your schedule. This can either be a calendar day, such as the 1st day of the month, or it can be described, such as the last Friday of the month.

To schedule a discovery to run on a specified day of the month

1   On the Schedule page of the Create Discovery wizard, click **Monthly**.

2   Select the start date and time.

3   Select the day of the month on which you want your discovery to run.

    - OR -

Select the weekly interval and day of the month.

4   Click ➕ **Add**.

Your schedule appears in the bottom pane of the Wizard. If it is not the schedule you want, click Remove.

5   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

6   Click **Finish** to complete the creation of your discovery.

Run Your Schedule Yearly

You can select month and day combinations to run your schedule. This feature offers a variety of options for timing a schedule including:

-   a single month and a single day, such as the 31st day December

-   single month and multiple days, such as the 1st and 31st day of January

-   multiple months and a single day, such as the 1st day of every other month

-   multiple months and multiple days, such as the 1st and 15th of every third month

To schedule a yearly run

1   On the Schedule page of the Create Discovery wizard, click **Yearly**.

2   Select the start date and time.

3   Select the months and days of the year on which you want your schedule to run.

4   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

Enabling and Disabling Discovery Schedules

You can enable or disable schedules that have been added to a discovery.

To enable or disable a discovery schedule

1   In the Manage Discoveries pane, select the discovery with the schedule to be modified.

2   Click **Edit**.

- OR -

Right-click on the discovery and click **Edit**.

3   Click the **Schedule** tab.

4   Select the schedule or schedules to enable, disable, or toggle.

5   Right-click **Enable** to set selected schedules to a status of Enabled.

- OR -

Right-click **Disable** to set selected schedules to a status of Disabled.

- OR -

Right-click **Toggle** (or click **Enable/Disable**) to change selected schedules from their current status to the opposite status.

6   Optionally, click **View Calendar** to display a calendar showing the schedule visually.

Step 4: Review the summary

For the NTFS Discovery, the final stage of creating a discovery is to review the summary. The summary displays the settings for each of the available options in the discovery. If any of the settings need to be changed, you can go back to the previous pages and make adjustments.

To review the summary of the discovery

1   After creating and configuring the discovery, click **Next**.

2   Review the settings for each of the available options.

3   If there is a schedule, review when the discovery is scheduled run.

4   Click **Finish** to accept the settings and finalize the discovery.

    - OR -

    Click Back to make changes to the settings.

Managing Discoveries

- How is a Discovery Processed?
- Viewing your Discoveries
- Working with Discoveries and Tasks

How is a Discovery Processed?

When a discovery runs, the process is as follows:

- The server dispatches the resolution task to an available node in the cluster.

- The resolution task reduces the scope to the smallest possible number of tasks. As the tasks are resolved, they go into the queue and are assigned sequentially to the nodes within the assigned cluster. Nodes can continue to accept tasks until they reach their maximum number of concurrent tasks.

  > **NOTE:** A task may be rejected if the target is already being collected by another processing discovery.

- If no nodes are available to process a task, the discovery is in the queue in a Pending state. When a node becomes available, the discovery will begin processing.

- If the task is a collection task, data is written to the database, the last collected time for the parent object is updated, and if there are changes to any data, the timestamp is updated. The collector examines existing data and compares it to the new data, and only sends changed information to the database. This keeps network traffic to a minimum.

  > **NOTE:** You can see the number of items that have been changed, and therefore updated, in the statistics for the task. For more information, see Viewing Statistics on page 116.

- Once all tasks have been processed, the discovery is finished.

  > **NOTE:** For information on configuring nodes, see Creating Your First Cluster and Node on page 27 and Improving the Performance of Your Discoveries on page 33.

See also:

- Types of Tasks
- Manually Running a Discovery

Types of Tasks

Each discovery is broken down into tasks for processing. You will see these task types as you manage your discoveries. There are several types of tasks.

Table 36. Types of Tasks in Configuration Manager

Task Type	Description
Resolution	Examines your scope and reduces it to the smallest number of targets. For an Microsfot® SQL discovery, a target is a SQL Server®.

Table 36. Types of Tasks in Configuration Manager

Task Type	Description
Discovery	Displays as the name of the target. A discovery task collects the data from the target and updates the Enterprise Reporter database.
Group Membership	When the collection requires information from other targets, to improve performance the work is performed at the end of the collection. This can prevent the same data from being repeatedly collected.

Manually Running a Discovery

You can manually run a discovery at any time — unless it is already running. Before you manually run a discovery, you should check the Next Run column in the Manage Discoveries main view to see if the next run is scheduled to start soon. A Discovery that has no enabled schedules will not show a Next Run time.

To manually run a discovery

1 On the Manage Discoveries pane, select the discovery.

2 Click **Run**.

See also:

Viewing your Discoveries

Ideally, you can configure and schedule your discoveries, and the Configuration Manager will run discovery jobs to provide fresh data to the Report Manager at the intervals you have specified. However, many things may arise that require your attention. For example, a SQL Server® may go offline, or another user may cancel your discovery. Or, you may receive inquiries from users about the freshness of the data, or whether a particular object is included in the data on which they are reporting. You can use the Manage Discoveries pane to view your discoveries and answer questions about data collection.

The main Manage Discoveries pane is a listing of all the discoveries created on any console connected to your Enterprise Reporter server. A snapshot of the current state of each discovery is shown. For each discovery, you can see:

- A red indicator if there were errors during the last run. This indicates that all of the requested data was not collected. If all errors have been suppressed for this discovery, the red indicator will also be suppressed.

- **Next Run**: The next run of the discovery. This is calculated based on the discovery's schedule.

  - If the discovery is currently running, this column contains a Processing link, which you can click to view the tasks for the discovery. For more information, see Viewing the Tasks for a Processing Discovery on page 113.

  - If the discovery is waiting to run, this column contains a Pending link that you can click to the view the tasks for the discovery.

  - If the discovery is not running, the date and time of the next run is shown.

  - If the discovery is not scheduled, and is not processing or pending, this column is empty.

- **Last Run Status**: The results of the last run of the discovery - whether it finished, finished but was unable to collect all of the data, or was canceled.

- **Last Run**: When the discovery last completed. This will be empty until the first successful completion of the discovery.

- **Last Run Time**: How long the most recent run of the discovery took to complete. This time is measured from the time the discovery is submitted to the server until the last task finishes running.

    > **i** | **NOTE:** The total time for the discovery to complete may actually be longer than the sum of its tasks, as the cluster may have been processing other discoveries at the same time, or there may have been issues on the node or server.

- **Average Run Time**: The average of the Last Run Time of the 10 most recent runs of the discovery. Only runs with a Last Run Status of *Finished* or *Finished with failures* are included.

- **Error Suppressions**: Indicates when errors have been suppressed for a discovery.

By default, discoveries are grouped by cluster and sorted by discovery name. The current grouping is displayed just above the column headers. An arrow in the column header indicates the current sort.

The Manage Discoveries pane can be customized to display columns in an alternate sequence or to display discoveries with revised grouping, sorting, and filtering. Changes made to the Manage Discoveries pane are automatically saved and used for subsequent logins of the same account.

To sort a column

- Click the up arrow ▲ to sort in descending alphabetical order (Z to A).

- Click the down arrow ▼ to sort in ascending alphabetical order (A to Z).

To filter a column

- Click the filter icon ▼ in a column heading to display a list of every entry in that column and select a value on which to filter.
    - (Blanks) displays rows with a blank entry in that column.
    - (Non-blanks) displays rows with an entry.
    - (Custom) opens a Custom AutoFilter where you can create a custom filter.

    - OR -

    If an empty row is displayed under the column headers, enter text in each column of this row to perform an exact-match search to narrow the results.

To clear a filter

- With a filter selected for a particular column, click the filter icon ▼ in that column heading and choose (All).

To move a column

- Click on the header of the column to be moved, drag it between two other columns (until arrows display indicating the new location for the column), and drop it into place.

To change the grouping

1. Optionally, drag any attributes from the grouping area above the column header down into the column headers to create a new column and stop grouping by that attribute.
   For example, drag the default grouping of Cluster down to create a Cluster column and stop grouping by Cluster.

2. Optionally, drag the header of any column into the grouping area above the column headers. This will remove the column and group by the attribute. For example, drag the Type column up to group by discovery type.

3. Optionally, drag other column headers into the grouping next to an existing group to add sub-groups.

4   Optionally, click on any column header to change the sorting within the grouping.
    For example, if grouping by discovery type, click the Last Run Time header to sort by last run time within each type of discovery.

5   Optionally, to clear the grouping, right-click in the grouping area just above the column headers and select **Clear Grouping**.

See also:

- Navigating the Manage Discoveries Pane
- Viewing the History of a Discovery
- What Does the Discovery Status Indicate?
- Viewing the Tasks for a Finished Discovery
- Viewing the Tasks for a Processing Discovery
- What Does the Task Status Indicate?
- Why is My View Empty?
- Viewing Errors and Error Suppressions
- Suppressing Discovery Errors
- Viewing Statistics
- Viewing a Cluster's Queue

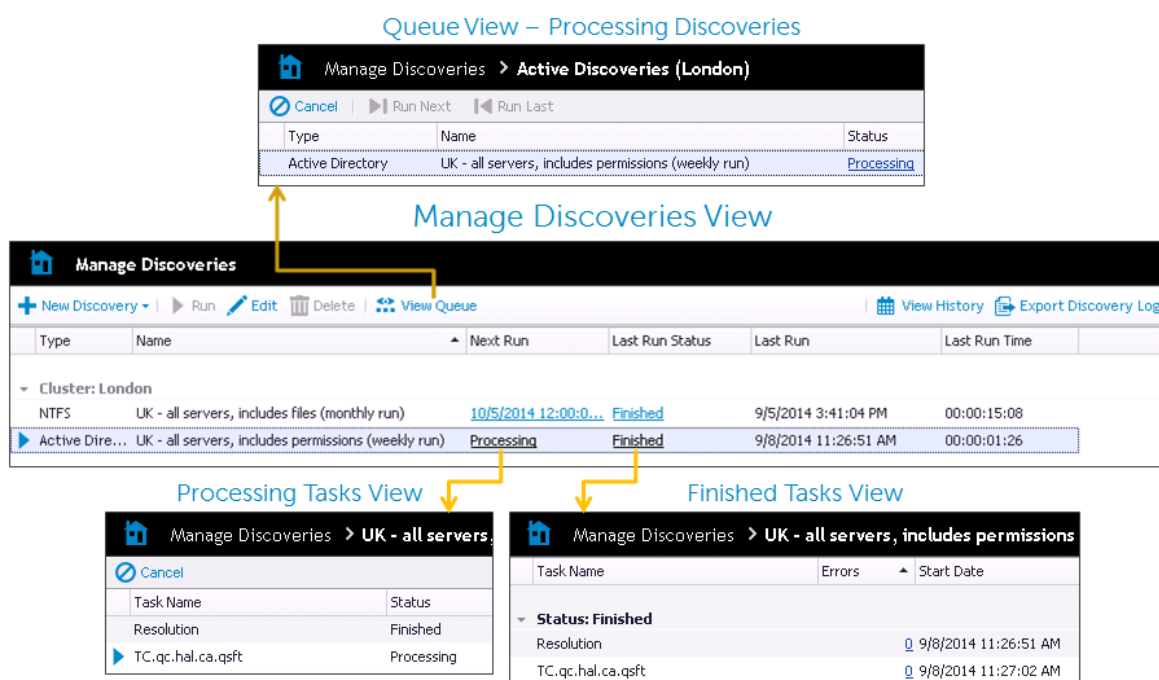# Navigating the Manage Discoveries Pane

The Manage Discoveries pane is actually a series of views designed to allow you to see the progress of your discoveries and to aid in troubleshooting any issues that arise during a collection. In the main Manage Discoveries view, you can create, edit, and run discoveries.

The Manage Discoveries pane uses a drill down approach:

- You can drill down into the last completed run of the discovery and see the details of each task that was processed. You can hover over any task to display the discovery options that were used during collection.

- You can drill down into a currently processing discovery, and see:

    ▪ The tasks for the discovery being processed. This lets you see all the work currently being processed or waiting to be processed.

    ▪ The activity currently taking place. For each task, you can see details of what is currently being processed by the assigned node.

    ▪ The current state of the tasks in your discovery, along with the errors and statistics for each task.

As you drill down into your processing or completed discoveries, a breadcrumb bar helps you understand the context of the information on your screen. At any time, you can click the Manage Discoveries breadcrumb to return to the main Manage Discoveries view. Figure 5 outlines how you can access the details of your completed or processing tasks.

Figure 5. A summary of the views in the Manage Discoveries pane

To display the main Manage Discoveries pane

• From within the Manage Discoveries pane, click the **Manage Discoveries** breadcrumb.

  - OR -

  From another pane, on the Navigation pane, click the **Manage Discoveries** button.

Viewing the History of a Discovery

Enterprise Reporter keeps a history of the last ten runs of each discovery. This can be an aid in troubleshooting — for example, you can see where an error first started appearing in a discovery.

