Symantec Data Insight Administrator's Guide

4.5

Symantec™

Symantec Data Insight 4.5 Administrator's Guide

4.5

Documentation version: 4.5 Rev 0

Legal Notice

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization

- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information

- Upgrade assurance that delivers software upgrades

- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis

- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level

- Hardware information

- Available memory, disk space, and NIC information

- Operating system

- Version and patch level

- Network topology

- Router, gateway, and IP address information

- Problem description:

  - Error messages and log files

  - Troubleshooting that was performed before contacting Symantec

  - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization

- Product registration updates, such as address or name changes

- General product information (features, language availability, local dealers)

- Latest information about product updates and upgrades

- Information about upgrade assurance and support contracts

- Information about the Symantec Buying Programs

- Advice about Symantec's technical support options

- Nontechnical presales questions

- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Contents

Getting started with Symantec Data Insight administration

This chapter includes the following topics:

- About Symantec Data Insight administration
- Handling changes in account password

About Symantec Data Insight administration

You administer the Symantec Data Insight system through the Management Console. The console has components for system administration, viewing data access information, configuring policies and alerts, and generating reports, which are accessible from the tabs located on the header panel. Navigate to the **Settings** tab on the console to carry out the various Data Insight administration tasks.

The Console is automatically installed with the Management Server. You access the Console through a Web browser that has a network connection to the Management Server. By default, the Management Server runs on HTTPS port 443. To access it, in the Web browser's address field, type `https://ms-host/`.

The Server Administrator user can see and access all parts of the administration console. Other users can see only the parts to which their roles grant them access. The user account under which you are currently logged on appears at the footer of the Management Console screen.

Operation icons on the Management Console

Table 1-1 shows the operation icons that are located on the console screen:

Table 1-1   Operation icons on the Management Console

Icon	Description
	Go up one level in the navigation control.
	Filter filers, Web applications, shares, site collections, users, and groups.
	The filter options depend on the current level of hierarchy.
	Clears the filter.
	The settings icon is used in assigning custodians.
	Screen refresh. Symantec recommends using this refresh button instead of your browser's **Refresh** or **Reload** button.
	Email the data on the current screen to one or more recipients. If the current screens data cannot be sent as an email, the icon is unavailable.
	Exports all data on a panel on the current screen to a `.csv` file.
Export	Exports all data on the current screen to a `.csv` file.
	Submits request to the Enterprise Vault server to archive the selected folders.
	The action selector icon displays a menu with the following two options: ■ Archive files using Enterprise Vault. ■ Submit request to invoke a custom action on selected paths.
	Submit request to invoke a custom action on selected paths.
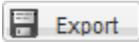
Data Insight administration tasks

Table 1-2 summarizes the tasks to be performed to set up a new Data Insight installation:

Table 1-2   Data Insight administration tasks

Action	Description
Configure SMTP server settings.	See "Configuring SMTP server settings " on page 35.
Setup notification policies.	See "Configuring email notifications " on page 272.
Configure directory service domain.	See "Adding a directory service domain to Data Insight" on page 65.
Configure data retention settings.	See "Configuring data retention settings" on page 48.
Configure Exclude Rules.	See "Adding exclude rules to Data Insight " on page 41.
Install license.	See "Managing Data Insight licenses " on page 59.
Configure Data Insight nodes either individually or configure multiple nodes by applying node templates.	See "About node templates" on page 221.
If monitoring events for NetApp file servers, configure Fpolicy service on collectors.	See "Preparing Symantec Data Insight for Fpolicy " on page 84.
If monitoring events for a clustered NetApp file server, install the DataInsightFPolicyCmod service on the collectors.	See "Preparing Data Insight for FPolicy in NetApp Cluster-Mode" on page 101.
If monitoring events for EMC Celerra file servers, configure Celerra service on collectors.	See "About EMC Common Event Enabler (CEE)" on page 104.
If monitoring events for EMC Isilon file servers, configure EMC Common Event Enabler (CEE) on the collectors.	See "Preparing Symantec Data Insight to receive event notifications from an EMC Isilon cluster" on page 116.
If monitoring events for Windows file servers , upload agent packages to collectors.	See "Configuring Windows File Servers " on page 119.
If monitoring events for a generic device, use web APIs to collect access event information.	See "About configuring a generic device" on page 129.
If monitoring events for SharePoint servers, install the Data Insight Web service on the SharePoint server.	See "Installing the Data Insight web service for SharePoint" on page 173.

Table 1-2   Data Insight administration tasks *(continued)*

Action	Description
Configure file servers.	See "Adding filers" on page 134.
Configure the SharePoint Web applications.	See "Adding web applications" on page 174.
Configure advanced analytics	See "Configuring advanced analytics " on page 53.
Configure archiving settings	See "About configuring archive options for Enterprise Vault" on page 228.
Configure permission remediation settings	See "Managing and configuring permission remediation" on page 225.
Configure workflows to remediate DLP incidents, review permissions on data resources, and confirm ownership of files and folders.	Do the following: 1 Install the Self-Service Portal. For details, see the *Symantec Data Insight Installation Guide*. 2 Configure workflow templates. See "About workflow templates" on page 242. 3 Create remediation workflows and submit them for further action on the Self-Service Portal. See "Creating a workflow using a template" on page 248.

Supported file servers and platforms

Table 1-3 lists the Network Attached Storage devices and SharePoint platforms that Data Insight supports.

Table 1-3   Supported file servers and platforms

Device	Version
NetApp ONTAP	From v7.3.5 to v8.1.x ONTAP 8.0.x and ONTAP 8.1.x are supported in 7 mode only. ONTAP 8.2.x is supported in 7-mode and Cluster-Mode.

Table 1-3   Supported file servers and platforms *(continued)*

Device	Version
EMC Celerra	5.6.45 or higher, VNX
EMC Isilon	OneFS version 7.1 or higher
Windows File Server	Windows Server 2003 or 2003 R2, 32 bit, and 64 bit
	Windows Server 2008, or 2008 R2, 32 bit and 64 bit
	Windows Server 2012, or 2012 R2 64 bit
Veritas File System (VxFS) server	6.0.1 or higher, configured in standalone or clustered mode using Veritas Cluster Server (VCS)
	Note: For VCS support, Clustered File System (CFS) is not supported.
Microsoft SharePoint	Microsoft Office SharePoint Server 2007
	Microsoft SharePoint 2010
	Microsoft SharePoint 2013
Symantec Data Loss Prevention (DLP)	Versions 12.0.1 and 12.5
Symantec Enterprise Vault	Versions 10.0.4 and 11.0

Note the following:

- Symantec strongly recommends that you upgrade your NetApp filer to the latest available firmware. Symantec recommends ONTAP 7.3.5 or higher.

- For all supported versions of 7-mode NetApp filers, Data Insight supports CIFS protocol over NTFS. The supported NetApp volume/qtree styles are NTFS and Mixed. For supported versions of Cluster-Mode NetApp filers, Data Insight supports only CIFS protocol over NTFS. NFS protocol is not supported.

- For all supported versions of EMC Celerra/VNX and EMC Isilon, Data Insight supports only CIFS protocol over NTFS. NFS protocol is not supported. Data Insight supports the latest Common Event Enabler (CEE), version 6.3.1. Data Insight still supports the older version of CEE and VEE, but Symantec recommends that you move to the latest EMC Common Event Enabler, which you can download from the EMC website.

- To use the Self-Service Portal to remediate DLP incidents, ensure that Symantec Data Loss Prevention (DLP) version 12.5 or higher is installed. Data Insight

uses the DLP Smart Response Rules to remediate incidents, which are
introduced in DLP version 12.5.

Handling changes in account password

You use various account credentials at the time of configuring the Data Insight
system. Accounts are used when configuring the following:

- Fpolicy Service
- Celerra Service
- Filers
- SharePoint Web applications
- Scanner
- Active Directory

Perform the following steps to ensure that updates to account passwords are
synchronized with the passwords used in Data Insight:

To handle changes in account password

1 Determine the places where the account is being used.

2 Log in to the Data Insight console and edit the saved credential password.

   For example, navigate to **Settings** > **Saved Credentials**, and edit the credential
   to update the password.

3 If the password of an account, which is used for Fpolicy service or Celerra
   service configuration, has changed, you must reconfigure the services as well.

   Navigate to Server details page for the corresponding nodes acting as
   Collectors, click the **Reconfigure Fpolicy** or **Reconfigure Celerra** sections
   on the page.

See "Managing saved credentials " on page 46.

Using Symantec Data Insight dashboards

This chapter includes the following topics:

About the Data Insight dashboard

Use the Data Insight dashboard to view a high-level summary of the configured storage devices from the perspective of space utilization, activity, and permissions assignment. You can view a separate dashboard for the configured storage devices as also for the shares and site collections.

The dashboard also provides an insight into the open shares on a file server, which can help you do an entitlement review. You can use the open shares information to ensure that sensitive content in your environment is not open to excessive number of users.

The Data Insight dashboard combines the statistics from various access, permissions, utilization, and consumption reports and provides a snapshot of relevant statistics on one page.

You can do the following with the information displayed on the dashboard:

- Sort and filter the data on the dashboard to view only information you are interested in.

- Email the data to desired recipients.

- Export the data on the dashboard to a `.csv` file.
  When you export the data on the **Devices** sub-tab, data pertaining to all the devices along with the configured shares and site collections on these devices is exported. When you export the data on the **Shares/Site Collections** sub-tab, only the data about the shares and site collections is exported.

See "Configuring advanced analytics " on page 53.

Viewing device summary dashboard

Use the **Dashboard** > **Devices** view to get a high level summary of the configured storage devices.

The device summary dashboard gives a snapshot view of all the configured storage devices that are being monitored by Data Insight. You can view the following details about a device on the dashboard:

- The host name or IP address of the configured device.

- The type of device.

- The storage capacity of the configured device.

- The size of the used space on the storage device.

- Total number of open shares on a filer.
  Open shares are shares that are accessible to global access groups, like Everyone, Domain Users, and Authenticated Users on the network. Such open shares may contain sensitive data. You can use the Device summary dashboard to monitor accesses on the open shares.
  See "About open shares" on page 55.

- Size of the data residing on open shares.

Note: Data Insight calculates the **Capacity** and **Used space** information at the volume level, across all volumes on the filer. While, the **Open Shares Data Size** is the sum of the sizes of all the files on all the open shares exported from that filer. For example, on a given volume, you have two shares *shareA* and *shareB*, such that *shareA* is exposed from *folderA* in the qtree and *shareB* is exposed from *folderB* that falls in the subtree of *folderA*. Then a file that is under *folderB* would appear in *shareA* and *shareB* in the Data Insight Management Console (assuming both shares are monitored by Data Insight) . Hence, some duplication of file-sizes could take place, giving rise to a larger value in the **Open Shares Data Size** value.

- Total number of files in open shares.

- A list of Data Loss Prevention (DLP) policies that have been violated.

- The size on disk. This size can be different from the logical size of the share or site collection. If a path is archived by Enterprise Vault, its on-disk size is much lower than its logical size.

- The primary attribute configured for the purpose of advanced analytics.

Note: Used space and capacity values are not shown for SharePoint Web applications, EMC Celerra filers, and Veritas File System (VxFS) filers.

To filter the dashboard data

1   Click the filter icon to the top right of the **Devices** page.

2   On the filter criteria pop-up, click **Add new clause**.

3   Select the filter criteria from the **Device** drop-down, and select an operator to build the filter query.

4   Enter a value for the condition, and click **Add Clause**.

5   Click **Apply** to view the filtered results on the **Devices** dashboard.

6   The expressions that you create to filter the data are saved for future use. Click the filter icon again to reuse the existing filter criteria.

Note: You can apply multiple filters to narrow your search of the dashboard data. All filters are applied simultaneously and all conditions must be satisfied to get the results.

Viewing summary of shares and site collections

The **Dashboard** > **Shares/Site Collections** tab displays the high-level summary of configured shares and site collections from the context of activity, security, and storage.

The data that is displayed on the dashboard changes depending on the context you choose. For example, if you choose **Activity** from the **Select view** drop-down, the dashboard displays the number of accesses on the share or site collection, the number of active users, and the date when the last activity took place. If you choose the **Security** view, the dashboard displays the number of sensitive files on the shares or site collections, and indicates whether a share is open or not.

The **Detailed** view displays the exhaustive data which combines the information from all perspectives. You can view the following information on the **Detailed** view:

- Name of the shares or site collections that are monitored by Data Insight.

- Path of the shares or site collections on the filers.

- Whether a share is open or not. The nature of activity on the share or site collection.

- The type of share - whether collaborative, accesed by multiple users, accessed by a single user.

- The total space occupied by the share or site collection.

- The size on disk. This size can be different from the logical size of the share or site collection. If a path is archived by Enterprise Vault, its on-disk size is much lower than its logical size.

- The total files on the share or site collection.

- The total number of sensitive files on the share or site collection. This information is fetched from Data Loss Prevention.

- The total folders on the share or site collection.

- The total number of access events on the share or site collection.

- The total number of active users on the share or site collection.

- The date of the last activity on the share or site collection.

- A list of Data Loss Prevention (DLP) policies that have been violated.

- The primary attribute configured for the purpose of advanced analytics. For example, if you configure Department as the primary attribute, the dashboard displays columns for the number and list of departments accessing the share or site collection.

- A warning symbol (!) that indicates some problem with computing the dashboard data. For example, a red exclamation mark indicates that the indexer node associated with the devices is down and cannot compute data. An orange exclamation mark indicates that the scan on that share or site collection has not taken place or the share or site collection is not added to the Data Insight configuration.

  A tootip indicating the cause for the error is displayed when you mouse over the exclamation mark.

You can create a custom view by selecting the columns that you want to display on the dashboard.

To filter the dashboard data

1   Click the filter icon to the top right of the **Shares/Site Collections** page.

2   On the filter criteria pop-up, click **Add new clause**.

3   Select the filter criteria from the **Device** drop-down, and select **equals** or **contains**.

4   Enter a value for the condition, and click **Add Clause**.

5   Click **Apply** to view the filtered results on the dashboard

6   The expressions that you create to filter the data are saved for future use. Click the filter icon again to reuse the existing filter criteria.

Viewing the system health overview

The **System Overview** dashboard available under **Settings** > **Health** > **System Overview** provides a quick-reference snapshot of the the health of your entire environment. Use the **System Overview** dashboard to view an inventory count and state of all monitored storage devices, Data Insight servers, and configured directory services. The **System Overview** dashboard also displays the state of scanning of all configured objects in your environment.

The various reports on the dashboard are laid out in a sequence. You must scroll down the content pane to view each report.

Use the **System Overview** dashboard to do the following:

- Review the count of inventory objects and their states.

- Review the count of the troublesome or potentially troublesome inventory.

- Identify the issues that need immediate attention.

The dashboard contains the following widgets:

Data Insight servers

The Servers pie-chart displays the graphical representation of the total number of Data Insight servers that are in the Faulted, At Risk, Unknown, and Healthy state.

The widget also provides an overview of the following:

- Inventory count of the number of Collectors, Indexers, and Windows File Server agents in your environment.
  Click a server type to view the filtered list for the selected server type on the **Settings** > **Data Insight Servers** page.
  Click **More details** to view the entire list of configured Data Insight servers.
- Notifications and warnings that indicate the reasons for the health status of the configured nodes.
  For more information about the factors that affect the health of a node, See "Managing Data Insight product servers" on page 195.

Scanning

The Scanning graph displays a color bar chart representing the number of scans during the last seven days from the current date.

The color bar chart represents the different states of the scans- Failed [Red]; Successful [Green]; and Partially successful [Yellow].

To the right of the chart, you can also view the following data:

- An overview of the consolidated scan status of all configured shares or site collections based on the latest scans that have taken place on these shares and site collections.
  Click on a consolidated status to view its details on the **Settings** > **Scanning** > **Overview** tab.
- Notifications pertaining to the state of scanning and related alerts and warnings.

See "Viewing the scanning overview" on page 28.

See "Viewing the scan status of storage devices" on page 31.

See "Viewing the scan history of storage devices" on page 32.

See "Viewing in-progress scans" on page 200.

Devices	The Devices pie-chart provides a snapshot view of the number of storage devices in the Faulted, At Risk, Unknown, and Healthy states.
	To the right of the chart, you can view the following data:
	■ Inventory of filers and Web applications that are being monitored by Data Insight. Click a device type to view the filtered list of configured devices of that type on the **Settings** > **Data Insight Filers** or the **Settings** > **SharePoint Web applications** page.
	■ Alerts associated with the storage devices that indicate the status, error conditions, and warnings that can help when troubleshooting issues with the device.
Directory services	The directory services widget provides an inventory count of the configured directory services. The widget also displays information and alert notifications associated with the directory services.

See

Viewing the scanning overview

The **Settings** > **Scanning** > **Overview** tab of the **Scanning** dashboard displays the scanning statistics for the configured storage devices. It enables you to quickly and visually evaluate the status of the scans running in your environment. The **Scanning** dashboard displays the following charts that show the various aspects of the scan statistics:

Consolidated status

The pie-chart provides a summary of the consolidated status of all scans on configured shares and site collections in percentage terms.

The consolidated status represents the combined status of scans, full and incremental, since the last full scan.

For example, a full scan successfully completes on a share. If there is subsequent failure of incremental scan on the same share, the consolidated status of scans for that share would display as **Failed**. Similarly, if a full scan on a share was partially successful, and all incremental scans succeed, the consolidated status stays partial. The consolidated status essentially indicates if any recent scans have failed for a share or site collection.

Below the pie chart, you can view the summary of the total failed, successful and partially successful scans on configured shares/site collections. The summary also displays the number of shares/site collections that have never been scanned till date.

Click on an area of the graph to view the filtered list of paths which have the selected consolidated status.

Scan History

The chart provides a graphical representation of the number of scans during a specified duration. Use the drop down to the right of the graph to select the time period for which you want to view the summary.

The color bar chart represents the different states of the scans - Failed [Red]; Successful [Green]; and Partially successful [Yellow]; for a more visual comparison of the data. In each bar, the number of failed, successful, and partially successful states are indicated by the length of the corresponding colors.

Click an area on the bar graph to view detailed information of the scans for the selected scan state on the **Scan History** sub-tab. For example, clicking on the green area of the bar graph displays all successful scans. You can further narrow the list by using other filters on the **Scan History** sub-tab.

Age Of Last Known Good State	The pie-chart provides a high-level overview of the age of last known good state of scan data for a share or site collection.
	The last known good state represents a consistent state of scan meta-data and comprises of a successful full scan followed by zero or more successful incremental scans. This state helps you decide if the scan meta-data can be used for reporting purposes. For example, if the last known good state for a share is 6 months old, it is recommended that you run a full scan for that share before you generate a report.
	In summary, consolidated status indicates if there have been failed scans on a share, which helps you troubleshoot environmental issues, where as the last known good state indicates how recent and consistent the scan data is in the Data Insight index.
	The chart is rendered for one of the following durations:

- More than three months old
- Three months old
- One month old
- One week old
- One day old
- Never

Clicking an area on the pie-chart displays the **Scan Status** sub-tab with detailed information of the scans for the selected age. For example, clicking on the green area of the pie-chart displays the shares or site collections which have a good scan state that is less than a day old.

Viewing the scan status of storage devices

The **Scan Status** sub-tab of the **Settings** > **Scanning** tab provides an overview of the status of the scans for all storage devices that are being monitored by Data Insight. The scans include the scans running on the configured file servers and SharePoint Web applications.

Use the provided search filter, to search for scans based on various criteria, for example, the type of the device. Since it is a dynamic search, the displayed list of scans changes automatically as you select the check box for a filter criteria. For instance, when you select NetApp in the **Device Type** category, and Successful in the **Consolidated status** filter, the application displays a list NetApp shares which have Successful as the consolidated status. Similarly, when you select a Collector node in the **Node Type** filter, Data Insight displays a list of scans running on devices associated with the selected Collector node.

Use the **Scan Status** page to do the following:

■   Review status of the scans on configured storage devices. You can review the status of the last full and incremental scans and also view the scan history and list of paths for which the last scan failed.

■   Filter the list of scans.

■   Start full scans.

To view the scan status

1   Click **Settings** > **Scanning**.

2   The **Overview** sub- tab displays by default. Do one of the following:

■   Click on the appropriate region of the **Consolidated Status** pie-chart to filter the scans based on the status of the scans. Or click on a region of the **Age Of Last Known Good State** pie chart to display the list of shares or site collections with the selected last known good state.

■   Click the **Scan Status** sub-tab.

To start a scan

◆   On the Scan Status page, click **Scan**.

Do one of the following:

■   Click **Scan selected records** to scan the selected objects.

■   Click **Scan all filtered records** to scan all objects that are displayed after applying a filter.

Viewing the scan history of storage devices

Use the Scan History page to view the history of full and incremental scans of all storage devices that are being monitored by Data Insight.

Use the provided search filter in the left-side panel, to view the scans based on various criteria, for example, the type of entity or the status of the scans. Since it is a dynamic search, the displayed list of scans changes automatically as you select

the check box for a filter criteria. For example, when you select Share in the **Type** category and Failed in the **By Status** category, the application displays a list of failed scans on all configured shares for the selected duration. You can also use the free-form Filter text box to enter the search criteria.

To view the scan history

1   Click **Settings**> **Scanning**.

2   The **Overview** sub-tab displays by default. Do one of the following:

- Click on the appropriate region of the **Scan History** chart to filter the scans based on the status of the scans.

- Click the **Scan History** sub-tab.

- Navigate to the share or site collection for which you want to view the scan history. See "Managing shares" on page 161. or See "Managing site collections" on page 181.

Configuring Data Insight global settings

This chapter includes the following topics:

- Configuring SMTP server settings

- Configuring scanning and event monitoring

- About filtering certain accounts, IP addresses, and paths

- About saved credentials

- About archiving data

- Configuring Symantec Data Loss Prevention settings

- Importing sensitive files information through CSV

- Configuring advanced analytics

- About open shares

- Configuring file groups

- Configuring Workspace data owner policy

- Managing Data Insight licenses

- Configuring Management Console settings

- About bulk assignment of custodians

Configuring SMTP server settings

Before Data Insight can send email notifications for events, reports, workflows, and alerts you must configure SMTP details for the Management Server.

To edit the SMTP settings

1   In the Management Console, click **Settings** > **SMTP Settings**.

2   On the SMTP settings page, click **Edit**.

3   Enter the following details:

- A valid SMTP server hostname or IP address.

- The port number for the SMTP mail server used to send email notifications. The default is 25.

- The username for the email server (optional).

- The password for the email server (optional).

- The address from which emails are sent (optional).

- Maximum attachment size. This information is used when Data Insight sends report notifications. Data Insight will not send reports as attachments, if the size of the report is over the specified limit.

4   Click **Save**.

Configuring scanning and event monitoring

Data Insight collects access events using asynchronous APIs, namely, Fpolicy for NetApp filers, the CEE framework for EMC filers, and filter driver for Windows File Servers. You can configure Data Insight to globally turn on or off receipt of event notifications and safeguards related to Fpolicy communication.

Data Insight allows you to disable scanning of all file systems. The scans can be enabled at any convenient time.

You can also configure whether you want the Scanner to fetch the Access Control Lists (ACLs) defined on folders and ownership information for files and folders. For Windows File Servers, you can configure safeguards that prevent an agent from taking up too much disk space.

To configure scanning and event monitoring

1   In the Management Console, click **Settings** > **Scanning and Event Monitoring**.

    You can view the state of scanning and event monitoring.

2   To change the state of a process to Enabled or Disabled, click **Edit**.

3 Do one of the following:

- Select the check box for the process that you want to enable.

- Clear the check box for the process that you want to disable.

4 Click **Save** to save the changes.

Table 3-1 Scanning and Event Monitoring options

Option	Description
Scan File System meta-data	Clear the check-box to turn off all future file system scanning on all filers. Once you save the setting, it will also stop all currently running scans.
Get Folder ACLs	Clear the check box if you do not want Scanner to fetch Access Control List (ACLs) for folders during scanning.
	If you disable this option, the **Workspace** > **Permissions** tab in the Console is disabled and permission related reports will not produce any data. If you do not need permissions data, you can disable this option to make the scans run faster.
Get Ownership information for files and folders	Clear the check box if you do not want Scanner to fetch the Owner attribute for files and folders.
	Ownership information is used to determine ownership for data when access events are not available. If you do not need this information, you can disable this option to make scans run faster.

Table 3-1   Scanning and Event Monitoring options *(continued)*

Option	Description
Throttling for NetApp filers	Select **Throttle scanning based on latency of the filer** to enable throttling of Data Insight scans for NetApp 7-mode and Cluster-Mode file servers. This option is not selected by default.
	Data Insight collects latency information from NetApp file servers. It can use this information to throttle scanning, if latency of the file server increases above a certain level. This ensures scanner does not put additional load on the file server during peak load conditions.
	You can configure the following parameters to enable throttling for NetApp file servers:
	■ Latency threshold - Specify latency in milliseconds, which when crossed, should throttle scanning for the file server.
	■ Minimum pause - Specify the minimum duration (in milliseconds) for which the scanner should pause between paths when in throttling mode.
	■ Back off value - If increased latency is sustained, pause interval will be increased by the Back off value specified (in milliseconds).
	■ Maximum pause - Specify the maximum pause interval for the scanner (in milliseconds). If exceeded, pause interval is no longer incremented by Back off value.
Monitor file system access events	Clear the check box to stop Data Insight from collecting access events from all file servers. In case of NetApp, it means all collector nodes will disconnect their Fpolicy connections to file servers.

Table 3-1   Scanning and Event Monitoring options *(continued)*

Option	Description
Disk Safeguard settings for Windows File Server agents	Select **Enable node safeguard** check box to monitor the disk usage on the Windows File Server node, and prevent it from running out of disk space by implementing safeguards.
	You can specify the threshold for disk utilization in terms of percentage and size. The DataInsightWatchdog service initiates the safeguard mode for the Windows File Server node if the free disk space falls under the configured thresholds.
	The DataInsightWatchdog service automatically resets the safeguard mode when the free disk space is more than the configured thresholds.
	You can edit the threshold limits as required. If you specify values in terms of both percentage and size, then the condition that is fulfilled first is applied to initiate the safeguard mode.
	For more information on the DataInsightWatchdog service, see the *Symantec Data Insight Installation Guide*.