To view the history of a discovery

1   On the Manage Discoveries page, select a discovery.

2   Click **View History**.

3   To view the tasks for a discovery, click the status link.

What Does the Discovery Status Indicate?

The status of a discovery is largely dependent on the status of the tasks within the discovery. For more information, see What Does the Task Status Indicate? on page 114. There are several locations you can view the status of a discovery on the Manage Discoveries page:

• The Last Run Status column tells you the status of the last completed run of the discovery.

• If a discovery is currently being processed, the Next Run column indicates its current status.

• After the first run, when a discovery is running you have access to both the status of the last run and the current run.

- You can see the status of any completed discovery run when you view the history of a discovery.

The following table outlines the available discovery statuses.

Table 37. Discovery Statuses in Configuration Manager

Discovery Status	Description
Pending	The discovery is in the queue, but has not yet started processing. This can be as a result of manual or scheduled run of the discovery.
Processing	Once the first task of a discovery is processed (the resolution task), the discovery goes into a processing state, and remains there until the discovery is canceled or all tasks complete.
Canceled	Indicates that the discovery has been successfully canceled.
Finished	Indicates that all tasks have finished with no errors or errors that did occur were suppressed. The Error Suppressions column will indicate when errors are suppressed for the discovery.
Finished with Failures	Indicates that all tasks have finished, but at least one task failed.
Failed	Indicates that all tasks have finished, but they all failed.

Viewing the Tasks for a Finished Discovery

Each discovery that has finished at least once has a list of tasks associated with it. Viewing this list is particularly useful if there were errors in your discovery. You can pinpoint exactly which targets are causing errors. The completed tasks are grouped by status, so you can easily see the outcome for each task. For each task, you can also see:

- The type of task. For more information, see Types of Tasks on page 108.

- The time the task started running.

- The total time it took to process the task.

- Errors and statistics for the task. For more information, see Viewing Errors and Error Suppressions on page 114 and Viewing Statistics on page 116.

- The discovery options that were used during collection.

To view a finished discovery

1 Display the main Manage Discoveries pane.

2 Select the discovery.

3 Click the status link in the Last Result column for the discovery.

The breadcrumb bar indicates the date and time the discovery was submitted for processing.

4 Optionally, hover over any task to display the discovery options that were used during collection.

Viewing the Tasks for a Processing Discovery

If a discovery is currently running, you can view all tasks for that discovery. You can use this view to troubleshoot issues with your discovery, and to cancel a running task. For each task, you can see:

- A red indicator if errors occurred during the processing of the task.

- An arrow beside a task if it is currently processing.

- The status of the task. This indicates whether the task is currently processing, pending, or finished. For more information, see What Does the Task Status Indicate? on page 114.

- The number of errors that occurred during the processing of the task.

- The time the task started processing.

- The time it took to process the task, or the current elapsed time.

- The node to which the task was assigned. A node can be configured to process multiple tasks at once. For more information, see Improving the Performance of Your Discoveries on page 33.

- The current activity on the node. This is a live stream of the work the node is performing, and you can use this to keep track of the progress of the task.

To view tasks for a processing discovery

1   Display the Manage Discovery pane main view.

2   Select the discovery.

3   In the Next Run column for the discovery, click the **Processing** link.

    The breadcrumb bar indicates that you are viewing the processing tasks for the discovery.

    You can see all tasks for the discovery, including tasks that have finished processing.

What Does the Task Status Indicate?

The status of the tasks in your discovery give you information about the how your discovery is being processed. You can view the status of a task in both the processing and finished task views of a discovery. By default, tasks are grouped by status. The following table outlines the statuses you may see:

Table 38. Task Statuses in Configuration Manager

Task Status	Description
Pending	When a discovery is run, each task has to be assigned to a node. If the node is already running its maximum number of concurrent tasks, the task is Pending. It is in the queue, and will be assigned to a node when one becomes available.
Dispatching	When a node becomes available, the task is sent to it for processing. While this is happening, the task status is Dispatching.
Processing	The task is running on the assigned node
Canceling	The server has received your request to cancel the task, and is communicating this to the node.
Canceled	The task has been canceled.
Finished	The task has successfully completed.
Failed	The task has completed, but was unable to collect all of the data you requested. For more information, see Viewing Errors and Statistics for Tasks on page 130.
Rejected	A task is rejected if the same target is already being accessed by a node within the cluster. This can happen when the same discovery is run more than once within a short time.

Why is My View Empty?

Two views in the Configuration Manager show currently processing tasks or discoveries: the cluster's queue, and the active tasks view. If there is nothing currently being processed, these views will be empty.

Viewing Errors and Error Suppressions

When there are errors during the discovery, a red indicator appears beside the discovery and a check mark indicates if any error suppression rules have been defined for that discovery. If you drill into a task view, you can identify the exact task that caused the error. Then, for that task, you can view a list of errors, which explains what

was happening at the time of the error, and the problem that was encountered. If you are searching for a specific error, you can filter the errors to help narrow your search.

> **i** | **NOTE:** If you have used alternate credentials on a discovery, and your resolution task fails, ensure that you have administrator rights on all node host computer in the assigned cluster.

To view the errors and error suppressions for a task

1   Click the discovery's **Last Run Status** to display the tasks for an active or finished discovery.

   Tasks with a red indicator have errors and display the number of errors as a link in the Errors column.

2   Select a task with errors.

3   Click the link in the Errors column.

   - OR -

   In the bottom pane, click **View Errors** in the red status bar.

4   Optionally, expand the Suppressions section at the bottom of the Error window to see all error suppression rules that have been applied.

To sort a column

- Click the up arrow to sort in descending alphabetical order (Z to A).

- Click the down arrow to sort in ascending alphabetical order (A to Z).

To filter a column

- Click the filter icon in a column heading to display a list of every entry in that column and select a value on which to filter.

   - (Blanks) displays rows with a blank entry in that column.

   - (Non-blanks) displays rows with an entry.

   - (Custom) opens a Custom AutoFilter where you can create a custom filter.

   - OR -

   If an empty row is displayed under the column headers, enter text in each column of this row to perform an exact-match search to narrow the results.

To clear a filter

- With a filter selected for a particular column, click the filter icon in that column heading and choose (All).

Suppressing Discovery Errors

Recurring or known errors can be suppressed so that they are not shown for future runs of a discovery thus highlighting any remaining errors for review. Suppression rules can be defined by selecting an existing discovery error and by adding customized filtering based on object name, object type, message, or problem. Errors matching suppression rules will be excluded from the Enterprise Reporter database the next time the discovery runs. Should a suppression rule require changes, the rule can be deleted and recreated at any time.

To add a suppression rule

1   Click the discovery's **Last Run Status** to display the tasks for an active or finished discovery.

   Tasks with a red indicator have errors and display the number of errors as a link in the Errors column.

2   Select a task with errors.

3   Click the link in the Errors column.

- OR -

In the bottom pane, click **View Errors** in the red status bar.

4   Select an error to use as the basis of a new error suppression rule and click **Add Suppression**.

5   Define the error suppression rule by selecting the appropriate filtering options.

6   Click **OK**.

To delete a suppression rule

1   Click the discovery's **Last Run Status** to display the tasks for an active or finished discovery.

Tasks with a red indicator have errors and display the number of errors as a link in the Errors column.

2   Select a task with errors.

3   Click the link in the Errors column.

- OR -

In the bottom pane, click **View Errors** in the red status bar.

4   Expand the Suppressions section at the bottom of the Error window to see all error suppression rules that have been applied.

5   Select the rule to be removed and click **Delete Suppression**.

Viewing Statistics

Statistics provide information about what was collected during the discovery. Statistics are displayed once all information for that discovery is collected. When the discovery is complete, a full listing of the objects found and a summary of the database changes appears.

Only items that have changed since the last time the discovery ran are updated in the Enterprise Reporter database. This keeps your database up to date, while enhancing performance. You can see how many items were updated by examining Total Added, Total Changed and Total Deleted. If you see that objects were discovered but not updated, this means that they have been previously added to the database.

i   **NOTE:** When processing the ACEs in a discovery, only unique ACEs are processed. For example, if 20 ACEs are discovered across all objects, but ten of those ACEs are identical, only one copy of that ACE is actually added to the database.

To view the statistics for a task

1   Display the tasks for an active or finished discovery.

2   In the bottom pane of the view, click the **Statistics** tab.

Viewing a Cluster's Queue

Each cluster maintains a queue for currently running discoveries. You can see the queue for a cluster whenever you have a processing discovery. The queue is a live view, so only discoveries that are currently processing or waiting to be processed are shown. As discoveries finish processing, they disappear from the queue.

For each discovery being processed, the queue shows the current status, and the number of errors encountered to date during the collection. You can drill into a discovery and see the status of individual tasks. You can also cancel a running discovery. For more information, see Canceling a Task or Discovery on page 118.

i   **NOTE:** If a discovery is in a cluster's queue, but not yet processing (in a Pending state) it means that the node is already processing the optimal number of tasks based on CPU load or it has reached its maximum allowable concurrent tasks. If you find that you often have discoveries queued, check that all nodes are set to optimize performance, or consider adding another node. For more information, see Improving the Performance of Your Discoveries on page 33.

To view the queue for a cluster

1   Click **Discovery Management | Manage Discoveries**.

2   Select the discovery or cluster.

3   On the task bar, click **View Queue**.

    If there are no discoveries being processed, the queue is empty.

Working with Discoveries and Tasks

Occasionally, you may need to duplicate a discovery, modify a discovery, or stop a discovery or task from running. If you want to permanently stop the discovery from running, remember to remove the schedule.

See also:

- Duplicating a Discovery
- Modifying a Discovery
- Canceling a Task or Discovery
- Deleting a Discovery

Duplicating a Discovery

You can duplicate all elements of an existing discovery using the **Duplicate** option. Once a copy of the discovery has been created, you can edit the contents to meet your needs. For more information, see Modifying a Discovery on page 117.

To duplicate a discovery

1   Click **Discovery Management | Manage Discoveries**.

2   Select one or more discoveries.

3   Click **Duplicate**.

4   Click **OK** to confirm that you want to duplicate the selected discoveries.

Modifying a Discovery

You can modify all elements of your discovery using the same pages that you used to create it.

ⓘ   **NOTE:** To change the Assigned Cluster of a discovery, the discovery must not be running. Discoveries should be assigned to the cluster located closest to its targets.

To modify a discovery

1   Click **Discovery Management | Manage Discoveries**.

2   Select the discovery.

3   Click **Edit**.

4   Navigate to the desired page of the Edit Discovery dialog box by clicking the desired button. Make your changes.

5   Click **OK**.

If you want to save the changes before moving to another page, click Apply. For example, you must save your scope changes before you can set scope exclusions. Once you click Apply, you cannot cancel the saved changes.

Canceling a Task or Discovery

If you want to stop a discovery or a task in a discovery from running, you can cancel it.

To cancel discoveries

1 Click **Discovery Management | Manage Discoveries**.

2 Select the discovery you want to cancel.

3 Click **View Queue**.

4 Ensure your discovery is selected and click **Cancel**.

5 Click **Yes** to confirm the cancelation.

To cancel tasks

1 Click **Discovery Management | Manage Discoveries**.

2 Select the discovery, and click the **Processing** link for the discovery to access the tasks.

3 Select the tasks to cancel.

4 Click **Cancel**.

5 Click **OK** to confirm the cancelation.

Deleting a Discovery

You can only delete a discovery when there are no tasks currently running. Deleting a discovery does not delete any previously collected data. If the last run status is processing, you must wait until it is finished before you can delete the discovery. Alternatively, you can cancel the running tasks in the discovery, or cancel the discovery, and then delete it. For more information, see Canceling a Task or Discovery on page 118.

> **i** | **NOTE:** If there is more than one Enterprise Reporter administrator in your organization, use caution when deleting a discovery, as it will no longer be available to any user.

To delete a discovery

1 Click **Discovery Management | Manage Discoveries**.

2 Select the finished discovery.

   You can select more than one discovery using Ctrl+click.

   If any selected discovery has tasks running, the Delete button is unavailable.

3 Click **Delete**.

4 Click **Yes** to confirm.

Troubleshooting Issues with Enterprise Reporter

Problems Opening the Consoles

If you have UAC enabled, ensure that you have Administrator permission to open the console at an elevated level.

To open a console, you must be assigned one of the Enterprise Reporter roles.
For more information, see Role Based Security in Enterprise Reporter in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

If you are unable to log into the Configuration Manager, verify the type of groups you have selected during installation and how you are adding accounts to those groups to give them access to Enterprise Reporter.

For more information, see Configuring the Database and Security Groups in the Quest Enterprise Reporter Installation and Deployment Guide in the Technical Documentation.

Troubleshooting Connectivity Issues

Each console maintains connections to the Enterprise Reporter server and to the SQL Server$^®$ database that stores Enterprise Reporter data. A loss of either connection causes problems. Figure 6 outlines the connections between the components and the server and database.

Figure 6. Connections between components and the server and database.

Discovery Console / Reporting Console

See also:

- Restoring a Connection to the Enterprise Reporter Server
- Restoring a Connection to the Enterprise Reporter Database

Restoring a Connection to the Enterprise Reporter Server

There are a number of reasons why a Enterprise Reporter server may be down. When a console loses its connection to the server, it becomes unusable and must be restarted. All users connected to the Enterprise Reporter server are affected. You should check the following connections:

- Ensure that the computer hosting the server is turned on and running properly.
- Ensure that the Enterprise Reporter server service is running. If necessary, restart it using the Services console.
- Ensure that you can reach the host computer over your network.
- Ensure that the server host computer meets the minimum system requirements.

If the server has gone down and been restored since you last logged in, then the next time you connect, you will be informed that the server went down. If you are the main Enterprise Reporter administrator, this allows you to be aware that your server has had issues. Intermittent failures over time may be due to instability in your network, problems on the server's host computer, or your SQL Server$^{®}$ deployment.

Restoring a Connection to the Enterprise Reporter Database

If your server has lost its connection to the database, you can still open a console and connect to the server, but functionality will be limited. You will be unable to create discoveries, run reports or modify your configuration. Ensure that the SQL Server$^{®}$ hosting the Enterprise Reporter database is running, and that the server can access it.

The Report Manager maintains a direct connection to the SQL Server$^{®}$ database, so ensure that the console's computer can also access the SQL Server$^{®}$.

Troubleshooting Connection Timeouts

As Enterprise Reporter processes your requests, constant communication with the database is required. Depending on your network configuration, your Enterprise Reporter deployment, and the power of your SQL Server® host, the solution for timeout issues may vary.

You can fix timeout issues by either increasing the timeout in Enterprise Reporter, or by investigating any systemic or deployment issues. For example, perhaps your SQL Server® where the database is hosted is underpowered, or you have located your Enterprise Reporter server physically distant from your SQL Server®.

There are the following settings for each timeout configuration:

- Connection timeout
  This is the amount of time given to make the initial connection to the database each time communication is needed. This is less likely to need adjustment. Timeouts are more likely due to SQL Server® or network issues than Enterprise Reporter specific problems. However, if you continually are seeing timeout errors, try increasing this setting.

- Command timeout
  This is the amount of time allowed for the database to process requests. If you are getting timeout error messages during data collection, increase this setting.

There are two types of database timeout settings in the Configuration Manager:

- Enterprise Reporter Server timeout
  You can increase the timeout between the Enterprise Reporter server and the database. If a timeout occurs, you will see a warning dialog box, indicating that this has occurred.

- Cluster timeout
  You can increase the timeout between the nodes in the cluster and the database. This is useful when a collection fails due to a timeout, which is indicated by an error on the discovery task. For information on viewing the error, see Viewing Errors and Error Suppressions on page 114. For information on changing the cluster timeout, see Modifying a Cluster on page 29.

To change the database timeout settings for the Enterprise Reporter Server

1 Click **System | Configuration**.

2 Click **Manage database settings used by Enterprise Reporter Server**.

3 To change the time allowed to establish a connection, modify the Connection Timeout.

4 To change the time allowed to process a database command, modify the Command Timeout.

To change the database timeout settings for a specific cluster

1 Click **Manage Discovery Clusters**.

2 Select the cluster to update.

3 Click the **Cluster Details** tab.

4 To change the time allowed to establish a connection, modify the Connection Timeout.

5 To change the time allowed to process a database command, modify the Command Timeout.

Troubleshooting Credential Change Failures

Each credential in the Credential Manager has three parts—an account name, a password and an optional description—and you can change any of them. While most changes should be processed smoothly, occasionally

issues in the network environment may prevent changes from being applies. When a change fails, you need to determine the reason, and then manually make the changes.

If you have to manually change a credential on a node, you should ensure that there are no discoveries running or queued before making the change. Change the credentials using the Services console on the host computer, then restart the service. Verify that the node started in the bottom pane of the Manage Discovery Clusters page. Restart any discoveries you canceled.

Once you have changed credentials:

- If a node fails to start, ensure the credentials have local administrator access on the node host computer, and check that the credentials you provided are valid.

- If a discovery fails, ensure that the new credentials have read access on the targets of any discovery. Check the discovery to see if it using the default node credentials or if credentials are specified. Ensure that the credentials you provided are valid.

- If an Azure or Office 365 discovery fails, (indicating that the tenant cannot be null), ensure that the related tenant application has been reconfigured. For more information, see To reconfigure application for a tenant on page 48.

Auditing Enterprise Reporter Activity

User activity from the Configuration Manager, the Reporter Manager console, and Encryption is stored in the Quest Enterprise Reporter Windows Event Log on the machine for the Enterprise Reporter Server. Using this information, you can audit the following user activity for compliance purposes.

Configuration Manager Event Log

Table 39. Configuration Manager Activity in the event log

Event	Event ID	Event Format
User Login (successful)	2000	User <USERNAME> on machine <MACHINENAME> successfully logged-in to the ConfigurationManager
Create Discovery	2002	User <USERNAME> on machine <MACHINENAME> created discovery '<DISCOVERYNAME>' (<DISCOVERYID>)
Run Discovery (manual)	2003	User <USERNAME> on machine <MACHINENAME> ran discovery '<DISCOVERYNAME>' (<DISCOVERYID>)
Modify Discovery	2004	User <USERNAME> on machine <MACHINENAME> modified discovery '<DISCOVERYNAME>' (<DISCOVERYID>)
Delete Discovery	2005	User <USERNAME> on machine <MACHINENAME> deleted discovery '<DISCOVERYNAME>' (<DISCOVERYID>)
Create Cluster	2006	User <USERNAME> on machine <MACHINENAME> created cluster '<CLUSTERNAME>' (<CLUSTERID>)
Modify Cluster	2007	User <USERNAME> on machine <MACHINENAME> modified cluster '<CLUSTERNAME>' (<CLUSTERID>)
Delete Cluster	2008	User <USERNAME> on machine <MACHINENAME> deleted cluster '<CLUSTERNAME>' (<CLUSTERID>)
Enable Cluster	2009	User <USERNAME> on machine <MACHINENAME> enabled cluster '<CLUSTERNAME>' (<CLUSTERID>)
Disable Cluster	2010	User <USERNAME> on machine <MACHINENAME> disabled cluster '<CLUSTERNAME>' (<CLUSTERID>)

Table 39. Configuration Manager Activity in the event log

Event	Event ID	Event Format
Add Node To Cluster	2011	User <USERNAME> on machine <MACHINENAME> added node <NODENAME> (<NODEID>) to cluster '<CLUSTERNAME>' (<CLUSTERID>)
Remove Node From Cluster	2012	User <USERNAME> on machine <MACHINENAME> removed node <NODENAME> (<NODEID>) from cluster '<CLUSTERNAME>' (<CLUSTERID>)
Enable Node	2013	User <USERNAME> on machine <MACHINENAME> enabled node <NODENAME> (<NODEID>) on cluster '<CLUSTERNAME>' (<CLUSTERID>)
Disable Node	2014	User <USERNAME> on machine <MACHINENAME> disabled node <NODENAME> (<NODEID>) on cluster '<CLUSTERNAME>' (<CLUSTERID>).

You can interrogate the event log directly for object user activity information. As shown in Table 39, each event is logged with the fully qualified user name of the Enterprise Reporter user who performed the event and the machine name where the event occurred.

Report Manager Event Log

Table 40. Report Manager Activity in the event log

Event	Event ID	Event Format
User Login (successful)	2000	User <USERNAME> on machine <MACHINENAME> successfully logged-in to the ReportManager
Create Report	2015	User <USERNAME> on machine <MACHINENAME> created report '<REPORTNAME>' (<REPORTID>)
Run Report (manual)	2016	User <USERNAME> on machine <MACHINENAME> created report '<REPORTNAME>' (<REPORTID>)
Modify Report	2017	User <USERNAME> on machine <MACHINENAME> modified report '<REPORTNAME>' (<REPORTID>)
Delete Report	2018	User <USERNAME> on machine <MACHINENAME> deleted report '<REPORTNAME>' (<REPORTID>)

Encryption Event Log

Table 41. Encryption Activity in the event log

Event	Event ID	Event Format
Security Audit Operation	2020	User <USERNAME> on machine <MACHINENAME> with IP <IPADDRESS> performed <OPERATIONNAME> operation using service contract <SERVICECONTRACTNAME>.
		Operations that may be performed are:
		• IsUserInRole - checks the role of the calling user
		• GETIvKey - returns Enterprise Reporter encryption IV key
		• ImportIvKey - sets a new security key
		• GenerateIvKey - generates and sets a new IV key
		• CheckOrUpgradeKey - check the status of the security key and performs an upgrade if possible
		• ResetAllPasswords - resets all passwords

Resolving Configuration Manager Issues

The Configuration Manager is used to configure your data collection. Collecting data involves your network security, which can occasionally cause problems.

See also:

- Node Issues
- Scope Enumeration Issues
- Data Collection Issues

Node Issues

A node is a computer assigned to a cluster and is responsible for processing discoveries. In this section, you will find troubleshooting for the following issues related to nodes:

- Node Deployment Issues
- Dealing with Unassociated Nodes
- Dealing with Faulted Nodes
- Problems Deleting a Node
- Changing the Node Logging Level
- Changing the Size of the Node Log Files

Node Deployment Issues

If something goes wrong with your node deployment or upgrade, you can manually install and configure the node. When you manually install a node, it appears in the Configuration Manager as an unassociated node.

> **NOTE:** You must have administrative permissions on the node host computer to install the node.

> **NOTE:** If the node host computer is behind a firewall, you must manually install the node.

> **NOTE:** Performance counters must be enabled on the node machine as the performance information is required by the Enterprise Reporter node performance optimization functionality.

Before you begin:

- Note the port being used by the Enterprise Reporter server service.

- For nodes requiring an alternate port, make note of an available port.

  The default port is 7737.

- Ensure the node account is a member of the Reporter_Discovery_Nodes group.

- If applicable, remove the node in the Configuration Manager. For more information, see What does the status of a node or cluster indicate? on page 35.

- If necessary, use the Control Panel to uninstall the node.

To manually install or upgrade a node

1   On a 64-bit operating system, locate the node installer,
    Enterprise ReporterNode 3.2.1.xxxxx x64.msi (where xxxxx is a unique 5-digit code),
    in the Quest\Enterprise Reporter\Server folder in the install location.

    > **NOTE:** If you are installing the node on a remote computer, copy the appropriate Enterprise Reporter Node 3.2.1.xxxxx x64.msi file to that computer.

2   Run the node installer.

3   On the Welcome screen of the Setup Wizard, click **Next**.

4   Accept the license agreement and click **Next**.

5   To install Quest Enterprise Reporter Node in the default folder, click **Next**.

    - OR -

    Click **Change** to choose another folder, then click **Next**.

6   Specify the credentials and port number that will be used by the Enterprise Reporter Node service, then click **Next**.

    This user account must be granted the 'Log on as a service' right and must have permission to access the Reporter server. For more information about the service credentials required by the node, see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17.

7   Enter or browse to the computer that hosts the Enterprise Reporter Server.

    Specify the fully qualified distinguished name of the computer.

8   Specify the port being used by the Enterprise Reporter Server service, then click **Next**.

9   Click **Install**, then click **Finish.**

10 When the node is installed, associate it with a cluster. For more information, see Dealing with Unassociated Nodes on page 126.

Dealing with Unassociated Nodes

An unassociated node is one that has been either manually installed, or left behind from a previous installation of Enterprise Reporter. You can either uninstall the node, or associate the node with a cluster.

To uninstall a node

- Use the Control Panel, and uninstall Quest Enterprise Reporter Node 3.2.1.xxxx (where xxxx is a unique 4-digit code)

To associate a node with a cluster

1   In the Manage Discovery Clusters pane, select the cluster.

2   In the Unassociated Nodes pane, select the node.

3   Click **Associate Node(s) with Selected Cluster**.

4   In the confirmation dialog box, click **Yes**.

5   If necessary, close the Unassociated Nodes pane.

    The node appears associated with the cluster in the Initializing state until it is deployed.

Dealing with Faulted Nodes

A faulted node has lost contact with the server.You should ensure that the host computer is available on the network and that the service is started.

Problems Deleting a Node

If you are deleting a node, you may see an error message indicating that the discovery node installation failed, for example:
"An error occurred copying the discovery node installation program Quest.Reporter.Core.Server.MsiInstaller.exe to \\servername\ADMIN$\Quest.Reporter.Core.Server.MsiInstaller.exe."

This indicates that there was a problem connecting to the host computer. Check your node credentials, ensure that the firewall is not enabled on the host, and ensure that the computer can be reached on the network. Once you have resolved the connection issue, you can attempt to remove the node again.

Changing the Node Logging Level

If you are experiencing difficulty, the support staff may ask you to change the logging level for nodes in a cluster. The default setting for node logging is Warning, which also includes Fatal and Error. You can increase the logging level to Information or Debug to help them troubleshoot your issue.