Table 3-1   Scanning and Event Monitoring options *(continued)*

Option	Description
Fpolicy Safeguard settings	Select **Enable Fpolicy Safeguard settings** to initiate safeguard mode for the NetApp file servers.
	Select **Enable Fpolicy Safeguard for vfilers** to initiate safeguard mode for virtual file servers by monitoring latency statistics for the associated physical NetApp filers.
	The safeguard mode for the virtual filers is initiated only if the details of the physical filer corresponding to the virtual filer are provided while adding the virtual filer.
	This option is not selected by default.
	Data Insight collects latency information from NetApp file servers. It can use this information to initiate safeguard mode, if latency of the file server increases above or falls below a certain level. When the safeguard is in effect, Data Insight drops its Fpolicy connection to the filer. This ensures event collection does not put additional load on the file server in peak load conditions.
	If the latency on the physical file server increases above the configured threshold, Data Insight disconnects from the associated virtual file server. This information is also displayed on the Data Insight System Overview dashboard.
	Configure the following values:
	■ The high and low thresholds for CIFS and NFS. ■ The number of samples to be considered to calculate the average latency. ■ The minimum time to wait before Data Insight reconnects to the filer after a disconnection.
	See "Add/Edit NetApp filer options" on page 135.

About filtering certain accounts, IP addresses, and paths

You can configure Symantec Data Insight to filter data accesses by specific users, IP addresses, file system paths, and URLs. You can combine these criteria together or use them individually to create a filter.

You can create separate exclude rules for file servers and SharePoint servers. For each of these, Data Insight supports two types of filters:

- Exclude rules for access events

- Exclude rules for Scanner

About exclude rules for access events

You can configure the following types of exclude rules for access events:

Filters for account names or SIDs	Typically used to mask service accounts from registering data accesses into Symantec Data Insight. For example, if an antivirus software performs scans on a mounted share using a specific user account, you can add that user account to a filter. Data Insight omits all accesses made by that service user account.
Filters for IP addresses	Used to filter data accesses from specific IP addresses. Such filters are useful if you have file system scanners configured on certain machines in your environment, whose accesses you want to ignore.
	Note: IP addresses for events are not available for Windows File Servers and SharePoint.

Filters for path names	Filters for path names are of two types, file extension based and path based.
	The file extension based filter specifies the file extensions to be filtered.
	The path based filter specifies the path of a folder and filters out all events which have that path prefix. For path-based filtering, you must specify a fully qualified path prefix or a path relative to the root of each share. You can also use the wildcard (*) with a prefix string to exclude paths that have the string in the path name.
Filter for URLs	Used to filter data accesses from specified Web applications or from SharePoint sites. Use fo wildcard (*) in exclude rules for access events is not supported.

About Exclude rules for Scanner

Scanner supports filtering out a top-level folder for all shares. You can define rules to exclude the scanning of the specified share, or SharePoint URL by the Scanner process.

Scanner does not support excluding folders under a top-level folder. Also, Use of wildcard (*) in exclude rules for access events is not supported.

Adding exclude rules to Data Insight

You must create a rule for every filter you want to add to Symantec Data Insight. The rule must contain a value for at least one criterion that you want to exclude.

To add an exclude rule:

1   In the Console, click **Settings** > **Exclude Rules**.

2   Click **Add Exclude Rule for File System** or **Add Exclude Rule for SharePoint**, as the case may be.

From the drop-down, select **Exclude access** or **Exclude scanning**.

3   On the Add Exclude Rule screen, enter the Exclude rule properties.

4   Click **Save**.

Add/Edit Exclude rule for access events options

Use this dialog box to add a new exclude rule for access events to Symantec Data Insight or to edit the an existing exclude rule.

Table 3-2 Add/Edit file system Exclude rule for access events options

Field	Description
Rule name	Enter a logical name for the Exclude rule.
Usernames/SIDs	Enter the username or SIDs that you want to exclude.
	Note: The usernames must be present in the Data Insight users database, before they can be added to a exclude rule.
IP Addresses	Enter the IP addresses that you want to exclude.
	This filter only applies to NetApp and EMC Celerra file servers.

Table 3-2   Add/Edit file system Exclude rule for access events options
(continued)

Field	Description
Exclude patterns	

Table 3-2   Add/Edit file system Exclude rule for access events options
(continued)

Field	Description
	When defining a file system rule, enter the file extensions or paths that you want to exclude. A CIFS file system path must be fully qualified path in the format, `\\filer\share\folder` or relative to each share, for example, `<name of folder>`. A NFS path must be a fully qualified physical path on the actual file system in the format, `/path/in/the/physical/filesystem`.

The logical operator OR is used create a rule with multiple values of the same dimension and the logical operator AND is used to combine values across dimensions in a rule. For example, if you create a rule to ignore user_foo1, user_foo2, and IP_10.209.10.20, it means that all accesses from IP_10.209.10.20 AND (user_foo1 OR user_foo2) will be ignored.

You can also specify the wildcard (*) in an exlude rule for paths. Data Insight allows the use of wildcard (*) in the following formats in an exclude rule:

- <prefix string> - Events for paths that start with the specified <prefix string> are excluded .
- <prefix string>* - Events on paths that start with the specified <prefix string> are excluded.
- *<string> - Events on paths which start with anything followed by the specified string are excluded.
- *<string>* Events on paths which have the specified string somewhere in the path name are excluded.

For example, if you specify *`<abc>`*, events on all paths that have the string `abc` anywhere in the path name will be excluded.

When defining a SharePoint rule, enter the URL of the SharePoint Web application or |

Table 3-2   Add/Edit file system Exclude rule for access events options
(continued)

Field	Description
	the site.
	You can use the wildcard (*) to exclude events for URLs that contain a specified sting in its name. For example, if you specify <abc>*, events on all URLs that have the string `abc` anywhere in the path name are excluded.
Pattern Type	Select PREFIX or EXTENSION from the **Pattern Type** drop-down. This field is only available for a file system rule.
Rule is enabled	Select the **Yes** radio button to enable the rule and the **No** radio button to disable it.

Add/Edit Exclude rule for Scanner options

Use this dialog box to add a new exclude rule for access events to Symantec Data Insight or to edit the an existing exclude rule.

Table 3-3   Add/Edit file system Exclude rule for Scanner options

Field	Description
Rule name	Enter a logical name for the Exclude rule.
Exclude Patterns	When defining a CIFS file system rule, specify the name of the folder to exclude as `/<name of first level folder>`. For NFS file system rule, specify the name of the folder to exclude as `/<name of first level folder>`
	When defining a SharePoint rule, enter the URL of the SharePoint Web application or the site.
Rule is enabled	Select the **Yes** radio button to enable the rule and the **No** radio button to disable it.

About saved credentials

An authentication credential can be stored as a saved credential in a central credential store. It can be defined once, and then referenced by any number of filers, shares, and Active Directory servers. Passwords are encrypted before they are stored.

The saved credential store simplifies management of user name and password changes.

You can add, delete, or edit stored credentials.

See "Managing saved credentials " on page 46.

Managing saved credentials

You can add saved credentials to Data Insight, view details of the configured credentials and delete one or more saved credentials on the Saved Credentials details page.

You can add new credentials to the credential store. These credentials can later be referenced with the credential name.

To add a saved credential

1   In the Management Console, click **Settings** > **Saved Credentials**, and click **Create Saved Credentials**.

2   Enter the following information:

Saved Credential Name	Enter your name for this stored credential.
	The credential name must be unique within the credential store. The name is used only to identify the credential.
Access Username	Enter the user name for authentication.
Access Password	Enter the password for authentication.
Confirm Password	Re-enter the password.
Domain	Enter the name of the domain to which the user belongs.

3   Click **Save**.

4   You can later edit or delete credentials from the credential store.

You can delete or edit a saved credential.

To delete a saved credential

1   In the Management Console, click **Settings** > **Saved Credentials**.

2   Locate the name of the stored credential that you want to remove.

3   Click the **Delete** to the right of the name.

     A credential can be deleted only if it is not currently used for filers, shares,
     Active Directory, Fpolicy service, EMC Celerra service, permission remediation
     scripts, custom action scripts, Enterprise Vault server, and as Data Insight
     server credentials..

To edit a saved credential

1   Locate the name of the saved credential that you want to edit.

2   Click the **Edit** to the right of the name.

3   Update the user name or password.

4   If you change the password for a given credential, the new password is used
     for all subsequent scans that use that credential.

5   Click **Save**.

For the purpose of access control, only a user assigned the role of Server
Administrator can add, edit, and view all saved credentials. A user assigned the
Product Administrator role can add new saved credentials, but can only view and
edit those credentials which the user has created.

About archiving data

Data Insight stores system events, alerts, and access events on the Indexer worker
node in a pre-determined folder. You can configure Data Insight to automatically
archive access events older that the specified interval to another folder to save
space. Once the data is archived, it is no longer available for querying. You can,
however, restore the data back to the original location on the Indexer node, if
needed.

By default, archived data is automatically moved to `$data/indexer/archive` folder
on each Indexer worker node. You can also configure a different archive folder on
the Indexer nodes. The archive folder is organized by YEAR/MONTH to make
restoring easy. Once data is moved to this folder based on the configured archive
policy, you can do one of the following:

■   Backup the archive folder and delete the archived files from the Indexer node.

■   Or, configure a file system archiving solution like Symantec Enterprise Vault
     File System Archiving to archive all files in the archive folder.

If you want to restore archived data at a later time, you must bring back the appropriate segments from the backup folder to their original location in the archive folder and use the indexcli utility to restore these segments. Once restored, segments are not archived or purged by the data retention policy till they are marked for re-archiving.

Note: You can disable archiving of access events information by enforcing a legal hold on shares and site collections.

See "About purging data" on page 48.

See "Configuring data retention settings" on page 48.

About purging data

If you want Data Insight to automatically delete data, such as access events, system events, and alerts older than specified interval, you can configure a purging policy. Use the indexcli utility if you want to purge data at a more granular level than what you can configure on the Data Retention page on the Management Console. Purged data cannot be restored back at a later time.

Note: You can disable purging of access events information by enforcing a legal hold on shares and site collections.

See "About archiving data" on page 47.

See "Configuring data retention settings" on page 48.

Configuring data retention settings

Data Insight enforces the data retention policy twice a month. Archived index segments can be restored using a command line utility called indexcli.exe. The utility is also useful to enforce a more granular archiving or purging policy, if the global option is not sufficient for your needs.

See indexcli.exe on page 292.

You can configure the duration for which you want Data Insight to retain various types of data and the duration after which you want to purge data. Automatic archiving and purging of data is not enabled by default.

To configure the data retention period

1   Click **Settings** > **Data Retention**.

2   On the Data Retention details page, click **Edit**.

3 Enter the following information:

Archive access data automatically	Do the following:

Do the following:

1 Select the check box to enable archiving of file system or SharePoint events.

2 Enter the age of the data (in months) after which the data must be archived.

3 Enter the path of the archive folder.

In case, you want to archive access data immediately, execute the following command on the Indexer node:

```
INSTALL_DIR/bin/configcli execute_job
DataRetentionJob
```

Purge access data automatically

Select the check box to enable purging of file system or SharePoint access events, and enter the age of the data (in months) after which the data must be deleted.

Purge Data Insight system events automatically

Select the check box to enable purging of Data Insight system events, and enter the age of the data (in months) after which the data must be deleted.

Data Insight system events are displayed on the **Settings** > **Events** page.

Purge alerts automatically

Select the check box to enable purging of alerts, and enter the age of the alerts (in months) after which they must be deleted.

Automatically purge data for deleted shares or site collections

Select the check box to enable purging of data pertaining to deleted shares. This option is enabled by default.

If you clear the check box , the index data for the share or site collection remains on the disk until the global archive or purge takes place. Thus, if you decide to decommission a filer, or site collection, the data continues to be available for reporting.

4 Click **Save**.

See "About archiving data" on page 47.

See "About purging data" on page 48.

Configuring Symantec Data Loss Prevention settings

Data Insight pulls information about sensitive files in a storage environment from Symantec™ Data Loss Prevention (DLP). Data Insight uses this information to raise alerts in response to configured DLP policies. The information about sensitive files and DLP policies is also displayed on the Data Insight Dashboard and ContextMap views. You can use this information to find high-risk shares and folders that violate important DLP policies.

For more information about DLP policies, see the *Symantec™ Data Loss Prevention Administration Guide*.

You can retrieve an incident list that flags sensitive files in your storage environment and create a saved report using the Enforce Server Administration Console. Data Insight uses the DLP Reporting API Web Service to request a list of incident IDs by specifying a saved report ID. A Data Insight process then fetches the sensitive files corresponding to the incident IDs.

Data Insight runs a job at 12:00 A.M. every night to retrieve a list of sensitive files from DLP.

You must configure the settings that allow Data Insight to communicate with Symantec Data Loss Prevention.

To configure Data Loss Prevention settings

1   In the Management Console, click **Settings** > **Data Loss Prevention**.

2   Click **Edit**, and enter the following details:

- The host name in the format *<Domain\User name>* or IP address of the DLP server

- The name of the domain to which the user belongs.

- Users who are assigned more than one role can only log on under one role at a time. Specify the role you want to use to log in to Data Loss Prevention

- The port through which Data Insight connects to the DLP server.

- The user name and password of the account that is used to access the DLP server.

Note: Ensure that the credentials belong to an existing DLP user assigned the Reporting API Client role.

- The ID of the Saved Report.

3   Select the **Configure resources automatically** check box. Selecting this
    option enables Data Insight to import incidents on resources (files and folders)
    that are not monitored by Data Insight. For such files and folders, Data Insight
    displays information about the DLP policies that are violated, but does not
    display any activity information.

The DLP Enforce Server administration console requires SSL transport for all
communication. Data Insight must be able to negotiate the SSL connection with
the Enforce Server. For this purpose, you must import the certificate to the keystore
used by Data Insight.

**To import the SSL certificate from the DLP Enforce Server to Data Insight using
Firefox**

1   Type the URL to connect to a DLP Enforce Server Administration console.

2   On the security certificate warning page, click **I understand the risks**.

3   Click **Add Exception**.

4   On the Add Security Exception page, click **View** to view the certificate details.

5   Click the **Details** tab and click **Export**.

6   From the Save as type drop-down, select X.509 Certificate (DER).

7   Click **Save**.

**To import the SSL certificate from the DLP Enforce Server to Data Insight using
Internet Explorer**

1   Type the URL to connect to a DLP Enforce Server Administration console.

2   On the security certificate warning page, click **Certificate Error** next to address
    bar.

3   Select **View certificates**.

4   Click the **Details** tab, and select the appropriate certificate.

5   Click **Copy to File**

6   In the Certificate Export Wizard, select **DER encoded binary**.

7   Click **Next**.

8   Enter the name of the file and browse to the location where you want to save
    the file.

9   Click **Next**

10 Click **Finish** to save the file.

After the SSL certificate is imported, complete the following steps to import the SSL certificate on the Data Insight server.

To import the SSL certificate on the Data Insight server

1   From the Windows Start menu, select **Run** and type `cmd` in the dialog box to open a command prompt window.

2   Run the following command:

    ```
    cd C:\Program Files\Symantec\DataInsight\jre\bin
    ```

    ```
    .\keytool -importcert -alias dlp -keystore c:\
    DataInsight\data\keys\commd.keystore -trustcacerts -file <file
    path of SSL certificate>
    ```

    Specify `changeit` as the password for the keystore.

    You can now pull a list of sensitive files from Symantec Data Loss Prevention (DLP).

Importing sensitive files information through CSV

Data Insight pulls information about sensitive files in your storage environment from Symantec™ Data Loss Prevention (DLP).

However, if you use a third-party application other than DLP, Data Insight provides you the ability to identify sensitive files in your storage environment using a .csv file. A scheduled Data Insight process reads the .csv file to retrieve the list of sensitive files from Data Insight.

To use a .csv file to classify sensitive files in Data Insight

1 Log in to the Data Insight Management Server.

2 Create a .csv file, in which each line indicates the path of the sensitive file, and policy names which that particular file violates.

For example, you have a `file /foo/bar/info.txt` which violates the policies *Personal Information* and *Hipaa*. And another file `/foo/ssn.pdf` which violates the policy *Personal Information* In this case, create a .csv file as follows:

```
data /foo/bar/info.txt,Personal Information, Hipaa /foo/ssn.pdf,
Personal Information
```

3 Edit the dlp_db.conf file to add the following lines:

```
dlp.csv.enabled=true

external.file.path=<full path to the location of
.csv file on Management Server>
```

The dlp_db.conf file is located in the <Installdir>\conf directory, where `Installdir` is the installation path for Symantec Data Insight.

Note: If the `dlp.csv.enabled` property is set to true in the `dlp_db.conf` file, the Data Insight process uses the .csv file to identify sensitive files, even if DLP is configured in Data Insight.

Configuring advanced analytics

Data Insight provides several advanced data analytics features in the form of Dashboard reports, Social Network Maps, and ContextMap. These features let you make intelligent decisions about your data.

Use the **Advanced Analytics Configuration** page to configure the following criteria:

- The period for which the activity information should be considered to calculate the device and share statistics. This period is also considered for analyzing collaborative activity on a share when displaying Social network Map, and for generating data displayed on the **ContextMap** view.

- The depth of the folder heirarchy with respect to the root of the share that must be evaluated to compute the control points within a share. The default folder depth for computing control points within a share is 5. This means that by default Data Insight evaluates the folder hierarchy 5 levels deep to calculate the control points within a share.

Data Insight displays the information about control points on the **ContextMap** view on the **Workspace** tab of the Management Console.

- The interval for refreshing the data that is displayed on the **Device** and **Shares and Site collections** tabs of the Data Insight dashboard. This interval is also considered for generating data that is displayed on the **ContextMap** view on the **Workspace** tab of the Management Console.

  You must ensure that you decide the frequency of refreshing the data judiciously, because the statistics are calculated for all the configured devices, shares, and site collections.

Note: The Data Insight dashboard does not display any data, if a summary report has not run even once.

The page also displays information about refresh cycles that have failed. Click on the **Failed** link to download the logs that you can use to troubleshoot the cause for failure.

To configure advanced analytics

1   In the Management Console, click **Settings** > **Advanced Analytics**. The existing analytics settings display by default.

2   Click **Edit** to change the appropriate settings.

3   Click **Save** to save the changes.

4   Click **Compute Now** to refresh the data on the Data Insight dashboard.

See "About open shares" on page 55.

See "Viewing summary of shares and site collections" on page 24.

Choosing custom attributes for advanced analytics

Data Insight lets you configure the custom attributes that you can use for the purpose of advanced analytics, reporting, and remediation.

See "Configuring attributes for advanced analytics" on page 73.

The configured attributes are available for selection on the Advanced Analytics Manager. Use this page to mark the attributes that you want to consider for analytics and for grouping of users based on the primary attribute.

The primary attribute is used to best identify users in the Social Network Map cluster groups. It is also used to distinctly categorize the users in the cluster groups. Ensure that the attribute that you choose as the primary attribute is the one that is populated for the majority of the users in your directory service. For example, good choices for primary attribute would be department or cost code. Attributes such as, employee

ID or title are different for each user, and should ideally not be used as the primary attribute.

To choose attributes for advanced analytics

1   On the **Settings** tab, click **Advanced Analytics**.

2   Click **Attributes** to display the **Advanced Analytics Manager**.

The **Available Attributes** panel displays all the configured custom attributes.

3   Select an attribute, and click the right arrow to select an attribute. Similarly, you can click the left arrow to remove an attribute from the list of selected attributes.

4   Use the up and down arrows to set the priority of the attributes for computing the analytics data.

5   From the **Primary grouping attribute** drop-down, select the attribute that you want to use as the primary attribute for identifying users and for creating attribute-based filters.

About open shares

A share is considered to be open due to the permissions assigned on it. You can define the parameters that determine whether a share should be considered as open. However, the open share policy does not allow you to define specific permissions that render a share open.

One of the following parameters determine whether a share is open:

■   If certain group has access to a share, either directly or as a nested group.

■   If more than a certain number of users have access to the share.

Additionally, you can also specify granular criteria for examination, such as:

■   The level at which Data Insight should start examining the paths.

■   The number of levels deep to examine the path permissions.

For example, you can define the following criteria to consider a share to be open:

■   List of groups: Domain Users or Everyone.

■   No of users who have access to the share: 500.

■   Level to start examining permissions: 1.

■   Depth to examine: 3.

According to the above specification, any share that is accessed by the Domain Users/Everyone group or by more than 500 users is considered to be open. For

this purpose, the ACLs are examined from level 1 (root being level 0), and all folders three levels down are examined.

Defining an open share policy helps you to review the permissions assigned at share level and revisit entitlement decisions. You can view the details of the open shares, the size of the data in open shares, and the number of files in open shares on the **Dashboard** tab on the Management Console.

See "Configuring an open share policy" on page 56.

See "About the Data Insight dashboard" on page 21.

Configuring an open share policy

You can define the parameters for an open share on the **Open Share Policy** page of the Management Console. The paths on a storage device that meet the criteria specified in the policy are considered as open.

See "About open shares" on page 55.

To configure an open share policy

1   Click **Settings** > **Advanced Analytics**. The **Configuration** page opens by default.

2   Click the **Open Share Policy** sub-tab.

     The page displays the existing policy that defines the parameters that determine whether a share is open.

3   To change the existing policy, click **Edit**.

4   Enter the criteria for considering a share to be open:

     ■   Enter the threshold for users that have access to the share. Or select the user(s) or group(s) that have access to the share. You can choose to select both conditions simultaneously.

     ■   You can choose to specify the users and groups that have permissions on the share. Use the users and groups selection widget or upload a .csv file with the users and groups information. The selected entities are listed in the Selected Users/Groups pane.
         You can type a name in the search bar to search for a user or group. You can also type a domain name in the **Domain Filter** field to narrow your search to users in a specific domain. You can also filter a user or group from the **Select Filter** drop-down. Select the **All Filtered Users** check box in the Selected Users/Group pane to include all filtered users in the policy definition.
         All the configured users and groups are displayed on the open share policy definition page.

5   Use the Up and Down arrows to define the level in the share hierarchy that the policy should be applied. You can also examine the depth starting from level 0, that is the root of the share.

6   The depth in terms of number of levels of the folder heirarchy for which the permissions should be examined.

7   Click **Save** to save the policy and close the window.

You can use the report.exe utility to exclude certain paths from the open share computation for the dashboard.

See reportcli.exe on page 298.

Configuring file groups

By default, Data Insight sorts files into 18 file groups based on the extension of the files. The file group information is used for reporting on ownership, access pattern, and space consumption on storage devices. Data Insight uses file group information when generating the following reports:

- Inactive Data by File Group

- Consumption by File Group

- Consumption by File Group and Owner

You can modify the default file groups by adding custom extensions or by deleting extensions from the file groups. You can also create new file groups in **File Groups** view from the Management Console or use the `fg.exe` command from the command line interface. When a new extension is added, all indexes update their extension information. This process can take several hours depending on size of your indexes. Thus, new extensions will not show in your report output till all indexes have been updated.

You can issue fthe ollowing command on your Indexer nodes to monitor the status of updating the indexes:

```
indexcli -j
```

You can search for extensions and file groups by using the filter at the top right corner of the screen.

For detailed information on `fg.exe`, see the *Command File Reference*.

To configure file groups

1   In the Management Console, click **Settings** > **File Groups**.

2   To add a new file group, click **Add new file group**.

3   Enter a logical name for the file group.

4   In the Extension field, enter the extension that you want to include in the file group.

5   Click **Add Extension**.

You can add multiple file extensions to a single file group.

6   To modify an existing file group, click the file group.

7   Click the delete icon corresponding to the extension you want to remove, or enter the extension that you want to include in the file group, and click **Add Extension**.

8   To delete an entire file group, click the corresponding delete icon.

Configuring Workspace data owner policy

By default, Data Insight infers owners of files or folders based on the access history. The most active user of a file is considered to be the data owner for the purpose of efficient remediation and data management.

However, you can define a global policy to infer the owners of files or folders based on one of the following criteria:

- The number of read events on the file or folder.

- The number of write events on the file or folder.

- The cumulative count of read and write events on the file or folder

- The creator of the file or folder.

- The user account which last accessed the file or folder.

- The user account which last modified the file or folder.

- Owner of the parent folder.
  If Data Insight is not able compute the owner on the basis of read, write, last accessed, last modified, or create events, the owner of the immediate parent folder is displayed as the owner of the file or folder.

For example, you can define a policy to consider the count of read and write events on a file to determine the data owner. In this case, the user with the most read and write accesses is considered to be the data owner.

The criteria that is defined in the global policy are also considered for determining the data owner in reports. However, you can choose to override the policy when you create an Inferred Data Owner report.

To configure a Workspace data owner policy

1   In the Management Console, click **Settings** > **Workspace Data Owner Policy**.

2   The Data Owner Policy details page displays the current policy.

3   Use the up and down arrow keys to arrange the criteria in the preferred order.

    Data Insight serially evaluates the criteria in the list to compute the data owner.

4   Select the check boxes to exclude the following types of users when Data Insight calculates the data owner:

    ■   Disabled users - These are the users whose accounts are disabled in the directory service, but are still present in the Data Insight configuration.

    ■   Deleted users - These are the users whose accounts are removed from the directory service.
        Accounts of users are typically disabled or deleted in the directory service when a user leaves an organization.

    ■   Unresolved SIDs - These are accounts for which the corresponding user information is missing in the directory service.

    For example, you define a policy that considers the number of write events on a file to compute the data owner. If the user account with the most write events on a particular file is either deleted or disabled, Data Insight excludes such user from the data owner computation. It then considers the next criteria in the list to determine the owner of the file.