> **i** | **IMPORTANT:** Use caution when increasing the logging level. We recommend that you do not increase the level permanently, as it may affect node performance. The logging levels are cumulative.
> - Fatal contains fatal errors.
> - Errors contains errors and fatal.
> - Warnings contain warnings, errors, and fatal errors.
> - Information contains information, warnings, errors, and fatal errors.
> - Debug contains debug, information, warnings, errors and fatal errors.
>
> If you increase the node logging level, you might want to increase the node file size temporarily as well. For more information, see Changing the Size of the Node Log Files on page 127.

To change the node logging level

1   In the Manage Discovery Clusters pane, select the cluster.

2   On the Cluster Details tab in the bottom pane, change the node logging level.

> **i** | **NOTE:** The cluster must be enabled to change the node logging level.

3   Click **Apply**.

Changing the Size of the Node Log Files

Enterprise Reporter writes to log files. By default, the node logging level is set to Warning and the cumulative size of the log files is set to 1000 MB. You can manage the log files within each discovery cluster by setting the node logging level and the cumulative size of the log files. For information on changing the node logging level, see Changing the Node Logging Level on page 126.

To change the size of the node log files

1   In the Manage Discovery Clusters pane, select the cluster.

2   On the Cluster Details tab in the bottom pane, change the size of the log files.

The default is 1000 MB and the maximum is 2000 MB. The increments can be set from 100 MB to 2000 MB.

3   Click **Apply**.

Scope Enumeration Issues

When selecting scopes in an on-premises discovery such as NTFS or Active Directory, the credentials in use determine the available scopes. Depending on your discovery, you may be:

- using the default node credentials, in which case your logged in user account determines the available scopes.

- using alternate credentials, in which case these credentials determine the available scopes.

If you are using alternate credentials, and no scopes are available, this may indicate a DNS issue. Ensure the credentials you are using are fully qualified.

Data Collection Issues

You may run into situations where not all of your data is collected, or even no data at all. The first thing you need to determine is what tasks in the discovery are failing. Once you have located the problem tasks, you can use the errors and statistics generated to pinpoint the problem. There are several other things you can examine:

- The errors generated for a task provide a good starting point for troubleshooting. For more information, see Viewing Errors and Statistics for Tasks on page 130.

- If your discovery fails for all tasks, it is possible that your shared data location is the problem. The shared data location may no longer exist, or the node may not have adequate access to it. Check the errors on the discovery task to investigate. For more information, see Viewing Errors and Error Suppressions on page 114. If this is the issue, ensure the shared data location belonging to the cluster exists and is properly permissioned. Shared data locations are not used for Active Directory, Exchange, or NTFS discoveries.

- If your discovery fails for a particular task:

  - The node may not have access to that server. Check your credentials, and change them if necessary. For more information see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17 and Modifying Node Credentials on page 30.

  - If you have used alternate credentials for the discovery, ensure that they are permissioned properly. For more information, see Node Credential and Alternate Credential Details for On-Premises Discoveries on page 17 and Detailed Permissions for Enterprise Reporter Discoveries on page 20.

- WMI and the SQL Server Browser service may be disabled or your credentials are inadequate. WMI and the SQL Server Browser service are used to query for SQL instances that are not broadcasting.

- The task may have been rejected. If a task is rejected, it means that it is currently being collected by another discovery. Due to the way the Enterprise Reporter collects data, collecting from a SQL Server® in more than one discovery can result in data loss. You could only create one discovery for each SQL scope.

- A discovery can fail if it runs at the same time attributes are being extended for that discovery type. Run the discovery again once the extension has been processed.

- If a particular task is timing out, you can increase the amount of time allowed to connect to the database or process a command. A task may fail because the target computer cannot be pinged. The ping setting is available for computer, Exchange, NTFS and registry discoveries. If a target computer cannot be pinged, for example due to network settings or firewall configurations, or if you know that all computers in the discovery are online and available, you can disable the ping. However, if you have added a domain or OU as your scope, and there is a chance that any computer in the container is not available, setting the ping time ensures that no time is spent preparing to collect from these computers. If a computer unexpectedly fails a ping check, try increasing or disabling the ping for the discovery.

- If your reporting users are experiencing unexpected data fluctuation, check your discovery configuration. If the same target (computer) is in more than one discovery, the data available for reporting reflects the last configuration that was run. Enterprise Reporter's recommended practice is to include a target in only one discovery of a given type. If you have accidentally included a target in more than one place, remove it from all but the desired discovery, and then run that discovery. If for some reason you choose to leave the target in more than one discovery, you can mitigate this issue by using the same settings in both discoveries.

- The node may be running an unsupported operating system. Check the system requirements, and if necessary, remove the node from the cluster, then rerun the discovery.

- If your Enterprise Reporter database is hosted on a SQL cluster which has experienced a node failure, this can occasionally result in a task that cannot finish processing. In this case, you may need to recreate the discovery.

- Try running the discovery, and monitoring the Activity column in the Processing Tasks view, or looking at the history of the discovery. This may help you identify the specific activity that is causing performance or data collection issues with the discovery. For more information, see Viewing the Tasks for a Processing Discovery on page 113, and Viewing the History of a Discovery on page 112.

- If your scheduled discovery does not run, there may be system issues that prevent the job from being created based on the schedule. In this case there is no error reported in the Configuration Manager. To address this issue check that your Enterprise Reporter server service is running, validate your license and check the state of any nodes on the system.

- If a discovery disappears, it is likely that another administrator deleted it. You will have to recreate the discovery.

See also:

- Computer Discovery Errors

Computer Discovery Errors

This section outlines some of the errors that may occur during collection and some suggested approaches for resolving them.

Table 42. Computer Discovery Errors

Error Message	Resolution
Retrieving addresses: access is denied - OR - Retrieving cached logon count: Requested registry access is not allowed	The service account, or the alternate credentials used in the Computer discovery does not have administrator access on the remote server. The account requires local Administrator rights on each target. To verify, run regedit.exe as the service account (by holding shift and right-clicking on the executable to run with different credentials) and attempt to connect to the remote server.

Troubleshooting Features in Enterprise Reporter

There are several features in Enterprise Reporter to help you solve problems.

- Exporting Logs from the Configuration Manager
- Viewing Information About Your Enterprise Reporter Configuration
- Viewing Errors and Statistics for Tasks

Exporting Logs from the Configuration Manager

Exporting Discovery Management Logs

Discovery logs can be used to troubleshoot issues with discoveries. Discovery logs are collected from the Reporter server and all of the nodes within a selected cluster, and zipped into files that can be sent to Quest Support to help resolve certain collection problems. The discovery log files are all sent to the Exported Logs folder on the Reporter server. You may have several different .zip files, which may take some time to appear, depending on your configuration:

- A ServerLogs.zip file containing the logs from the server.
- A <Computer Name>_NodeLog.zip file for each node in the cluster.

To export discovery logs

1   On the Manage Discoveries pane, select a discovery.

By selecting a discovery first, the correct cluster for the discovery is automatically chosen.

2   Click the **Manage Logs** button and select **Export Logs**.

3   If necessary, change the selected cluster.

4   Click **Export**.

5   Click the link to locate your zip files.

Zip files are all located in the \ProgramData\Quest\Enterprise_Reporter\Exported_Logs folder.

You can now email your log files to your Quest Support representative.

6   Click **Close**.

---

i | **NOTE:** To help diagnose issues, the ServerLogs.zip contains additional files including the DatabaseMaintenance.log file and a SQLLite file containing Enterprise Reporter system configuration information.

Exporting the Configuration Manager Logs

The Configuration Manager logs can be used to troubleshoot issues with the Configuration Manager service. Information is collected from the Configuration Manager service and is zipped into log files that can be sent to Quest Support to help resolve certain Configuration Manager problems. The log files are sent to the desktop on the Configuration Manager computer and may take some time to appear, depending on your configuration:

To export Configuration Manager logs

1   Click **System | Information**.

2   Under Client Logging Information, click **Export Configuration Manager logs**.

3   Click **Export**.

4   Click the link to locate your zip file.

    You can now email your log files to your Quest Support representative.

5   Click **Close**.

Viewing the Configuration Manager Logs

You can unzip and view the Configuration Manager logs using the Log Viewer.

To view the Configuration Manager logs

1   Click **System | Information**.

2   Under Log Viewer, click **View logs** to open the Log Viewer.

Viewing Information About Your Enterprise Reporter Configuration

Understanding your system setup can be useful when troubleshooting. You can use the System Information page to determine where your console, Reporter server and Reporter database are hosted, what port the server is using to communicate, your software version, and other similar information you may find helpful in resolving issues.

To view system information in the Configuration Manager

•   On the Navigation pane, click **System | Information**.

Viewing Errors and Statistics for Tasks

For each task of a discovery, you can view collection options, errors, and statistics. These may be helpful when you experience failed collections, data that does not match your expectations, or when working on performance issues. To display the collection options for a task, hold the mouse over any column of data for that task.

For more information, see .

Moving the Enterprise Reporter Database

The following summary outlines line how to move the Enterprise Reporter Database from one SQL Server® to another SQL Server®.

* Backup the database on SQL Server®

* Restore the database from a backup to the new SQL Server®

* Connect Enterprise Reporter to the new database location

* Verify that the database has been moved successfully

The following procedures assume that you have the following permissions:

* SQL permissions to access SQL Server® Management Studio on both the current and new SQL Server® to backup and restore the SQL database

* Windows account permissions to copy the database file from one server to another, and to stop and start services on the Enterprise Reporter server

To back up the Enterprise Reporter database (dbReporter) on SQL Server

1   Open Enterprise Reporter and stop all node services in the Configuration Manager.

2   On the Enterprise Reporter Server, stop the Quest Enterprise Reporter Server services (and the Quest Enterprise Reporter Node services if a node was deployed on the Enterprise Reporter server).

3   Start SQL Management Studio and connect to the SQL Server® where the dbReporter database resides.

4   Expand the database node.

5   Right-click on the dbReporter database and select **Tasks | Back up**.

6   In the Back Up Database dialog, note the location and name of the dbReporter database backup and click **OK**.

To restore the Enterprise Reporter database (dbReporter) from a backup to the new SQL Server

1   Copy the .BAK file(s) to the new SQL Server®.

2   Start SQL Management Studio on the new SQL Server®.

3   Right-click on the database's node and select **Restore Database**.

4   Under Source, select **Device**, click the ellipsis (…).

5   Click **Add** and browse to the backup copies.

6   Click **Add**, select the .BAK file, and click **OK** twice.

7   Under Backup sets to restore, select **Restore** beside the backup name.

8   From the Restore database window, click **OK** to restore the database.

To connect Enterprise Reporter to the new Database location

1. On the Enterprise Reporter server, stop the Quest Enterprise Reporter Server (and the Quest Enterprise Reporter Node if a node was deployed on the Enterprise Reporter server).

2. You may have more than one node deployed, so be sure to stop all node services.

3. On the Enterprise Reporter server, select **All Programs | Database Wizard**.

4. Choose **Select/Upgrade Database** and click **Next**.

5. Browse to (or type in the name of) the new SQL Server® and make sure the Enterprise Reporter database name is correct (for example, dbReporter), and click **Next**.

6. Accept the defaults on the Security Groups screen and click **Next**.

7. Click **OK** on the Database Maintenance Wizard popup that appears regarding replication.

8. Click **Finish** to initiate the database configuration.

9. Click **OK** upon completion.

10. Once you are returned to the Main Menu of the Database Wizard, click **Close**.

11. Start the Enterprise Reporter Server service.

12. Start all Enterprise Reporter node services.

To verify that the database has been moved successfully

1. Open Configuration Manager **System | Information** and confirm the database location.

Disaster Recovery

The following backup/restore procedure is the Enterprise Reporter 3.2.1 strategy for disaster recovery. This strategy will help ensure that Enterprise Reporter will be available for use as soon as possible. With regularly scheduled backups of the Enterprise Reporter database, recovery requires re-installing Enterprise Reporter, restoring the data, restoring a registry key, and restarting Enterprise Reporter.

See also:

- Back Up of Enterprise Reporter
- How to Deploy Enterprise Reporter to Another Computer After a Disaster
- Checking the Enterprise Reporter Configuration After a Recovery

Back Up of Enterprise Reporter

One SQL Server® database is created and used by Enterprise Reporter and should be included with the regular SQL Server® backup. This Discovery Management database has a default name of dbReporter.

How to Deploy Enterprise Reporter to Another Computer After a Disaster

If the original computer is unavailable due to disaster or hardware failure, Enterprise Reporter may need to be deployed on a new computer. The Enterprise Reporter database will be required.

To deploy Enterprise Reporter on a new computer

1   Build a recovery computer with the same name as the previous computer on which to install Enterprise Reporter.

  The recovery computer must have the same name as the previous Enterprise Reporter computer so that the agents and nodes that are still active in the environment can continue to use the computer name to contact the Enterprise Reporter services.