5   Select the groups or users that you want to exclude from the scope of the policy. Double-click the group or user to select it. The selected data set is listed in the **Exclusion List** pane.

    Use the Domain Filter drop-down to search for specific users, users in a particular domain, built-in user accounts, and SIDs that have been migrated to another domain.

6   Click **Save**.

Managing Data Insight licenses

When you purchase Symantec Data Insight, you must install the Data Insight license file. License files have names in the format *name*.slf.

If you do not have a valid license, Data Insight displays a warning in red in the footer of the Management Console screen.

To install a license

1   Obtain the new license file.

2   In the Management Console, click **Settings** > **Licensing**.

3   On the Licensing page, click **Add/Update License**.

4   On the Add new license page, browse to the new Data Insight license file that you have downloaded, and click **Upload**.

Configuring Management Console settings

In the Console Settings view, you can configure global settings that apply to various tasks that you carry out on the Management Console.

To configure the Console settings

1   Click **Settings** > **Console Settings**.

2   Click **Edit**.

You can edit any of the following settings:

Session Timeout	Your login session on the Management Console times out after certain period of inactivity. The default timeout period is one hour.
	To configure the session timeout period, enter the time in minutes.
Report Footer Text	You can choose to add a footer to all the reports that you run in the Console. Enter the sentence string that you want to appear in the footer of the report. For example, Proprietary and Confidential.

3   Click **Save**.

For more information about creating reports, see the *Symantec Data Insight User's Guide*.

About bulk assignment of custodians

Data Insight uses the information about data custodians to distribute remediation requests and reports. The custodians are responsible for taking action on the data

assigned to them. To know more about the Data Insight custodians, you can refer to the *Symantec Data Insight User's Guide*.

You can assign custodians to a data resource from the **Workspace** tab. However, assigning custodians on a path at a time can be tedious and time-consuming. You can use the **Custodian Manager** to easily assign multiple custodians using only a few steps.

You can bulk-assign custodians by using any of the two options:

- **Assign by CSV** - You can use a CSV file that contains the information about the paths and their respective custodians for to assign custodians on paths.
  See "Assigning custodians in bulk using a CSV file" on page 61.

- **Assign by owner method** - You can specify the criteria for computing the possible owner of the selected paths, and assign the computed owners as the custodians.
  See "Assigning custodians based on data ownership" on page 62.

Assigning custodians in bulk using a CSV file

To assign custodians by using a CSV file

1   Prepare a `.csv` file containing the following details separated by a comma:

    - Paths for which a custodian is to be assigned.
      For example: `\\30.219.81.23\Managers\Reportees\`

    - Custodian to be assigned to that path.
      To specify a custodian, you can use the user name along with the domain name. For example: *Custodian_name@domain_name*. Alternatively, you can specify the SID of the user.

2   On the Data Insight Management Console, navigate to **Settings** > **Global Settings** > **Custodian Manager**.

3   From the drop-down menu, select **Assign by CSV**.

4   Click **Choose File**.

5   Browse to the location where you have saved the CSV file. Select the CSV file and click **Open**.

6   Click **Assign**.

Note: The custodian assignment in Data Insight can take some time depending on the number of paths. You can view the status of the operation on the **Settings** > **Events** page of the Management Console.

Assigning custodians based on data ownership

Data Insight computes the possible owners for any path based on parameters such as read counts, write counts, last modifiers, last accessors, creators etc. Based on your requirements, you can select any of these parameters and assign them priorities for ownership computation. For computing the owner, Data Insight by default considers the activity period configured on the **Advanced Analytics** page.

To assign custodians in bulk based on data ownership

1   From the Data Insight Management Console, navigate to **Settings** > **Custodian Manager**.

2   From the drop-down menu, select **Assign by owner method**. The pane expands to display the **Custodian Selection Criteria** panel.

3   Note that the **Use default data owner policy** check-box is selected by default. Do any of the following:

   ■   Click the **Use default data owner policy** check-box to let Data Insight use the data owner policy and exclusion rules as defined under the **Workspace Data Owner Policy** setting.

   ■   Clear the **Use default data owner policy** check-box to define your own set of criteria to calculate the data owner. Use the right arrows and the left arrows to select or deselect a criterion. Use the up arrow and the down arrow to change the priority of the criteria.

   Note: When you clear the check-box for **Use default data owner policy**, Data Insight still enforces the exclusion rules for deleted, disabled, and unresolved users as defined under the **Workspace Data Owner Policy** setting.

   ---

4   From the **Paths** panel, select the paths for which you want to assign custodians. The selected paths are displayed under the **Selected resources** tab.

5   Click **Assign** to complete the request.

Note: The custodian assignment in Data Insight can take some time depending on the number of paths. You can view the status of the operation from the **Settings**> **Events** page of the Management Console.

See "About bulk assignment of custodians" on page 60.

Configuring directory service domains

This chapter includes the following topics:

- About directory domain scans

- Adding a directory service domain to Data Insight

- Managing directory service domains

- Fetching users and groups data from NIS+ scanner

- Configuring attributes for advanced analytics

- Deleting directory service domains

- Scheduling scans

- Configuring business unit mappings

- Importing additional attributes for users and user groups

About directory domain scans

Symantec Data Insight periodically scans the configured directory service domains in your organization to fetch information about users and user groups. Data Insight correlates this information with file and folder access logs to provide access and usage reports. This information is stored on the Management Server in the user database. Symantec recommends that you add each such domain to Data Insight whose users access file system and SharePoint resources of your organization. The time it takes to scan a directory service domain depends on the number of users and groups in the domain.

Data Insight supports the following implementations of a directory service:

■ Microsoft Active Directory

■ Network Information Service

■ LDAP

By default, Data Insight also automatically scans local users of all Windows File Server agents, all NetApp and Celerra filers, and SharePoint site collections.

See "Adding a directory service domain to Data Insight" on page 65.

Adding a directory service domain to Data Insight

You can configure Data Insight to scan one or more directory service domains.

To add a directory service domain to Data Insight

1   In the console, click **Settings** > **Directory Services** to display the configured directory service domains in Data Insight.

2   From the **Add New Directory Service** drop-down, select the type of directory service domain you want to add - Active Directory, LDAP, NIS, or NIS+.

3   On the **Add New Directory Service** page, enter the server properties.

4   Click **Save**.

5   On the **Directory Services** listing page, click **Scan Now**.

Once the initial scan is complete, the users and groups appear under the **Workspace** tab.

Add/Edit Active Directory options

Use this dialog box to add an Active Directory server to Data Insight, or edit the properties of an existing Active Directory server.

Table 4-1   Add/Edit Active Directory options

Field	Description
Domain Name	Enter the name of the domain which you want to scan.
	The domain name is used for display purpose only. The domain name that appears on the **Workspace** tab depends on the name set in the domain.

Table 4-1   Add/Edit Active Directory options *(continued)*

Field	Description
Domain Controller IP	Enter the hostname or IP address of the Active Directory domain controller.
Scanning Details	Do the following: 1 Select the saved credentials from the drop-down or specify new credentials. 2 Click **Test Credentials** to test the availability of network connection between the Management Server and the Active Directory Server, and also to verify that the credentials given are correct. Symantec recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the Active Directory domain controller.
Bind Anonymously	Select the check box if you want to allow Data Insight to connect to the Active Directory server without a credential. Use anonymous binding only when connecting to an Active Directory server belonging to a trusted domain.
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit LDAP domain options

Use this dialog box to add a LDAP directory service server to Data Insight.

Table 4-2   Add/Edit LDAP properties options

Field	Description
Fully Qualified Domain Name	Enter the fully qualified name of the domain that you want to scan. Entering the FQDN will automatically populate the User and Group search Base DN fields.

Table 4-2 Add/Edit LDAP properties options *(continued)*

Field	Description
LDAP server address	Enter the hostname and the port of the LDAP server. By default, the LDAP server runs on HTTPS port 389. If TLS is enabled, the LDAP server runs on port 636, by default.
Type	The type of LDAP schema used by the directory service. Data Insight extracts the attributes from the schema attribute file when scanning the domain. Select one of the following: ■ OPENLDAP ■ Sun ONE You can also create a schema attribute file with customized attributes for each LDAP implementation that does not match the defaults. Ensure that you name the file as `ldap_<ldap_type>.conf` and save it at `$data\conf\ldap` on the Management Server.
Search base DN	The Organization Unit (OU) in which all users and groups have been defined.
This directory uses secure connection (TLS)	Select this check box if the LDAP server uses the TLS protocol.

Table 4-2   Add/Edit LDAP properties options *(continued)*

Field	Description
Scanning details	

Table 4-2   Add/Edit LDAP properties options *(continued)*

Field	Description
	Select the saved credentials from the drop-down or specify new credentials.
	The credentials should belong to an LDAP user who has appropriate privileges to scan the LDAP domain.
	If you are specifying scanning credential other than the directory administrator, then make sure that you have specified the correct DN for that user. For example, uid=ldapuser,ou=People,dc=openldap,dc=com. You can connect to the LDAP database to verify the DN for an LDAP user.
	The example below shows the DN of a sample user, ldapuser, created on a Linux openLDAP server: uid=ldapuser,ou=People,dc=openldap,dc=com.
	The DN string may change depending upon the LDAP schema used. Refer to the LDAP schema to get correct DN for the user.
	If specifying a user other than the directory administrator, ensure that the following limits have been set appropriately on the LDAP server:
	■ nsSizeLimit - Specifies maximum entries that are returned in response to a Data Insight scan. Set this attribute to -1 to return unlimited entries. ■ nsLookThroughLimit - Specifies the maximum number of Data Insight user entries checked for matches during a search operation. Set this attribute to -1 to indicate that there is no time limit. ■ limit.nsIdleTimeout - Specifies the time a Data Insight connection to the server can be idle before it is terminated. The value is given in seconds. Set this attribute to -1 to indicate that there is no time limit.
	The schema attribute names for setting these limits may vary depending upon the LDAP implementation. The above example is for

Table 4-2   Add/Edit LDAP properties options *(continued)*

Field	Description
	Sun ONE.
Test Credentials	Click to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the LDAP server. Symantec recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the LDAP server.
Bind anonymously	Select the check box if you want to allow Data Insight to connect to the LDAP server without a credential. Select this option only if the LDAP server permits anonymous connections.
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit NIS domain options

Use this dialog box to add a NIS directory service server to Data Insight.

Table 4-3   Add/Edit NIS properties

Field	Description
Fully Qualified Domain Name	Enter the name of the domain that you want to scan.
Hostname/IP address	Enter the hostname or IP address of the NIS server.
Scanning Details	Click **Test Credentials** to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the NIS server. Symantec recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the NIS server.

Table 4-3 Add/Edit NIS properties *(continued)*

Field	Description
Disable scanning	Select the check box to disable the scanning of the directory server.

Add/Edit NIS+ domain options

Use this dialog box to add a NIS+ directory service server to Data Insight.

Table 4-4 Add/Edit NIS+ properties

Field	Description
Fully Qualified Domain Name	Enter the name of the domain that you want to scan.
Hostname/IP address	Enter the hostname or IP address of the NIS+ server.
Configured in NIS compatibility mode	This check box is only available when adding a NIS+ server. When configuring a NIS+ server, select the **Configured in NIS compatibility mode** check box if the NIS+ server is configured in the NIS compatibility mode. In this mode, Data Insight can fetch the users and groups data from the NIS+ server remotely in most cases. In non NIS-compatible mode or when Data Insight cannot scan users and groups remotely, you must manually fetch the users and groups data from the NIS+ server. See "Fetching users and groups data from NIS+ scanner" on page 73.
Scanning Details	Click **Test Credentials** to verify that the given credentials are correct and to test the availability of network connection between the Management Server and the NIS+ server. Symantec recommends that you test the connection before proceeding to ensure that the Management Server is able to scan the NIS+ server.

Table 4-4	Add/Edit NIS+ properties *(continued)*
Field	**Description**
Disable scanning	Select the check box to disable the scanning of the directory server.

Managing directory service domains

You can add directory service domains to Data Insight, view details of the configured domains, and scan one or more domains on the **Directory Services** listing page.

To manage the directory service domain servers

1   In the Console, click **Settings** > **Directory Services** to display the configured directory service domains in Data Insight.

2   Filter the list by using the **Domain Type** filter or by specifying the name of the domain in the search bar at the top of the page.

3   Review the following information about the configured domains:

   ■   The name of the domain.

   ■   The address of the domain server hosting the domain.

   ■   The type of directory service - Microsoft Active Directory, LDAP, or NIS+.

   ■   The number of users and groups in the directory service domain.

   ■   The additional attributes that Data Insight extracts for the domain.

4   To scan all domains, click **Scan Now**.

   Note: Data Insight scans all domains together because dependencies might exist between the different domains.

5   To edit the scan schedule for the configured domains, click **Edit Schedule**.

   By default, Data Insight scans all domains at 3:00 A.M. everyday.

   On the **Set Directory Scan Schedule** dialog, change the schedule, and click **Update Schedule**.