2   Recover the Enterprise Reporter database.

  This step may or may not be needed depending on how the initial configuration of Enterprise Reporter was done. For example, if the database was created on a common SQL Server® and the Enterprise Reporter server was on a separate computer, then the database is still available for use. If SQL Server® was installed on the same computer where Enterprise Reporter was installed, and that computer was damaged, then SQL Server® must be installed on the recovery computer and the Enterprise Reporter database must be restored on the recovery computer or on another SQL Server®.

3   If you have recovered the Enterprise Reporter database, you will need to import the encryption key from the backup file using the Enterprise Reporter Encryption Key Manager and the password that was entered when the encryption key was created. For more information, see Appendix: Encryption Key Manager on page 99 and Importing a Key File on page 100.

  > **i** | **NOTE:** If the encryption key backup file is unavailable, you may use the Enterprise Reporter Encryption Key Manager to erase the encrypted passwords used by the Enterprise Reporter Credential Manager. After using this feature, the passwords for all credentials must be re-entered using the Enterprise Reporter Credential Manager.

4   Install Enterprise Reporter on the recovery computer.

5   Start the Database Wizard.

6   Click **Select/Upgrade Existing Database** in the Database Wizard to allow Enterprise Reporter to make all of the necessary connections to the database and click **Next**.

7   Enter the database server and the database name (or accept the default of dbReporter). Select the connection type (Windows or SQL, depending on the initial configuration) and click **Next**.

8   In Configure Security Groups, it is recommended to leave the default setting unless another configuration was selected during the initial install. Click **Next**.

9   Once the database processing has finished, click **Finish**.

  > **i** | **NOTE:** If SQL Server® is installed on the same recovery computer as Enterprise Reporter, review the popup message about upgrading Enterprise Reporter. Select "**I understand and wish to continue**".

Checking the Enterprise Reporter Configuration After a Recovery

Start the Configuration Manager and check the health of the recovered Enterprise Reporter configuration.

To check the Enterprise Reporter Configuration after a recovery

1   Start the Configuration Manager.

2   Select **System | Information**.

  Review and confirm all of the settings that Enterprise Reporter is currently using.

3   Select **Discovery Management | Manage Discoveries** and then click on the **Discovery Nodes** tab.

  All the nodes are displayed.

4   Remove any node with a status of Faulted by selecting the node and clicking **Remove Node**.

5   Select **Yes** on the popup message.

6 Select **Discovery Management** and then click **Manage Discoveries**.

All of the discoveries should be available for use.

> **ℹ** | **NOTE:** If a Shared Data Location is being used with the Clusters, then delete files located in the share as the data in this share will be out of date and will cause errors in the data in reports.

Appendix: PowerShell cmdlets

- What is Microsoft Windows PowerShell?

- What are cmdlets?

- Registering Enterprise Reporter cmdlets

- Adding the snap-ins automatically to new sessions

- Enterprise Reporter cmdlets

- Enabling Enterprise Reporter cmdlets

- Loading the Enterprise Reporter cmdlets

- Using cmdlets to manage clusters and nodes

- Using cmdlets to manage jobs (discoveries)

- Using cmdlets to run reports

What is Microsoft Windows PowerShell?

Microsoft® Windows PowerShell® is a Windows® command-line shell and scripting language designed specifically for system administrators and built on top of the Microsoft .NET Framework. Windows PowerShell can be installed on Windows® XP, Windows Vista®, and Windows Server® 2003, and is included with Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2.

What are cmdlets?

Windows PowerShell® has the concept of cmdlets. A cmdlet is a simple, single-function command that manipulates objects and is designed to be used in combination with other cmdlets.

If you already had Windows PowerShell installed on your computer before you installed Quest Enterprise Reporter, the Enterprise Reporter cmdlets were automatically installed and registered with Windows PowerShell.

The examples in this section show you leverage the cmdlets available in Enterprise Reporter version 3.2.1. These cmdlets allow you to perform many of the functions of Enterprise Reporter in an automation environment. The cmdlets also can be of great use in any environment where a repetitive process involving Enterprise Reporter is needed.

Registering Enterprise Reporter cmdlets

If you installed Windows PowerShell® on your computer after you installed Quest Enterprise Reporter, you must register the cmdlets before you can start using them in Windows PowerShell.

To register the Enterprise Reporter cmdlets

1   Open Windows PowerShell and type the following at the command prompt:

```
Add-PSSnapin Quest.Reporter.Configuration
Add-PSSnapin Quest.Reporter.Reporting
```

2   Type the following at the command prompt to verify that the snap-in was added:

```
Get-PSSnapin
```

All registered snap-ins are listed.

Adding the snap-ins automatically to new sessions

If you do not want to add the Quest™ Enterprise Reporter snap-ins manually each time you start a new Windows PowerShell® session, you can modify the Windows PowerShell profile file so that the snapins are added automatically.

To add the Enterprise Reporter snap-ins automatically when you start a new Windows PowerShell session

* Add the following lines to the Windows PowerShell profile file (profile.ps1) file:

```
Add-PSSnapin Quest.Reporter.Configuration
Add-PSSnapin Quest.Reporter.Reporting
```

* The location of the Windows PowerShell profile file is as follows:

WINDOWS\system32\windowspowershell\v1.0

> **i** | **NOTE:** If you get the error message "...profile.ps1 cannot be loaded because the execution of scripts is disabled" the next time you start a new Windows PowerShell session, type the following at the Windows PowerShell command prompt:
> * Set-ExecutionPolicy RemoteSigned
>
> Then, type the following at the Windows PowerShell command prompt to confirm that the execution policy has been changed:
> * Get-ExecutionPolicy RemoteSigned

Enterprise Reporter cmdlets

This table lists the cmdlets included with Quest™ Enterprise Reporter.

Table 43. Enterprise Reporter cmdlets for use with Windows PowerShell®

Cmdlet	Module
Add-ERADDiscoveryAttribute	Quest.Reporter.Configuration
Add-ERComputerDiscoveryWMIClass	Quest.Reporter.Configuration
Add-ERNode	Quest.Reporter.Configuration
Add-ERSystemCredential	Quest.Reporter.Configuration
Connect-ERConfigurationServer	Quest.Reporter.Configuration
Connect-ERReportingServer	Quest.Reporter.Reporting
Disable-ERCluster	Quest.Reporter.Configuration
Disable-ERNode	Quest.Reporter.Configuration
Enable-ERCluster	Quest.Reporter.Configuration
Enable-ERNode	Quest.Reporter.Configuration
Export-ERReportDefinition	Quest.Reporter.Reporting
Get-ERCluster	Quest.Reporter.Configuration
Get-ERDirectoryConfiguration	Quest.Reporter.Configuration
Get-ERItSearchConfiguration	Quest.Reporter.Configuration
Get-ERJobDefinition	Quest.Reporter.Configuration
Get-ERJobRun	Quest.Reporter.Configuration
Get-ERLastJobRun	Quest.Reporter.Configuration
Get-ERLogs	Quest.Reporter.Configuration
Get-ERNasConfiguration	Quest.Reporter.Configuration
Get-ERNode	Quest.Reporter.Configuration
Get-ERReport	Quest.Reporter.Reporting
Get-ERReportSQL	Quest.Reporter.Reporting
Get-ERSystemCredential	Quest.Reporter.Configuration
Get-ERUnassociatedNode	Quest.Reporter.Configuration
Initialize-Services	Quest.Reporter.Reporting
Invoke-ERReport	Quest.Reporter.Reporting
New-ERCluster	Quest.Reporter.Configuration
New-ERCredentialMapping	Quest.Reporter.Configuration
New-ERJobDefinition	Quest.Reporter.Configuration
Register-ERDirectoryDiscoveryType	Quest.Reporter.Configuration
Register-ERLicense	Quest.Reporter.Configuration
Remove-ERCluster	Quest.Reporter.Configuration
Remove-ERDirectoryConfiguration	Quest.Reporter.Configuration
Remove-ERJobDefinition	Quest.Reporter.Configuration
Remove-ERNode	Quest.Reporter.Configuration
Send-ERItSearchData	Quest.Reporter.Configuration
Set-ERADChangeHistory	Quest.Reporter.Configuration
Set-ERCommonChangeHistory	Quest.Reporter.Configuration
Set-ERComputerChangeHistory	Quest.Reporter.Configuration

Table 43. Enterprise Reporter cmdlets for use with Windows PowerShell®

Cmdlet	Module
Set-ERConfigurationManagerOptIn	Quest.Reporter.Configuration
Set-ERConfigurationServer	Quest.Reporter.Configuration
Set-ERDirectoryConfiguration	Quest.Reporter.Configuration
Set-ERDirectoryDiscoveryTypeConsent	Quest.Reporter.Configuration
Set-ERItSearchConfiguration	Quest.Reporter.Configuration
Set-ERJobDefinitionSchedule	Quest.Reporter.Configuration
Set-ERNodeCluster	Quest.Reporter.Configuration
SetERNTFSChangeHistory	Quest.Reporter.Configuration
Set-ERRegistryChangeHistory	Quest.Reporter.Configuration
Set-ERReportManagerOptIn	Quest.Reporter.Reporting
Set-ERSQLChangeHistory	Quest.Reporter.Configuration
Set-ERSuppress32BitWarning	Quest.Reporter.Configuration
Set-ERSystemCredential	Quest.Reporter.Configuration
Submit-ERJobDefinition	Quest.Reporter.Configuration
Update-ERNode	Quest.Reporter.Configuration

Enabling Enterprise Reporter cmdlets

Before you can use any of the Enterprise Reporter cmdlets, you need to configure the systems on which you will be running the cmdlets.

To enable the Enterprise Reporter comdlets

- Within the folder C:\Windows\System32\WindowsPowerShell\v1.0, create a new file called powershell.exe.config that contains the following lines:

```
<?xml version="1.0"?>
<configuration>
  <startup useLegacyV2RuntimeActivationPolicy="true">
    <supportedRuntime version="v4.0" />
    <supportedRuntime version="v2.0.50727"/>
  </startup>
</configuration>
```

Loading the Enterprise Reporter cmdlets

You can create a profile file that displays the list of Quest™ Enterprise Reporter cmdlets each time you open Windows Powershell®.

To create a profile to display the list of Enterprise Reporter cmdlets

1   On the system where Enterprise Reporter is installed, open Windows PowerShell®.

2   Enter this command to set the registry value so you can run scripts:

```
Set-Executionpolicy RemoteSigned
```

3   Enter this command to create a Windows PowerShell profile file:

```
New-Item -path $profile -type file –force
```

4   Enter this command to open the profile file in Notepad:

```
notepad $profile
```

5   Add the following commands to load the Enterprise Reporter snapins, connect to the server, and display the list of cmdlets:

```
add-pssnapin Quest.Reporter.Configuration

add-pssnapin Quest.Reporter.Reporting

Connect-ERConfigurationServer localhost

Connect-ERReportingServer localhost

get-command -Module quest*
```

6   Save the profile and exit Notepad.

7   Exit Windows PowerShell, and then open it.

    The snapins load and a list of commands display.

Using cmdlets to manage clusters and nodes

The examples in this section deal with the basics of Enterprise Reporter, which are clusters and nodes for the clusters. Without nodes, clusters cannot direct any work to be done and without clusters, nodes cannot do any work.

This section contains the following examples:

- Creating a cluster
- Creating a node
- Disabling a node
- Enabling a node
- Finding a node by name
- Piping cmdlets
- Finding a cluster by name
- Disabling a cluster
- Enabling a cluster

Creating a cluster

The New-ERCluster cmdlet creates a new cluster in Enterprise Reporter Configuration Manager on which discoveries can execute. Nodes are associated with this cluster, which can be installed on systems in remote locations to allow jobs to execute closer to the physical location.

Syntax

```
New-ERCluster [-Name] <String> [[-Description] <String>] [[-SharedDataLocation]
<String>] [[-ConnectionTimeout] <String>] [[-CommandTimeout] <String>]
```

> **i** | **NOTE:** The only parameter that is required is Name. All other variables can be added at a different time. Any
> parameter that contains a space must be enclosed in quotation marks.

Example

In this example, the new cluster named Second Cluster is created.

```
New-ERCluster -Name "Second Cluster" -Description "This is a test description" –
SharedDataLocation  C:\Shared -ConnectionTimeout 100 -CommandTimeout 500
```

Creating a node

The Add-ERNode cmdlet creates a new node that is associated with a cluster in Enterprise Reporter Configuration
Manager. Nodes execute the jobs assigned to them by the cluster. A cluster can have numerous nodes installed
on different systems, which allows for more efficient processing of jobs and returns quicker results.

Syntax

```
Add-ERNode [-Cluster] <String> [-ComputerName] <String> [-Credential] <PSCredential>
[[-MaxJobSlots] <Int32>]
```

> **i** | **NOTE:** Only the -MaxJobSlots parameter is optional; all other parameters must be supplied.