   The updated schedule is used for all subsequent scans of the configured domains.

6   To edit the properties of a directory service domain, from the **Select Actions** drop-down, select edit **Edit**.

On the directory service properties screen make the necessary changes, and click **Save**.

7   To delete a configured directory service domain, from the**Actions** drop-down, select **Delete**.

Select **OK** on the confirmation message.

Fetching users and groups data from NIS+ scanner

If the NIS+ server is configured in a NIS non-compatible mode, you must manually fetch users and groups data from the NIS+ server. However, you must still add the NIS plus domain with correct information like domain name and IP address to Data Insight.

To get users and groups data

1   Log in as root to the NIS+ server.

2   Open the command prompt, and type the following commands:

To get users data	`niscat passwd.org_dir > users.txt`
To get the groups data	`niscat group.org_dir > groups.txt`

3   Save the `users.txt` and `groups.txt` files at `$data\users\nisplus\`*`example.com`* on the Management Server, where example.com is your domain name.

4   On the Directory Services listing page, click **Scan Now** to import the NIS+ domain data. You can also run the scan job from the command line by executing the following command on the Management Server:

`configcli execute_job ADScanJob`

Configuring attributes for advanced analytics

Data Insight lists the custom attributes along with their display names or aliases that are configured for a domain. It also provides you the ability to edit configured attributes and add new custom attributes to Data Insight. Certain well-known custom

attributes such as Email are added to Data Insight by default. The Email attribute is used to send emails for reporting and remediation.

Data Insight also uses the configured attributes for the following purposes:

- For advanced analytics, included in reports, and displayed on the console.

- To sort users in a Social Network Map cluster based on the attribute distribution.

- To group users who have activity on any given level in the hierarchy. This attribute-based grouping is displayed in the **ContextMap** view of the **Workspace** tab.

- To create attribute-based filters. These filters are used to drill down the Social Network Map to identify specific user(s) with the selected attribute.

To configure user attributes

1   In the Management Console, click **Settings** > **Directory Services**. The **Domains** page that lists all configured domains opens by default.

2   Click the **Attributes** sub-tab.

The page lists certain well-known attributes, such as Email and Manager. You can edit these attributes, but you cannot delete them.

3   Click **Add New Custom Attribute** to add the additional attributes that you want Data Insight to extract from the directory service.

4   On the **Custom Attributes** pop-up, enter the following details:

- The logical name for the attribute as displayed in Data Insight. You can enter a display name of the attribute.

- From the **Type of custom attribute** drop-down, select whether the attribute applies to a user or a group.

Data Insight extracts the following attributes from Active Directory for users :

- displayName

- distinguishedName

- givenName

- objectSid

- sAMAccountName

- memberOf

- primaryGroupID

- userAccountControl

- sn

Whereas, for user groups,Data Insight extracts the following attributes from Active Directory:

- distinguishedName

- sAMAccountName

- memberOf

- objectSid

5   To edit a well-known custom attribute, click **Edit** and do the following:

- Enter the default name of the attribute as used by your directory service.

- Select the domain for which you want to specify an alias for the attribute.

- Enter the name of the attribute. An attribute can be referred to by many different names across domains. Specify the alias by which the attribute is named in the selected domain.
  The **Domain Specific LDAP Attibute Names** field supports the auto-complete feature. Data Insight provides suggestions for custom attributes when you enter part of an attribute name in the field.
  For example, the attribute name, Department, can map to the alias *Dept* in domain A and *Department* in domain B. In this example, select A and enter the *Dept* as the alias.

6   Select the **Do not fetch attribute** check box if the value for the attribute does not exist in the selected domain.

7   Click **Add New LDAP Attribute Name** to add other domain-specific names for the custom attribute.

8   Click **Save**.

For detailed information about using the Social Network Map to analyze collaborative activity, see the *Symantec Data Insight User's Guide*.

See "Choosing custom attributes for advanced analytics" on page 54.

Deleting directory service domains

You can delete a configured directory service domain.

To delete a directory service domain

1   In the Console, click **Settings** > **Directory Services** to display the configured directory service domains.

2   Click **Select Action** > **Delete** for the domain that you want to delete.

3   Click **OK** on the confirmation message.

Note: Users from a deleted directory domain are removed from Data Insight only
after the next directory scan runs.

Scheduling scans

Symantec Data Insight scans configured domains everyday at 3:00 A.M., by default.
You can, however, configure the scanning schedule, as needed.

Data Insight also scans local users of all file servers and site collections that are
managed Data Insight. Information from these scans becomes visible in Data Insight
after the directory scan runs.

See "About directory domain scans" on page 64.

See "Managing directory service domains " on page 72.

Configuring business unit mappings

Symantec Data Insight lets you associate a business unit name and business unit
owner with each user imported from directory services. This information is later
included in the report outputs and also sent to Symantec Data Loss Prevention as
a part of ownership information.

To import business unit mappings

1   Create a .csv file, `bucsv.csv`, in the `users` folder in the Data Insight data
    directory on the Management Server. By default, the `users` directory on the
    Management Server is located at `C:\DataInsight\data\users`.

    The CSV file must contain the following information:

    ■   The name of the user in the format, user@domain name.

    ■   The name of the business unit.

    ■   The name of the business unit owner.

    For example, *John_Doe@mycompany.com,Sales,Greg Smith*

2   The information is imported into the users database when the next Active
    Directory scan runs. To do so immediately, run the following command:

    ```
    adcli.exe --mode importbu
    ```

Note: The domain name that is given in the .csv file must be among the domains
that are scanned by Data Insight.

Importing additional attributes for users and user groups

Data Insight periodically scans directory domains and fetches basic information about users, groups, and their relationships. You can also import the additional attributes for users and user groups from other user management systems into Data Insight users database. This attribute information is later included in the report outputs and also sent to Symantec Data Loss Prevention as a part of ownership information.

To import additional attributes

1   Create and save a `.csv` file, `customattr.csv` in the `users` folder in the Data Insight data directory on the Management Server. By default, the `users` directory on the Management Server is located at `C:\DataInsight\data\users`.

Note: When you save a `.csv` file with multibyte characters, you must select UTF-8 encoding instead of unicode or default encoding.

The .csv file must contain the following information:

- Whether the entity is a user or group.

- The name of the user or group in the format, logon_name@domain name.

- The type of custom attribute, for example, Email or DN. Four types of custom attributes are defined, Integer, String, DNString, and Email.

- The name of the custom attribute.

- The value of the custom attribute. In case a custom attribute has multiple values, list each value separated by commas.

<user/group>,<logon_name@domain>,<value_type>,<attribute_name>,<value1> <value2>,...<value_n>.

For example, *user, John_Doe@mycompany.com, integer, age, 26*

2   The information that is contained in the `.csv` file is imported into the users database when the next Active Directory scan runs. To do so immediately, run the following command:

```
adcli.exe --mode importcustomattr --csvfile <location of csv file>
```

Note: To troubleshoot issues related to import of such attributes, check the log file, `adcli.log` in the `log` folder on the Management Server

---

Configuring NetApp file server monitoring

This chapter includes the following topics:

Prerequisites for configuring NetApp file servers

Before you start configuring NetApp filers, verify the following:

- The filer is accessible from the collector node using the short name or IP address you plan to use when adding the filer.

- There is connectivity to the collector node from the filer using the short name and the Fully Qualified Host Name (FQHN) of the Collector node.

- The DNS lookup and reverse-lookup for hostname of the Collector node from the filer is working fine.

- The standard RPC ports are open in the firewall.

- The local security policies are set. The installer automatically registers the local security policies on Windows 2008 machines which are used as collector nodes. However, if the installer fails to register the security policies, you must set them manually. Click **Administrative Tools** > **Local Security Policy** > **Local Policies** > **Security Options** and change the following settings:

  - Network access: Named Pipes that can be accessed anonymously - Add NTAPFPRQ to the list.

  You must restart the machine after making these changes.

Credentials required for configuring NetApp filers

Table 5-1 lists the set of credentials that are required by Symantec Data Insight during system configuration.

Table 5-1   Credentials for configuring NetApp filers

Credential	Details
Credentials required to configure DataInsightFpolicy service. The DataInsightFpolicy service runs on the Collector and processes the events that are sent by the NetApp filer using the Fpolicy RPC interface. This service must be configured on each Collector that is used to connect to NetApp filers. DataInsightFpolicy service also performs miscellaneous tasks, such as gathering storage information from the filer.	The credential should belong to a user in the domain of which the Data Insight Collector node and the NetApp filers are a part. This user should be a part of the Backup Operators group on the filer. If the filers and the Collector belong to different untrusted domains, you can use the Local System account to run the DataInsightFpolicy service. However, when you add filers, the account you specify during filer configuration must have Backup Operator privileges on the filer.

Table 5-1 Credentials for configuring NetApp filers *(continued)*

Credential	Details
Credentials required during filer configuration through the Symantec Data Insight Management Console.	Required to discover shares and enabling Fpolicy on the NetApp filer. This credential belongs to the NetApp ONTAP user who has administrative rights on the NetApp filer (for example, root) or a domain user who is part of the Administrators group on the filer.
	Or, this credential belongs to the NetApp ONTAP user or a domain user who is a non-administrator user on the filer, but has specific privileges.
	See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 90.
	Note: The domain user can be the same user account that was specified when configuring the DataInsightFpolicy service.
	If you use the Local System account to configure the DataInsightFpolicy service on the Collector, the user you specify here must also belong to the Backup Operators group on the filer.
	See "Preparing Symantec Data Insight for Fpolicy " on page 84.

Table 5-1   Credentials for configuring NetApp filers *(continued)*

Credential	Details
Credentials required for scanning of shares.	

Table 5-1   Credentials for configuring NetApp filers *(continued)*

Credential	Details
	Required for scanning of shares from the NetApp filer.
	When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part. This user must belong to either the Power Users or Administrator's group on the NetApp filer. If the credential is not part of one of these groups, the scanner will not be able to get share-level ACLs for shares of this filer.
	You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system heirarchy.
	You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:
	■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions
	For scanning NFS shares, Data Insight needs a Unix account with at least read and execute permissions on all folders, alongwith at least read permission on all files. By default, Data Insight uses User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Settings** > **Advanced Settings** section of the Collector node.
	See "Configuring advanced settings" on page 203.
	When monitoring only NFS shares, you can specify **Use Local System account** from the scanning credentials drop-down, else you can specify credentials required to scan CIFS shares.

Table 5-1   Credentials for configuring NetApp filers *(continued)*

Credential	Details
	Note: Symantec recommends that the credentials used to discover shares on the NetApp file server must be the same as the credentials used to configure DataInsightFpolicy the service.

Configuring SMB signing

Ensure that Server Message Block (SMB) signing is either turned on or turned off on both the Collector node and the NetApp filer. If SMB signing is turned on, all packets of data sent over a network to a remote host are signed. A mismatch in the setting on the Collector node and the NetApp filer can cause the filer to drop the Fpolicy connection to the Collector node.

To configure SMB signing

1   Check whether the SMB signing option on the NetApp filer, `options.cifs.signing.enable` is set to off or on.

2   On the Collector node assigned to the NetApp filer, open the Windows' Registry Editor (**Start** > **Run** > **regedit**).

3   In Registry Editor, navigate to **HKEY_LOCAL_MACHINE** > **SYSTEM** > **CurrentControlSet** >**SERVICES** > **lanmanserver** > **Parameters**.

4   Modify the following registry entries:

   ■   `enablesecuritysignature` - Enter the value 0 to turn signing off and enter the value 1 to turn signing on.

   ■   `requiredsecuritysignature` - Enter the value 0 to turn signing off and enter the value 1 to turn signing on.

About FPolicy

Symantec Data Insight uses the FPolicy framework provided by NetApp to collect access events from the NetApp filers.

NetApp provides an interface called FPolicy which allows external applications to receive file access notifications from the NetApp Storage subsystem. FPolicy allows partner applications to perform tasks like file access screening and auditing. The FPolicy interface uses Remote Procedure Calls (RPC) and external applications

can use these tools to register with the NetApp Filer as FPolicy servers. FPolicy supports both CIFS and NFS.

The unit of FPolicy configuration on the NetApp filer is called a policy, which is identified by a user specified name. You can configure a policy to monitor all or a list of volumes on the NetApp filer along with a specified set of operations. The monitored operations are open, close, read, write, create, delete, rename, and set attribute. As soon as a file operation is performed on a file or folder on the filer which is being monitored, a notification is sent to the registered FPolicy server asynchronously.

Note: The policy created by Symantec Data Insight should not be shared by any other applications or clients.

By default, Data Insight does not register for read and close events from NetApp filers. Data Insight treats an open event as a read event. This behavior reduces the load on the filer in case of peak traffic loads from third party applications like backups over CIFS. It also does not have an adverse effect for most consumer applications because consumer applications seldom write to a file before first reading it. Data Insight assumes that an open event is almost always be followed by a read event and then optionally by a write event. However, you can customize the default behavior as per your requirements.

See "Enabling export of NFS shares on a NetApp file server" on page 93.

See "Preparing the NetApp filer for Fpolicy" on page 85.

Preparing Symantec Data Insight for Fpolicy

The Symantec Data Insight FPolicy server can reside on the Management Server and/or on each Collector worker node. The Management Server and/or the Collector worker node must register with the NetApp filer to receive audit information. Before you assign a Data Insight server as a collector for a NetApp filer, you must configure the FPolicy service on that server.

To set up the environment for Symantec Data Insight FPolicy service

1   Provision a Windows 2003 or 2008 server in the same Active Directory domain as the filers you wish to monitor using Fpolicy. This machine hosts the FPolicy server. If your filers belong to different untrusted domains, you can add the server to any one domain.

2   Install the Data Insight Collector worker node or the Data Insight Management Server on this server.

3   Login to the Data Insight Management Console.

4   In the Console, click **Settings** > **Data Insight Servers** to open the listing page
     for the server.

5   Select the server from the server list to open the details page for the server.

6   Click the **Services** tab.

7   Select **DataInsightFPolicy** to expand the section..

8   In the **Policy Name** field, enter the policy name that will be enabled on each
     filer, of this node Collector. The default name is *matpol*.

9   Under Credentials, select the saved credentials that the service needs to run
     as. The credentials must belong to the Backup Operators group on the NetApp
     filer that is being monitored by the Collector. If the filers in your organizations
     belong to multiple untrusted domains, select **LocalSystem account** from the
     drop-down.

     See "Credentials required for configuring NetApp filers" on page 79.

10   Click **Configure** to apply these settings to the server and start the FPolicy
     service.

See "Configuring SMB signing" on page 83.

See "About FPolicy" on page 83.

Preparing the NetApp filer for Fpolicy

The Symantec Data Insight Fpolicy server registers with the NetApp filer and
receives file access events from it. Fpolicy has to be enabled and configured on
that NetApp filer. Symantec recommends that you automatically enable auditing
when adding filers.

See "Adding filers" on page 134.

However if you want more control on the shares you want to monitor use the manual
steps. The manual steps are valid for Netapp ONTAP version 7.0 and higher.

Note: The steps below assume that the name of the policy is *matpol*.

To configure the Fpolicy on the NetApp filer using manual steps

1   Launch a Telnet session with the filer and run the following commands, as
     appropriate:

     ■ To create a policy:

     ```
     fpolicy create matpol screen
     ```

■ To enable a policy:

```
fpolicy enable matpol -f
```

2   Use the following optional commands for monitoring:

■ To set the Fpolicy for CIFS to monitor specific events:

```
fpolicy mon add matpol -p cifs -f write,
open,close,delete,rename,create
```

■ To set the Fpolicy for NFS to monitor specific events:

```
fpolicy mon add matpol -p nfs -f create,delete,rename,write,
open,link,symlink,setattr
```

■ To monitor specific events on NetApp filer versions 7.3 or higher:

■ Enable set attributes operation:

```
fpolicy options cifs_setattr on
```

```
fpolicy options nfs_setattr on
```

■ Add events to be monitored:

```
fpolicy mon add matpol -p cifs -f write,
open,close,delete,rename,create,setattr
```

```
fpolicy mon add matpol -p nfs -f create,delete,rename,write,
open,link,symlink,setattr
```

■ To see details of a configured policy:

```
fpolicy show matpol
```

■ To disable monitoring of specific events:

```
fpolicy mon remove matpol -p cifs -f write,
open,close,delete,rename,create
```

```
fpolicy mon remove matpol -p nfs -f create,delete,rename,write,
open,link,symlink,setattr
```

■ To disable use of a policy:

```
fpolicy disable matpol
```

- To delete a policy:

```
fpolicy destroy matpol
```

3 To add a domain user to the administrator's group:

```
useradmin domainuser add domain-username
-g Administrators
```

Note: The domain user is the user who is configured to run the Fpolicy service on the collector.

To configure a non-administrator user:

See "Preparing a non-administrator domain user on the NetApp filer for Data Insight" on page 90.

4 To display a list of users who are already configured:

```
useradmin domainuser list -g Administrators
```

A list with the SIDs of the configured domain users appears. To resolve the SIDs, run the following command:

```
cifs lookup SID
```

See "Configuring SMB signing" on page 83.

Preparing the NetApp vfiler for Fpolicy

The Symantec Data Insight Fpolicy server can register with the NetApp vfiler and receive file access events from it. Fpolicy has to be enabled and configured on that NetApp vfiler manually.

To configure the Fpolicy on the NetApp vfiler using manual steps

1 Launch a Telnet session with the filer and run the following commands, as appropriate:

- To get the vfiler name:

```
vfiler status
```

Choose the name of the vfiler that you want to configure and then perform the following operations for that vfiler. Ignore the name, *vfiler0*, which is the default name given to the physical filer by NetApp.

Note: Consult your system administrator to get the IP address of the vfiler. You will need this IP address while adding the vfiler from the Management Console.

See "Adding filers" on page 134.

- To create a policy:

```
vfiler run vfilername fpolicy create matpol screen
```

- To enable a policy:

```
vfiler run vfilername fpolicy enable matpol -f
```

2 Use the following optional commands for monitoring:

- To set the Fpolicy for CIFS to monitor specific events:

```
vfiler run vfilername fpolicy mon add matpol -p cifs
 -f write,open,close,delete,rename,create
```

To set Fpolicy for NFS to monitor specific events:

```
vfiler run vfilername fpolicy mon add matpol -p nfs -f create,
delete,rename,write,open,link,symlink,setattr
```

- To set the Fpolicy for CIFS to monitor specific events on NetApp filer versions 7.3 or higher:

  - Enable set attributes operation:

    ```
    vfiler run vfilername fpolicy options cifs_setattr on
    ```

    ```
    vfiler run vfilername fpolicy options nfs_setattr on
    ```

  - Add events to be monitored:

    ```
    vfiler run vfilername fpolicy mon add matpol -p cifs
    -f write,open,close,delete,rename,create,setattr
    ```

    ```
    vfiler run vfilername fpolicy mon add matpol -p nfs -f create,
    delete,rename,write,open,link,symlink,setattr
    ```

- To see details of a configured policy:

```
vfiler run vfilername fpolicy show matpol
```

- To disable monitoring of specific events:

  ```
  vfiler run vfilername fpolicy mon remove matpol -p cifs
  -f write,open,close,delete,rename,create

  vfiler run vfilername fpolicy mon remove matpol -p nfs -f create,
  delete,rename,write,open,link,symlink,setattr
  ```

- To disable use of a policy:

  ```
  vfiler run vfilername fpolicy disable matpol
  ```

- To delete a policy:

  ```
  vfiler run vfilername fpolicy destroy matpol
  ```

  where, *vfilername* is the name of the vfiler you want to configure.

3 To add a domain user to the administrator's group:

```
vfiler run vfilername useradmin domainuser
add domain-username -g Administrators
```

Note: The domain user is the user who is configured to run the Fpolicy service
on the collector. See "Preparing the NetApp filer for Fpolicy" on page 85.

To configure a non-administrator user:

See "Preparing the NetApp filer for Fpolicy" on page 85.

4 To display a list of users who are already configured:

```
vfiler run vfilername useradmin domainuser list
-g Administrators
```

A list with the SIDs of the configured domain users appears. To resolve the
SIDs, run the following command:

```
cifs lookup SID
```

Preparing a non-administrator domain user on the NetApp filer for Data Insight

To configure a NetApp filer from the Management Console, you can use an account which is not in the administrators group on the NetApp filer, but has some specific privileges.

Perform the following steps on the NetApp filer console to add a non-administrator user, for example, *testuser*.

To create a non-administrator user

1   Create a new role, for example *testrole*, using the `useradmin` utility on the filer.

2   Add the login-* and API-* capabilities to the role.

    For example, `useradmin role add testrole -a login-*,api-*`.

    You can also choose to assign specific capabilities to the role.

    Table 5-2 provides a detailed description of each capability.

3   Create a new group, for example, *testgroup* and apply the role *testrole* to it.

    For example, `useradmin group add testgroup -r testrole`.

4   Add the user *testdomain\testuser* to *testgroup*

    For example, `useradmin domainuser add testdomain\testuser -g testgroup`.

5   Add the user *testdomain\testuser* to *Backup Operators* group.

    For example, `useradmin domainuser add testdomain\testuser -g Backup Operators`.

    Note: For vfilers, append the above command-line examples with `vfiler run <vfilername>`.

Table 5-2 Additional capabilities for adding a non-administrator user account

Capability	Description
`login-http-admin`	Enables you to log into the NetApp filer and run commands. With this capability, you can get latency statistics (for scan throttling), volume size information, or discover shares.

Table 5-2 Additional capabilities for adding a non-administrator user account *(continued)*

Capability	Description

`api-system-get-ontapi-version`
 	
`api-system-get-version`
	Enables you to get the ONTAPI version number and the system version number respectively. These are required to set the login handle context properly. Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. Also, if these capabilities are absent, you cannot execute any APIs including those required to discover shares, and get latency statistics.

`api-fpolicy-set-policy-options`	Enables you to set a flag on the NetApp filer to enable ACL change notifications. If you choose not to supply this capability to Symantec, the filer administrator must set this property manually using the root telnet console (`fpolicy options <policyname> cifs_setattr on`).

`api-fpolicy-list-info`
 	
`api-fpolicy-create-policy`
 	
`api-fpolicy-enable-policy`
 	
`api-fpolicy-disable-policy`
 	
`api-fpolicy-destroy-policy`
	Required to enable Data Insight to automatically create and enable FPolicy on the NetApp filer. Optionally, the filer administrator can set up the policy using the root telnet console.
	
	Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.

`api-fpolicy-server-list-info`
 	
`api-fpolicy-list-info`
	Used to retrieve useful statistics from the NetApp filer, such as, total event count and event failures. These APIs are used every two hours so they do not load the system. However, absence of these capabilities does not cause any problems.

Table 5-2 Additional capabilities for adding a non-administrator user account *(continued)*

Capability	Description
`api-license-list-info`	Used to check if this NetApp filer has CIFS license and print the appropriate diagnostic message.
	Data Insight reports a failure when you create or enable FPolicy on the filer or when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.
`api-options-set`	Used to enable the global FPolicy flag on the NetApp filer.
	In the absence of this capability, FPolicy creation fails. Also, Data Insight reports a failure when you test the connection to the filer, if these capabilities are absent. However, you can save a filer with the policy that is already created and enabled by the filer administrator without testing the connection to the filer.

`api-cifs-share-list-iter-start`
 	
`api-cifs-share-list-iter-next`
 	
`api-cifs-share-list-iter-end`
	Used to discover shares on the NetApp filer. Absence of these capabilities can result in a failure to discover the shares. Optionally, you can add shares manually from the Data Insight console.

`api-perf-object-get-instances-iter-start`
 	
`api-perf-object-get-instances-iter-next`
 	
`api-perf-object-get-instances-iter-end`
	Used to get CIFS latency information from the NetApp filer, which enables the self-throttling scan feature of Data Insight. Absence of these APIs can cause scanner to fail if you enable the feature to throttle scanning.

`api-volume-list-info`	Used to periodically fetch size information for NetApp volumes.
`api-nfs-exportnfs-list-rules`	Used to discover all NFS shares that are exported from the NetApp filer. If this capability is absent, these NFS shares are not discovered.

Table 5-2   Additional capabilities for adding a non-administrator user account
(continued)

Capability	Description

`api-net-ping`
 	
`api-net-resolve`
	Used to check network connectivity from the filer to the Data Insight Collector. These APIs are useful to run some diagnostic checks on the filer. However, such checks can also be done manually by the NetApp administrator, and hence these APIs are not mandatory.

`api-fpolicy-volume-list-set`	Used to set the volume names on the filer which are to be excluded from being monitored by FPolicy. If this capability is absent, you cannot exclude volumes from being monitored by FPolicy from the Data Insight console . However, you can exclude volumes manually using the CLI on the NetApp console. See "Excluding volumes on a NetApp file server" on page 94.
`api-system-get-info`	Used to discover the NetApp system serial number. The information is used by external reporting tools like Veritas Operations Manager (VOM) to report about the NetApp filers that Data Insight monitors. This privilege is mandatory if VOM integration is required.

Enabling export of NFS shares on a NetApp file server

Before you add a NetApp filer to Data Insight, you must enable the export of NFS shares on the NetApp filer to allow Data Insight to discover the NFS shares on the filer.

To enable export of NFS shares on the NetApp filer

1   On the NetApp FilerView Web console, select **NFS** > **Manage exports**.

2   On the Export wizard, click **Add Export** or you can edit the existing exports to modify them.

3   On the first page of the wizard ensure that you have at least selected read only and root access, Other options can also be specified, as required, and click **Next.**

4   Define the export path and give read only access to the Data Insight Collector node, and click **Next**.

5   On the Read-Write Access page, enable read-write access for all clients or for specific hosts, as per your need.

6   Click **Next**.

7   On the Root Access page, define root access to the the Data Insight Collector node, and click **Next**.

8   On the Security page, accept the default options, and click **Next**.

9   On the Summary page, review the configuration and click **Commit** to save the changes.

See "Adding filers" on page 134.