Example

This example involves a three step process. The first step encrypts the password used by the service account
before sending it across the network. The second step combines the encrypted password with the service account
into a new system object containing the credentials for the service account. The third step indicates the cluster,
identifies the server where the node is to be installed, supplies the credentials, and defines how many jobs slots
the node is to use.

```
secpasswd = ConvertTo-SecureString 'pA$$w0d' -AsPlainText -Force
credentials = New-Object System.Management.Automation.PSCredential
('AMER\Administrator', $secpasswd)
Add-ERNode "Second Cluster" "AMERGEN02" $credentials 5
```

Now the cluster named Second Cluster, which was created in the previous example (see Creating a cluster), has a
node associated with it and is ready to run a job.

Disabling a node

There are times when a system may require maintenance or be taken down for some specific reason. During
these times you will want to disable the node installed on that system. Disabling the node allows the cluster to
manage the jobs based on the remaining nodes that are available for work.

> **i** | **NOTE:** Even though you may have disabled a node, any jobs running on the node continue to be processed
> until completed. Only new jobs are not assigned to the node. Therefore, if maintenance is planned for the
> system, consider disabling the node in plenty of time for any job to finish.

Syntax

```
Disable-ERNode [-Node] <Node> [-CancelTasks [<SwitchParameter>]] [-PassThru
[<SwitchParameter>]]
```

```
Disable-ERNode [-Cluster] <String> [-ComputerName] <String> [-CancelTasks
[<SwitchParameter>]] [-PassThru [<SwitchParameter>]]
```

Example 1

In this example, the node associated with the cluster named First Cluster that is installed on the computer named AMERGEN01 is disabled.

```
Disable-ERNode -Cluster "First Cluster" -ComputerName AMERGEN01
```

Example 2

In this example, the node information is stored in the variable $node. The information contained in $node is then used as input to the Disable-ERNode cmdlet.

```
$node = Get-ERNode -Node AMERGEN01.amer.sitraka.com
Disable-ERNode -Node $node
```

Enabling a node

Once any work has been done on the system and you want to bring the node back into use, you need to enable the node so that the cluster knows the node is available and ready for work. Once the node is enabled, the cluster will assign the jobs waiting to be processed.

Syntax

```
Enable-ERNode [-Node] <Node> [-PassThru [<SwitchParameter>]]
```

```
Enable-ERNode [-Cluster] <String> [-ComputerName] <String> [-PassThru
[<SwitchParameter>]]
```

Example 1

In this example, the node associated with the cluster named First Cluster that is installed on the computer named AMERGEN01 is enabled.

```
Enable-ERNode -Cluster "First Cluster" -Passthru -ComputerName AMERGEN01
```

Example 2

In this example, the node information is stored in the variable $node. The information contained in $node is then used as input to the Enable-ERNode cmdlet.

```
$node = Get-ERNode -Node AMERGEN01.amer.sitraka.com
Enable-ERNode -Node $node
```

Finding a node by name

As nodes are an important part of the job processing, knowing about the nodes is vital so that you can ensure they are functioning properly and that the cluster has enough nodes to process jobs. The Get-ERNode cmdlet retrieves information about a node. You In addition, a cluster can be specified to show all nodes associated with the cluster. A computer can also be queried to see if there is a node install on it.

Syntax

```
Get-ERNode [[-Node] <String>] [[-Cluster] <String>]
```

Example 1

This example returns information from all nodes in all clusters.

```
Get-ERNode
```

Example 2

This example returns all nodes on the computer named AMERGEN01.

```
Get-ERNode -Node AMERGEN01
```

Example 3

This example returns all nodes in the cluster named First Cluster.

```
Get-ERNode -Cluster "First Cluster"
```

Piping cmdlets

Cmdlets can pipe the output from one cmdlet into another cmdlet. This feature is useful and powerful when you pipe the Get-ERNode cmdlet into Enable-ERNode and Disable-ERNode cmdlets.

Example 1

This example disables all nodes associated with the cluster named First Cluster. The data for all nodes is retrieved by the Get-ERNode cmdlet, and then piped into the Disable-ERNode cmdlet.

```
Get-ERNode -Cluster "First Cluster" | Disable-ERNode
```

Example 2

This example enables all nodes associated with the cluster named First Cluster. The data for all nodes is retrieved by the Get-ERNode cmdlet, and then piped into the Enable-ERNode cmdlet.

```
Get-ERNode -Cluster "First Cluster" | Enable-ERNode
```

Finding a cluster by name

As with nodes you can retrieve information about clusters. This information includes whether the cluster is enabled or disabled, if the cluster is using a shared data location, and the path of the shared data location.

Syntax

```
Get-ERCluster [[-Cluster] <String>]
```

Example 1

This example returns information on all clusters associated with Enterprise Reporter.

```
Get-ERCluster
```

Example 2

This example returns information on the cluster named Second Cluster. Note that he cluster name is in quotes because Windows PowerShell® requires that any parameter containing spaces must be enclosed with quote marks.

```
Get-ERCluster -Cluster "Second Cluster"
```

Disabling a cluster

As with nodes, there are times when a system may require maintenance or you may want to stop the processing of specific jobs for some specific reason. During these times you will want to disable the cluster. Disable a cluster, effectively stopping new jobs and tasks from starting on any nodes within that cluster.

> **i** | **NOTE:** Even though you may have disabled a cluster, any jobs running on the cluster will continue to execute until completed. Disabling the cluster will disable all the nodes associated with it.

Syntax

```
Disable-ERCluster [-Cluster] <Cluster> [-CancelJobs [<SwitchParameter>]] [-PassThru
[<SwitchParameter>]]
```

```
Disable-ERCluster [-Name] <String> [-CancelJobs [<SwitchParameter>]] [-PassThru
[<SwitchParameter>]]
```

Example 1

In this example, the cluster named First Cluster is disabled.

```
Disable-ERCluster -Name "First Cluster"
```

Example 2

In this example, jobs scheduled to run on the cluster named First Cluster are canceled, and then the cluster is disabled. Jobs currently running will finish even though the cluster is disabled.

```
Disable-ERCluster "First Cluster" -CancelJobs
```

Example 3

In this example, the Get-ERCluster cmdlet first retrieves information about the cluster named First Cluster, and then stores it in the $cluster variable. Next, the cluster with the name stored in the $cluster.name variable is disabled.

If the Get-ERCluster cmdlet is executed without identifying a cluster, it returns the information on all clusters. This technique can be useful when there are a number of clusters since they can be looped through disabling each one.

```
$cluster = Get-ERCluster "First Cluster"
Disable-ERCluster $cluster.name
```

Example 4

In this example, the Get-ERCluster cmdlet retrieves information for the cluster named Second Cluster and pipes it into the Disable-ERCluster cmdlet.

```
Get-ERCluster "Second Cluster" | Disable-ERCluster
```

Enabling a cluster

Once a cluster is enabled, the nodes assigned to the cluster start processing jobs. As with disabling the cluster, enabling the cluster will enable all the nodes associated with it.

Syntax

```
Enable-ERCluster [-Cluster] <Cluster> [-PassThru [<SwitchParameter>]]
[<CommonParameters>]
```

```
Enable-ERCluster [-Name] <String> [-PassThru [<SwitchParameter>]]
[<CommonParameters>]
```

Example 1

This example enables the cluster named First Cluster.

```
Enable-ERCluster "First Cluster"
```

Example 2

This example first stores cluster information in the $cluster variable, and then uses it as input to the Enable-ERCluster cmdlet. Note that the cluster name does not need to be parsed out and passed to the Enable-ERCluster cmdlet.

```
$cluster = Get-ERCluster
Enable-ERCluster $cluster
```

Example 2

This example uses Get-ERCluster to retrieve information for the cluster named First Cluster and pipes it to the Enable-ERCluster cmdlet.

```
Get-ERCluster -Cluster "First Cluster" | Enable-ERCluster
```

Using cmdlets to manage jobs (discoveries)

An important aspect of Quest™ Enterprise Reporter, discoveries return information about the systems in your environment. Discoveries gather data about Active Directory®, computers, SQL Server®, and NTFS permissions for files and folders. When using Enterprise Reporter cmdlets, discoveries are referred to as jobs. The examples in this section demonstrate how to create and run jobs using cmdlets.

This section contains the following examples:

- Getting job information

- Creating a job

- Running a job

- Scheduling a job

- Deleting a job

Getting job information

You probably have configured and run discoveries in the Enterprise Reporter Configuration Manager. These job definitions are useful in understanding how the cmdlets work and provide good examples for you to follow when creating new jobs using cmdlets. The Get-ERJobDefinition cmdlet returns information on the jobs.

Syntax

```
Get-ERJobDefinition [-ClusterId <Nullable`1[Guid]>] [-ClusterName <String>] [-
JobDefinitionName] <String>
```

```
Get-ERJobDefinition [-Unassigned [<SwitchParameter>]]
```

```
Get-ERJobDefinition -JobDefinitionId <Guid>
```

```
Get-ERJobDefinition -JobDefinition <JobDefinition>
```

Example 1

In this example, information about a job identified with the name of Active Directory is returned.

```
Get-ERJobDefinition -JobDefinitionName "Active Directory"
```

Output

```
JobDefinitionId     : ecaae8bb-f0f8-48ee-91ff-da959a937dfa
JobTypeId           : bec47934-cadf-4b94-8ff7-68efadc88165
Name                : Active Directory
Description         :
Configuration       : <DiscoveryResolutionTask
subType="ActiveDirectory"><Scopes><Scope
                      type="ActiveDirectory"><Content><Root class="Domain"
inex="Include"><Ldap>LDAP:

//AMER.amer.sitraka.com/DC=AMER,DC=amer,DC=sitraka,DC=com</Ldap><Dns>

AMER.amer.sitraka.com</Dns><DC>AMERGENDC.AMER.amer.sitraka.com</DC>
                      </Root></Content></Scope></Scopes><Parameters><Parameter

key="GatherSharedEntities"><Value>False</Value></Parameter><Parameter
                      key="CollectDCs"><Value>True</Value></Parameter><Parameter
                      key="CollectAccounts"><Value>True</Value></Parameter><Parameter

key="CollectComputers"><Value>True</Value></Parameter><Parameter
                      key="CollectOUs"><Value>True</Value></Parameter><Parameter
                      key="CollectSites"><Value>True</Value></Parameter><Parameter
                      key="CollectTrusts"><Value>True</Value></Parameter><Parameter

key="CollectPhysicalComputerInfo"><Value>False</Value></Parameter><Parameter

key="FindForeignMember"><Value>False</Value></Parameter><Parameter
                      key="SyncLogon"><Value>False</Value></Parameter><Parameter

key="CollectPermissions"><Value>False</Value></Parameter><Parameter

key="CollectRemoteServices"><Value>False</Value></Parameter><Parameter
                      key="CollectPhotosForUsers"><Value>False</Value></Parameter>
                      </Parameters></DiscoveryResolutionTask>
AssignedClusterId   : af0e6da4-10b9-4b20-bf5a-f19f0d71b589
CredentialId        :
AssignedClusterName : First Cluster
IsTombstoned        : False
Schedule            : <Trigger Type="RunOnceTrigger" StartDateTime="2014-07-
14T13:50:00Z"
                      Expression="2014-07-14 13:50:00Z@0 50 13 14 7 ? 2014" />
NextRun             :
```

Example 2

In this example, information about all the jobs (the * wildcard is used in -JobDefinitionName) located on the cluster named Second Cluster is returned.

```
Get-ERJobDefinition -ClusterName "Second Cluster" -JobDefinitionName *
```

Output

```
JobDefinitionId     : 7a8d758b-15e8-4f84-a30e-1c8545eed4f5
JobTypeId           : bec47934-cadf-4b94-8ff7-68efadc88165
Name                : Computer
Description         :
```

```
Configuration          : <DiscoveryResolutionTask subType="Computer"><Scopes><Scope
type="Computer">

                         <Content><Rootclass="Domain"
inex="Include"><Ldap>LDAP://AMER.amer.sitraka.com

/DC=RPTCH,DC=dev,DC=hal,DC=ca,DC=qsft</Ldap><Dns>AMER.amer.sitraka.com
                         </Dns></Root></Content></Scope></Scopes><Parameters><Parameter

key="GatherSharedEntities"><Value>False</Value></Parameter><Parameter
                         key="CollectAccounts"><Value>True</Value></Parameter><Parameter

key="CollectEventLogConfig"><Value>True</Value></Parameter><Parameter
                         key="CollectPolicies"><Value>True</Value></Parameter><Parameter
                         key="CollectPrinters"><Value>True</Value></Parameter><Parameter
                         key="CollectServices"><Value>True</Value></Parameter><Parameter
                          key="CollectShares"><Value>True</Value></Parameter><Parameter
                          key="CollectVolumes"><Value>True</Value></Parameter><Parameter

key="CollectExtendedWMIEntities"><Value>True</Value></Parameter><Parameter
                         key="PingTimeout"><Value>0</Value></Parameter></Parameters>
                         </DiscoveryResolutionTask>
AssignedClusterId      : 390b0bd1-6488-4ace-b2cc-616925b55c06
CredentialId           :
AssignedClusterName    : Second Cluster
IsTombstoned           : False
Schedule               : <Trigger Type="RunOnceTrigger" StartDateTime="2014-07-
15T13:58:00Z"

                         Expression="2014-07-15 13:58:00Z@0 58 13 15 7 ? 2014" />
NextRun                :
```

As you can see in these examples, there is a lot of information contained in the job definition. The largest and seemingly most complicated part is the configuration, which contains all of the information about the jobs that you created using the Discovery Wizard. For more information about the configuration, see Creating a job.

Creating a job

Using cmdlets to create a new job requires planning as there is a lot information contained in a job definition. You would use cmdlets to automate a process, such as cloning a current job or creating a new job in an environment with limited resources.

Syntax

```
New-ERJobDefinition [-Name] <String> [-JobType] <String> [-Configuration] <String>
[[-ClusterId] <Guid>] [[-CredentialId] <Guid>] [[-Description] <String>] [[-
Schedule] <String>] [[-NextRun] <Nullable`1[DateTime]>] [-PassThru
[<SwitchParameter>]]
```

• [-JobType] values are: ActiveDirectory, Computer, Exchange, MS SQL, NTFS, and Registry.

• [-Configuration] is the XML representation of the job or discovery configuration. See any of the job examples in the Getting job information section, which discussed the Get-ERJobDefinition cmdlet. This is the best way to get a configuration to use in creating a job manually.

The configuration can be contained in an XML file, making it easier to navigate. Using Notepad, copy the configuration from an existing job and paste it into a file.   It is recommended that you create default files with the extension .XML and have the format as seen below which makes it easier to go thru and make the setting you want to make.

Do not use any special program that provides an XML format. Any additional data in the file will not be interpreted correctly and can cause errors with the job creation. Use Notepad to avoid unseen character formatting.

Example 1

As you can see in this example, each section of the XML file has an opening and a closing statement. When you are working with a copy of the configuration from the Get-ERJobDefinition cmdlet, pay attention to spaces, text, slashes, and other characters, as missing or extra characters will cause an issue with the job.

```
<DiscoveryResolutionTask subType="NTFS">
  <Scopes>
    <Scope type="NTFS">
      <Content>
        <Root class="Computer" inex="Include">

<Ldap>LDAP://AMER.amer.sitraka.com/DC=AMER,DC=amer,DC=sitraka,DC=com</Ldap>
          <Dns>\\AMERGEN02.AMER.amer.sitraka.com</Dns>
        </Root>
      </Content>
    </Scope>
  </Scopes>
  <Parameters>
    <Parameter key="CollectPublicShares">
      <Value>True</Value>
    </Parameter>
    <Parameter key="TreatSharesAsShares">
      <Value>True</Value>
    </Parameter>
    <Parameter key="CollectFolderPermissions">
      <Value>True</Value>
    </Parameter>
    <Parameter key="CollectChildACLDifferences">
      <Value>True</Value>
    </Parameter>
    <Parameter key="FolderDepth">
      <Value>-1</Value>
    </Parameter>
    <Parameter key="GatherSharedEntities">
      <Value>False</Value>
    </Parameter>
    <Parameter key="GlobalInclude"/>
    <Parameter key="GlobalExclude" />
    <Parameter key="CollectFiles">
      <Value>True</Value>
    </Parameter>
    <Parameter key="CollectFilePermissions">
      <Value>True</Value>
    </Parameter>
    <Parameter key="CollectOnlyExplicitFilePermissions">
      <Value>False</Value>
    </Parameter>
    <Parameter key="IncludeFile">
      <Value>*.*</Value>
    </Parameter>
    <Parameter key="PingTimeout">
      <Value>0</Value>
    </Parameter>
  </Parameters>
</DiscoveryResolutionTask>
```

Example 2

In this example, the contents of the XML configuration file is used with the New-ERJobDefinition cmdlet. The values for the -Configuration parameter are enclosed in a single quote mark (').

The parameters in this example are formatted to make it easier for you to read. When running the cmdlet, the parameters cannot contain any carriage returns in the command line. If you want to use this example, you must first paste it into NotePad and remove the carriage returns.

```
New-ERJobDefinition -Name "NTFS All" -JobType NTFS -Configuration
'<DiscoveryResolutionTask  subType="NTFS"><Scopes><Scope type="NTFS"><Content><Root
class="Domain"
inex="Include"><Ldap>LDAP://AMER.amer.sitraka.com/DC=AMER,DC=amer,DC=sitraka,DC=com
</Ldap><Dns>AMER.amer.sitraka.com</Dns></Root></Content></Scope></Scopes><Parameter
s>
<Parameter key="CollectPublicShares"><Value>True</Value></Parameter><Parameter
key="TreatSharesAsShares"><Value>True</Value></Parameter><
Parameter key="CollectFolderPermissions"><Value>True</Value></Parameter><Parameter
key="CollectChildACLDifferences"><Value>True</Value></Parameter><Parameter
key="FolderDepth"><Value>-1</Value></Parameter><Parameter
key="GatherSharedEntities"><Value>False</Value></Parameter><Parameter
key="GlobalInclude"/><Parameter
key="GlobalExclude" /><Parameter
key="CollectFiles"><Value>True</Value></Parameter><Parameter
key="CollectFilePermissions"><Value>True</Value></Parameter><Parameter
key="CollectOnlyExplicitFilePermissions"><Value>False</Value></Parameter><Parameter

key="IncludeFile"><Value>*.*</Value></Parameter><Parameter
key="PingTimeout"><Value>0</Value></Parameter></Parameters></DiscoveryResolutionTas
k>'
-ClusterId af0e6da4-10b9-4b20-bf5a-f19f0d71b589 -Schedule '<Trigger
Type="RunDailyTrigger"
StartDateTime="2014-07-17T04:00:00Z" Expression="2014-07-1704:00:00Z@0 0 4 1/2 * ?"
EveryNDay="2" />'
```

Example 3

In this example the configuration is in an XML file, which allows changes to be made to the configuration when using the cmdlets. The XML files are just simple files that can be edited using Notepad and do not require any special formatting. The configuration of any current job is in the XML format and can be used as a template.

```
$filepath = "c:\configuration.xml"
$configuration =[string]::join([environment]::newline, (get-content -path
$filepath))
$foundCluster = Get-ERCluster 'First Cluster'
New-ERJobDefinition -Name NTFS10 -JobType NTFS -Configuration $configuration –
ClusterId $foundCluster.ClusterId
```

Example 4

In this example, you want to clone a current job. The important item to note is that the -Name parameter needs to be changed to a unique value. The first cmdlet Get-ERJobDefinition gets the data on the job you wish to clone. In the New-ERJobDefinition cmdlet you use the data in the configuration of the cloned job by using $JobDefinition.Configuration to supply the needed configuration for the new job.

```
$JobDefinition = Get-ERJobDefinition "NTFS All"
$foundCluster = Get-ERCluster 'First Cluster'
New-ERJobDefinition -Name NTFS2 -JobType NTFS -Configuration
$JobDefinition.Configuration -ClusterId $foundCluster.ClusterId
```

Running a job

Now that you have a job or two you need to run them to retrieve data from your environment by sending a job to the Enterprise Reporter server for immediate execution. Depending on what is processing within the server, the job may be queued to run at the next available time. This is different than scheduling a job which is discussed later.

Syntax

```
Submit-ERJobDefinition [-JobDefinitionId] <Guid>

Submit-ERJobDefinition [-JobDefinition] <JobDefinition>
```

Example 1

In this example, the job or discovery identified by the JobDefinitionId ecaae8bb-f0f8-48ee-91ff-da959a937dfa is submitted for immediate processing. If the job starts, True is returned.

```
Submit-ERJobDefinition -JobDefinitionId ecaae8bb-f0f8-48ee-91ff-da959a937dfa
```

Example 2

In this example, the information about the job definition retrieved by the Get-ERJobDefinition cmdlet is piped to the Submit-ERJobDefinition cmdlet, so the job starts immediately. If the job starts, True is returned.

```
Get-ERJobDefinition "Active Directory" | Submit-ERJobDefinition
```

Scheduling a job

You may want to change the start time for a scheduled job because it conflicts with another job.

i | **NOTE:** The schedule must be in CRON format. The time are UTC or GMT (time zone Z (zulu)).

Syntax

```
Set-ERJobDefinitionSchedule [-JobDefinitionId] <Guid> [[-Schedule] <String>]

Set-ERJobDefinitionSchedule [-JobDefinition] <JobDefinition> [[-Schedule] <String>]

Set-ERJobDefinitionSchedule [-JobDefinitionName] <String> [[-Schedule] <String>]
```

Example 1

This is an example of a Run Once job set to start at a specific date and time. First, the job is placed into the $discovery variable using the Get-ERJobDefinition cmdlet. Second, the Set-ERJobDefinitionSchedule cmdlet is executed with a different time and date for the job. Note the single quote that encloses the time.

```
$discovery = get-ERJobDefinition "Active Directory"
Set-ERJobDefinitionSchedule $discovery.JobDefinitionId '<Trigger
Type="RunOnceTrigger"  StartDateTime="2014-07-15T16:50:00Z" Expression="2014-07-15
16:50:00Z@0 50 16 15 7 ? 2014" />'
```

Example 2

This is an example of a Run Daily job set to start at a specific date and time, and to run every day.

```
Set-ERJobDefinitionSchedule -Schedule '<Trigger Type="RunDailyTrigger"
StartDateTime="2014-05-27T14:40:00Z" Expression="2014-05-27 14:40:00Z@0 40 14 1/1 *
?" EveryNDay="1" />'
```

Example 3

This is an example of a Run Daily job set to start at a specific date and time, and to run every 4[th] day.

```
Set-ERJobDefinitionSchedule -Schedule '<Trigger Type="RunDailyTrigger"
StartDateTime="2014-05-28T16:04:00Z" Expression="2014-05-28 16:04:00Z@0 4 16 1/4 *
?" EveryNDay="4" />'
```

Example 4

This is an example of a Run Weekly job set to start at a specific date and time, and to run on Monday, Wednesday, and Friday.

```
Set-ERJobDefinitionSchedule -Schedule '<Trigger Type="RunWeeklyTrigger"
StartDateTime="2014-05-30T16:01:00Z" Expression="2014-05-30 16:01:00Z@0 1 16 ? *
MON,WED,FRI" RunOnWeekDays="21" />'
```

Example 5

This is an example of a Run Monthly job set to start at a specific date and time and to run on the 1st Wednesday of
the month.

```
Set-ERJobDefinitionSchedule -Schedule '<Trigger Type="RunMonthlyTrigger"
StartDateTime="2014-06-03T16:02:00Z" Expression="2014-06-03 16:02:00Z@0 2 16 3 * ?"
MonthlyScheduleType="NthDayOfTheMonth" WeekType="First" WeekDay="Wednesday"
DayOfTheMonth="3" />'
```

Example 6

This is an example of a Run Monthly job set to start at a specific date and time on the 3rd Wednesday of the month.

```
Set-ERJobDefinitionSchedule -Schedule '<Trigger Type="RunMonthlyTrigger"
StartDateTime="2014-06-18T16:02:00Z" Expression="2014-06-18 16:02:00Z@0 2 16 ? *
WED#3" MonthlyScheduleType="NthDayOfTheNthWeek" WeekType="Third"
WeekDay="Wednesday" DayOfTheMonth="28" />'
```

Deleting a job

Syntax

```
Remove-ERJobDefinition [-JobDefinitionId] <Guid>

Remove-ERJobDefinition [-JobDefinition] <JobDefinition>

Remove-ERJobDefinition [-JobDefinitionName] <String>
```

Example 1

In this example, all job definitions that start with the letter A are piped into the Remove-ERJobDefinition cmdlet for
deletion.

```
Get-ERJobDefinition a* | Remove-ERjobdefinition
```

Example 2

```
$job = Get-ERJobDefinition "AD Discovery"
Remove-ERJobDefinition $job
```

First, the job data is placed into the $job variable. Second, the Remove-ERJobDefinition is run with $job as the
input.

Using cmdlets to run reports

Now that you have collected data for your environment, you will want to produce reports with the data. Normally
this is done using the Report Manager, but there are a number of cmdlets available to provide reports. The
examples in this section demonstrate how to run reports using cmdlets.

This section contains the following examples:

- Connecting to the server
- Getting report information
- Exporting a report definition

- Generating a report with data

Connecting to the server

Before you can use the report cmdlets, you must establish a connection to the Enterprise Reporter server. If you do not have your profile set up (see Loading the Enterprise Reporter cmdlets), you will need to establish a connection.

Syntax

```
Connect-ERReportingServer [-Server] <String> [[-Port] <Int32>]
```

Example

In this example, a connection is established to the AMERGEN01 server through port 7738.

```
Connect-ERReportingServer -Server AMERGEN01 -Port 7738
```

Getting report information

Report Manager has a report library that is broken into a logical folder structure with reports in each folder. The functional cmdlets require the ID and path associated with each report and not just the report name. To get the report ID and path, use the Get-ERReport cmdlet. You can use wildcard expressions with this cmdlet. The wildcard search is performed on the full report path including folder names and report name.

Syntax

```
Get-ERReport [[-ReportName] <String>]
```

Example 1

This example returns all information on the Domain Accounts report.

```
Get-ERReport "*\Active Directory\Domain Accounts*"
```

Output

Id	CategoryPath	ReportName
58355095-75ee-4d21-9c11-1ccc15cbe3e3	Report Library\Active Directory	Domain Accounts

Example 2

This example returns all the reports that begin with Domain.

```
Get-ERReport "*\Active Directory\Domain*"
```

Output

Id	CategoryPath	ReportName
58355095-75ee-4d21-9c11-1ccc15cbe3e3	Report Library\Active Directory	Domain Accounts
e4067d44-b100-46e9-94da-8c3348584b50	Report Library\Active Directory	Domain Computer Information
59143e6b-cf55-4aa7-b70c-69c2b5b61659	Report Library\Active Directory	Domain Controller Information

```
584229c5-dc71-4107-a0e5-38a2a6d4b8de   Report Library\Active Directory   Domain Groups
with Members

e6cf50aa-6078-4be8-aa92-3b0850264855   Report Library\Active Directory   Domain Groups
without Members

c49948e9-0337-47bd-bd48-e5cbf861391f   Report Library\Active Directory   Domain Groups

e4067d44-b100-46e9-94da-8c3348584b50   Report Library\Active Directory   Domain Hidden
Computers

237e4518-b843-4af7-911d-2b0ed62f1001   Report Library\Active Directory   Domain Sites

92e8568c-850d-416a-8b2b-3b01ad0f1b43   Report Library\Active Directory   Domain
Summary

b41d8e07-d0f4-46da-a2ae-bdd96992a12b   Report Library\Active Directory   Domain Trusts

73d3a6b2-95d7-4690-9aad-3fdf39dac04e   Report Library\Active Directory   Domain Users
with Recent Logons

73d3a6b2-95d7-4690-9aad-3fdf39dac04e   Report Library\Active Directory   Domain Users
without Recent Logons

95ccdb73-b815-47d2-af98-bc359201b6cd   Report Library\Active Directory   Domain Users
```

Exporting a report definition

Reports in Report Manager can be modified and configured. The modified report can be exported to a designated location for disaster recovery or to share with others. The complete report information obtained from the Get-ERReport cmdlet is required when performing the export. For more information, see Getting report information.

Syntax

```
Export-ERReportDefinition [-Report] <Report> [-Destination] <String>
```

Example 1

In this example, the Get-ERReport cmdlet places the report information for a report in the Report Library into the $rpt variable. Next, the report definition is exported to the c:\ drive. The exported definition contains the report name and the report Id: Domain Users_95ccdb73-b815-47d2-af98-bc359201b6cd.xrd.

```
$rpt = Get-ERReport "Report Library\Active Directory\Domain Users"
Export-ERReportDefinition $rpt -Destination c:\
```

Example 2

In this example, the Get-ERReport cmdlet places the report information for a report located in My Reports into the $rpt variable. Next, the report definition is exported to the c:\ drive. The exported definition contains the report name and the report Id: Domain Computer Information_e4067d44-b100-46e9-94da-8c3348584b50.xrd.

```
$rpt = Get-ERReport "My Reports\Domain Computer Information"
Export-ERReportDefinition $rpt -Destination c:\
```

Generating a report with data

You can generate reports in either PDF or CSV format, with the CSV file as a comma delimited file. Reports can be useful as recordkeeping or for disaster recovery. The complete report information obtained from the Get-ERReport cmdlet is required when performing the export. For more information, see Getting report information.

Syntax

```
Invoke-ERReport [-Report] <Report> [-Type] <String> [-Destination] <String>
```

Example 1

In this example, the report information is placed into the $rpt variable with the cmdlet Get-ERReport. Next, the report in PDF format is written to the c:\ drive.

```
$rpt = Get-ERReport "Report Library\Active Directory\Domain Users"
Invoke-ERReport -Report $rpt -Type PDF -Destination c:\
```

Example 2

In this example, the information for the report located in My Reports is placed into the $rpt variable with the cmdlet Get-ERReport. Next, the report in PDF format is written to the c:\ drive.

```
$rpt = Get-ERReport "My Reports\Domain Computer Information"
Invoke-ERReport -Report $rpt -Type PDF -Destination c:\
```

Appendix: Encryption Key Manager

- Starting the Encryption Key Manager
- Generating a Key File
- Importing a Key File
- Exporting a Key File
- Resetting Credentials

Enterprise Reporter makes use of FIPS 140-2 compliant encryption to secure user credentials and includes an encryption key management tool. The Enterprise Reporter Encryption Key Manager can be started from the Windows Start menu. This tool allows you to perform the following tasks related to the Enterprise Reporter encryption key.

- Generating an encryption key
- Importing an encryption key from a backup file
- Exporting an encryption key to a backup file
- Resetting Enterprise Reporter user credentials

Starting the Encryption Key Manager

To start the Encryption Key Manager from the Windows Start menu

1   Click Programs | Quest | Enterprise Reporter | Encryption Key Manager

Generating a Key File

The Encryption Key Manager can be used to generate a new encryption key. If Enterprise Reporter contains credentials with passwords, selecting this option will force the decryption and re-encryption of all Enterprise Reporter user credentials. If the decryption of the existing passwords fails, the procedure is unsuccessful, errors are returned, and no key file is generated. If the decryption and re-encryption is successful, the procedure continues and the new encryption key is written to the secure Windows Credential Manager (not to be confused with the Enterprise Reporter Credential Manager). The user is prompted to export the new key to a backup file.

To generate a key file

1   Stop all Enterprise Reporter nodes.

2   Start the Enterprise Reporter Encryption Key Manager.

3   Click the **Generate Key** button.

4   Read and accept the warning.

5   Click **OK** to generate a new key file.

6   Click **OK** to continue to export the key file to a backup file.

Importing a Key File

The Encryption Key Manager can be used to import an encryption key from an Enterprise Reporter backup file. This option requires the user-supplied password that was used to create the backup file. If Enterprise Reporter contains credentials with passwords, this procedure will decrypt and re-encrypt all of them and store the imported encryption key in the secure Windows Credential Manager (not to be confused with the Enterprise Reporter Credential Manager).

To import a key file

1   Stop all Enterprise Reporter nodes.

2   Start the Enterprise Reporter Encryption Key Manager.

3   Click the **Import Key** button.

4   Enter the fully qualified filename of the backup file.

   - OR -

   Click the ellipsis to navigate to the Import Location of the backup file.

5   Enter the user-supplied password for the backup file.

6   Read and accept the warning.

7   Click **OK** to import the backup file.

8   Click **OK** to accept the successful import notification.

Exporting a Key File

The Encryption Key Manager can be used to export the current encryption key to a backup file encrypted with a user-supplied password.

> **i** | **IMPORTANT:** It is very important to remember this password as it is non-recoverable.

To export a key file

1   Click the **Export Key** button.

2   Enter a fully qualified filename as an export location for the backup file.

   - OR -

   Click the ellipsis to navigate to an Export Location for the backup file.

3   Enter a password with a minimum of 10 characters.

4   Enter the password again to confirm the password.

5   Click **OK** to create the backup file.

6   Click **OK** to accept the successful backup notification.

Resetting Credentials

The Encryption Key Manager can be used to erase all encrypted passwords used by the Enterprise Reporter Credential Manager when it is impossible to restore a valid encryption key. After using this feature, passwords for all credentials must be re-entered using the Enterprise Reporter Credential Manager. The Credential Manager will display a red key icon next to each account that requires a password.

> ! | **CAUTION: Use the Reset Credentials option only when it is impossible to restore a valid encryption key.**

To reset credentials

1 Click the **Reset Credentials** button.

2 Read and accept the warning.

3 Click **OK** to erase all passwords for credentials stored in the Enterprise Reporter Credential Manager.

4 Click **OK** to confirm that you wish to erase all passwords.

5 Click **OK** to accept the successful reset notification.

•

Appendix: Log Viewer

The Enterprise Reporter Log Viewer can be started from the Configuration Manager, the Report Manager, or the Windows Start menu. The Enterprise Reporter Database Log Viewer allows you to perform the following tasks on the log files generated by Enterprise Reporter.

- Browsing for log files

- Unzipping log files

- Drag and drop to open log files

- Correlating events from multiple log files and displaying them chronologically

- Searching within log files for specific events or errors

- Limiting the events displayed using filters

Starting the Enterprise Reporter Log Viewer

To start the Enterprise Reporter Log Viewer in the Configuration Manager

1   Click **System | Information | Log Viewer | View Logs**

To start the Enterprise Reporter Log Viewer in the Report Manager

1   Click **System Information | log viewer |View Logs**

To start the Enterprise Reporter Log Viewer from the Windows Start menu

1   Click **Programs | Quest | Enterprise Reporter | Log Viewer**

Finding and Opening Log Files

The first time the Log Viewer is started, it displays the contents of the default Enterprise Reporter log folder including date, time, and file size information.

\programdata\quest\enterprise_reporter

To navigate to different folder containing log files, click the ellipsis to the right of the log folder path. The files in the selected folder will be listed in the Log Viewer file browser. During the time the Log Viewer is open, the contents of the folder may be updated by clicking the Refresh icon next to the log folder path.

To open log files using the Log Viewer browser

1   Double-click the log file containing entries to be viewed.

- OR -

Select the log file containing entries to be viewed and click the **Open** icon.

- OR -

Drag the log file containing entries to be viewed onto the main log entry viewing panel.

- OR -

Drag a log file from Windows File Explorer onto the main log entry viewing panel.

To unzip log files in the Log Folder panel

1   Select the log file containing entries to be unzipped and click the **Unzip** icon.

- OR -

Right-click the log file containing entries to be unzipped and select the **Unzip** icon.

To clear log files in the Imported Log Files panel

1   Select the log file containing the entries to be cleared from the main log entry viewing panel and click the **Clear** icon.

- OR -

Right-click the log file containing entries to be cleared from the main log entry viewing panel and select the **Clear** option.

Viewing and Searching Log File Entries

The files listed in the Imported Log Files panel have their contents correlated and displayed in the main log entry viewing panel sorted by date and time.

To search for specific text within the log entries

1   Enter the text to locate in the **Find** text box.

2   Press **Enter** to locate the first occurrence of the text within the log entries being viewed.

The matching log entry will be highlighted.

3   Optionally, click the **Find Next** icon to find the next occurrence of the text within the log entries being viewed.

4   Optionally, click the **Find Previous** icon to find the previous occurrence of the text within the log entries being viewed.

To search for errors by error text within the log entries

1   Enter the error text to locate in the **Find** text box.

2   Optionally, click the **Next Error** icon to find the next occurrence of the text within the ERROR log entries being viewed.

3   Optionally, click the **Previous Error** icon to find the previous occurrence of the text within the ERROR log entries being viewed.

To browse for errors within the log entries

1 Click on the log entry from which you wish to browse.

2 Optionally, right-click the log entry and select **Next Error** to browse to the next error within the log entries being viewed.

3 Optionally, right-click the log entry and select **Previous Error** to browse to the next error within the log entries being viewed.

To view the details of a log entry

1 Double-click a log entry.

   - OR -

   Right-click a log entry and select **View Details**.

   - OR -

   Select a log entry and click the **View Details** icon above the main log entry viewing panel.

To clear all event log entries

1 Click the **Clear All** button above the main log entry viewing panel.

2 Accept the warning message to continue with removing all of the log entries.

Filtering Log File Entries

Once a listing of log file entries is displayed, the Filters option can be used to limit the entries by dates and other properties. For more information, see Viewing and Searching Log File Entries on page 158. Setting a Start Date will display entries with a time stamp that occurs on or after that date. Setting an End Date will display entries with a time stamp that occurs on or before that date. Selecting options for each property will display entries matching those options.

To filter log file entries

1 Once you are viewing a listing of log file entries, click the **Filters** button.

2 Optionally enter a Start Date.

3 Optionally, enter an End Date.

4 Optionally, select at least one option per property.

5 Optionally, click the reset button to clear all filters and start again.

6 Click **Apply** to display the log file entries that match the filters.

Index

cluster, 29, 121
server, 121
tombstoning, 56

U

uac, 119
unassociated node, 126
undeploying status, 36

V

view history, 112
viewing
errors, 114
queue, 116
statistics, 116

W

wmi, 128

About us

Quest provides software solutions for the rapidly-changing world of enterprise IT. We help simplify the challenges caused by data explosion, cloud expansion, hybrid datacenters, security threats, and regulatory requirements. We are a global provider to 130,000 companies across 100 countries, including 95% of the Fortune 500 and 90% of the Global 1000. Since 1987, we have built a portfolio of solutions that now includes database management, data protection, identity and access management, Microsoft platform management, and unified endpoint management. With Quest, organizations spend less time on IT administration and more time on business innovation. For more information, visit www.quest.com.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at https://support.quest.com.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.