Excluding volumes on a NetApp file server

To optimize the performance of a NetApp file server. you can choose to exclude a set of volumes from being monitored by FPolicy using NetApp commands.

If a configured NetApp filer hosts applications that create a huge number of access events on certain volumes, it may cause a high I/O latency for CIFS. Symantec recommends that you exclude such volumes from being monitored by FPolicy.

To exclude volumes on the NetApp filer

1   Log in to the Net App file server.

2   Execute the following command:

```
fpolicy vol exc add <name of policy> <comma separated name list
of volumes to be excluded>
```

Note that the list of volumes is not the path of the volumes, but names of the volumes.

Handling NetApp home directories in Data Insight

Data ONTAP enables you to create users' home directories on the NetApp storage system. You can access your home directory share the same way you can access any other share on the file system. You can connect only to your home directory matching your user name without seeing other users' home directories. The user name for matching a home directory can be a Windows user name, a domain name followed by a Windows user name, or a UNIX user name. Refer the NetApp documentation for more details about how Data ONTAP matches a home directory with a user,

Data Insight cannot parse FPolicy events coming from home directory shares because of insufficient information regarding mapping to existing shares. To monitor events on home directory shares, you must configure an artificial share in Data Insight. This share represents all home directories on a NetApp filer, and events from the users' home directories are directed to this share. Note that Data Insight does not scan the artificial share.

To configure an artificial share in Data Insight

1   Log in toData Insight Management Console.

2   Click **Settings** > **Filers**.

3   Click the filer for which you want to add a share.

4   On the filer details page, click **Monitored Shares**.

5   On the **Monitored Shares** page, click **Add New Share**.

6   On the New Share pop-up, enter the following details:

   ■   The name of the share that you want to add - `CIFS.HOMEDIR`.

   ■   The physical path of the share on the filer -`/CIFS.HOMEDIR.`.

7   Select the **Define custom cron schedule** radio button, and select **Never** from the schedule drop-down.

8   To verify that the share is added, on the Management Console, navigate to **Settings** > **Filers**.

9   Click the relevant filer. On the filer details page, click the **Monitored Shares** tab to review the list of configured shares.

10   Repeat the steps for each NetApp filer for which you want to add an artificial share.

Chapter **6**

Configuring clustered NetApp file server monitoring

This chapter includes the following topics:

- About configuring a clustered NetApp file server

- About configuring FPolicy in Cluster-Mode

- Pre-requisites for configuring NetApp file servers in Cluster-Mode

- Credentials required for configuring a clustered NetApp file server

- Preparing a non-administrator local user on the clustered NetApp filer

- Preparing Data Insight for FPolicy in NetApp Cluster-Mode

- Preparing the ONTAP cluster for FPolicy

About configuring a clustered NetApp file server

Symantec Data Insight supports the monitoring of clustered NetApp 8.2 NAS devices along with standalone NetApp file servers.

In Data ONTAP Cluster-Mode (C-mode), a Vserver represents the individual nodes in the cluster. A Vserver is a logical unit within an ONTAP cluster which contains a CIFS server. A Vserver facilitates data access. For the purpose of monitoring, a CIFS server within a Vserver represents a filer for Data Insight. Every Vserver can have multiple physical nodes. Data Insight uses the FPolicy interface on ONTAP to receive event notifications, which are sent in XML format over a TCP connection to the Data Insight's FPolicy server.

Data Insight automatically discovers the following entities in the cluster environment:

- The Vservers and the CIFS server underneath each Vserver.

- The CIFS shares for each CIFS server.

To enable Data Insight to discover shares in the cluster, you must ensure the following:

- Complete all the pre-requisites.
  See "Pre-requisites for configuring NetApp file servers in Cluster-Mode" on page 98.

- Provide correct cluster user credentials when adding the clustered file server. You can also use the credentials of a local cluster administrator user.
  See "Credentials required for configuring a clustered NetApp file server" on page 99.

- Add the clustered NetApp file server to the Data Insight configuration.
  See "Adding filers" on page 134.

After you add a CIFS server from the ONTAP cluster to Data Insight, Data Insight automatically enables FPolicy on the corresponding Vserver on the ONTAP cluster. This operation helps the ONTAP cluster register with the Data Insight FPolicy server. Note that in ONTAP 8.2 Cluster-Mode, the connection is initiated by ONTAP.

See "Preparing the ONTAP cluster for FPolicy" on page 102.

Once a TCP connection is established with the Data Insight FPolicy server (Collector node), it starts receiving access event information from the ONTAP cluster.

Note: Data Insight does not support scanning of NFS shares on a clustered NetApp file server.

About configuring FPolicy in Cluster-Mode

FPolicy is a file access notification framework which allows screening of NFS and CIFS accesses. See "About FPolicy" on page 83.

In case of a Cluster-Mode configuration, FPolicy requires that all the nodes in the NetApp cluster are running Data ONTAP 8.2 or later.

To receive access events from the ONTAP cluster, you must first install the DataInsightFPolicyCMod service on the Data Insight Collector node that is configured to monitor the NetApp cluster. The service communicates with the NetApp cluster and fetches information about the file operations that Data Insight monitors. Data Insight automatically configures the list of operations to monitor when you add the

NetApp cluster to Data Insight. However, you can also choose to configure the operations to be monitored manually.

Note: The FPolicy server for a NetApp standalone (7-mode) configuration and C-Mode configuration can co-exist on the same Data Insight Collector node.

For detailed information about the DataInsightFPolicyCmod service, see the *Symantec Data Insight Installation Guide*.

See "Preparing Data Insight for FPolicy in NetApp Cluster-Mode" on page 101.

See "Preparing the ONTAP cluster for FPolicy" on page 102.

Pre-requisites for configuring NetApp file servers in Cluster-Mode

Before you can start using Data Insight to monitor the NetApp file servers operating in Cluster-Mode, verify the following:

- ONTAP 8.2 cluster is configured in accordance with NetApp documentation.

- The DataInsightFPolicyCmod service is installed on the Data Insight Collector that is configured to monitor the NetApp filers operating in Cluster-Mode.

- FPolicy is configured for every Vserver in the ONTAP cluster.
  See "Preparing the ONTAP cluster for FPolicy" on page 102.

- The ONTAP Cluster Management host is accessible from the Collector using the short name or IP address.

- Data Insight Collector should be able to communicate with port 80 on the ONTAP cluster management host.

- Data Insight should be able to communicate with the CIFS server hosted within the ONTAP cluster.

See "About configuring a clustered NetApp file server" on page 96.

Credentials required for configuring a clustered NetApp file server

Table 6-1	Credentials for configuring NetApp file servers in Cluster-Mode
Credential	**Details**
Credentials required during filer configuration through the Symantec Data Insight Management Console.	Required to discover shares and enabling FPolicy on the NetApp filer. This credential belongs to the NetApp ONTAP cluster administrator user who is a local user on the ONTAP cluster. Or, this credential belongs to the ONTAP cluster non-administrator user with specific privileges. See "Preparing a non-administrator local user on the clustered NetApp filer " on page 100.
Credentials required for scanning of shares.	Required for scanning of shares from the NetApp filer. When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part. This user must belong to either the Power Users or Administrator's group on the NetApp filer. If the credential is not part of one of these groups, the scanner is not able to get share-level ACLs for shares of this filer. You do not need this privilege if you do not want to get the share-level ACLs. In this case you only need privileges to mount the share and scan the file system hierarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions

Preparing a non-administrator local user on the clustered NetApp filer

To configure a NetApp cluster file server from the Data Insight Management Console, you can use a local user account which is not in the administrators group on the NetApp cluster, but has some specific privileges.

To create a non-administrator user

1   Launch a Telnet session with the NetApp Cluster Management host.

2   Create a new role, for example *testrole*, using the `useradmin` utility on the filer.

3   Run the following commands to create the role with specific privileges:

```
security login role create -role testrole -cmddirname "version"
-access all

security login role create -role testrole -cmddirname "vserver cifs"
-access readonly
```

4   `security login role create -role testrole -cmddirname "vserver"`
     `-access readonly`

   `security login role create -role testrole -cmddirname`
   `"vserver cifs share" -access readonly`

   `security login role create -role testrole -cmddirname "volume"`
   `-access readonly`

   `security login role create -role testrole -cmddirname`
   `"vserver fpolicy policy" -access all`

   `security login role create -role testrole -cmddirname`
   `"vserver fpolicy" -access readonly`

   `security login role create -role testrole -cmddirname`
   `"vserver fpolicy enable" -access all`

   `security login role create -role testrole -cmddirname`
   `"vserver fpolicy disable" -access all`

   `security login role create -role testrole -cmddirname`
   `"statistics" -access readonly`

5   Run the following command to create a local user, for example, *testuser*, and
     assign the role that you created in 3 to the user:

   `security login create -username testuser`
   `-application ontapi -authmethod password -role testrole`

Preparing Data Insight for FPolicy in NetApp Cluster-Mode

The Symantec Data Insight FPolicy server is represented by DataInsightFpolicyCmod service which runs on each Collector worker node. Before you assign a Data Insight server as a Collector for a clustered NetApp filer, you must configure the Cluster-Mode FPolicy service on that server.

To configure the DataInsightFPolicyCmod service

1   Provision a Windows 2003 or 2008 server. Symantec recommends a minimum requirement of a Windows 2008 64-bit server with 4 to 8GB RAM and a quad core processor. A higher configuration may be required if the load on the FPolicy server is high.

    This computer hosts the FPolicy server.

2   Install the Data Insight Collector worker node or the Data Insight Management Server on this server.

3   Log in to the Data Insight Management Console.

4   In the Console, click **Settings** > **Data Insight Servers** to open the listing page for the server.

5   Select the server that is configured to monitor the NetApp clustered file server to open the details page for the server.

6   Click the **Services** tab.

7   Select **DataInsightFPolicyCmod** to expand the section.

8   To configure the service, enter the following details:

    ■   The user-configured name to create FPolicy the ONTAP cluster.

    ■   The IP address of the Data Insight Collector running the FPolicy server. The NetApp ONTAP Cluster-Mode filer connects with the DataInsightFPolicyCmod service running on the Data Insight Collector node on this IP address.

    ■   The TCP port used on the DataInsightCollector server. The NetApp ONTAP Cluster-Mode filer connects on this port to the Data Insight Collector. Ensure that this port is not blocked by firewall.

9   Click **Configure**.

See "Configuring Data Insight services" on page 201.

Preparing the ONTAP cluster for FPolicy

The Symantec Data Insight FPolicy server registers with the ONTAP cluster and receives file access events from it, if FPolicy is enabled and configured on the corresponding Vservers in the cluster. Symantec recommends that you automatically enable auditing when adding the clustered filers.

When you enable FPolicy from the Data Insight console, Data Insight automatically does the following:

- Creates an FPolicy with a unique name.
- Creates an FPolicy engine by specifying the server IP address and the server port.
- Creates a CIFS event object.

Once you enable FPolicy on the Vserver, it initiates a TCP connection to the Data Insight FPolicy server.

To configure FPolicy on the ONTAP cluster using manual steps

1   Launch a Telnet session with the Vserver on which you want to configure FPolicy.

2   Run the following command to create an External Engine Object on the ONTAP shell:

```
diontapclust::> vserver fpolicy policy external-engine create
-vserver <Vserver name> -engine-name <choose an external engine
name> -primary-servers <IP address of Data Insight FPolicy server>
-port <port number on which Data Insight FPolicy server is
listening> -extern-engine-type asynchronous -ssl-option no-auth
```

3   Run the command to configure FPolicy to monitor specific CIFS events:

```
diontapclust::> vserver fpolicy policy event create -vserver
<Vserver name> -event-name <choose an event name> -protocol cifs
-file-operations create, create_dir, delete, delete_dir,
read,close, rename, rename_dir -filters first-read,
close-with-modification
```

4   Run the command to create a resident FPolicy on the Vserver:

```
diontapclust::> vserver fpolicy policy create -vserver <Vserver
name> -policy-name <choose a policy name> -events <specify list
of events> -engine <specify engine name> -is-mandatory false
```

5   Run the command to configure the scope of FPolicy:

```
diontapclust::> vserver fpolicy policy create -vserver <Vserver
name> -policy-name <choose a policy name> -events <specify list
of events> -engine <specify engine name> -is-mandatory false
```

7

Configuring EMC Celerra monitoring

This chapter includes the following topics:

- Configuring EMC filers
- Credentials required for configuring EMC Celerra filers

Configuring EMC filers

Symantec Data Insight supports EMC Celerra, EMC VNX, and EMC Isilon file servers. Data Insight uses the EMC Common Event Enabler (CEE) framework to collect access events from the EMC filers.

As a prerequisite, you must download and install the CEE framework from the EMC website.

See "Adding filers" on page 134.

See "Add/Edit EMC Celerra filer options" on page 143.

About EMC Common Event Enabler (CEE)

The CEE framework provides a working environment for the following mechanisms:

- VNX/Celerra Antivirus Agent (CAVA)
- VNX/Celerra Event Publishing agent (CEPA)

Symantec Data Insight uses the CEPA functionality of the CEE framework to receive event notifications. The EMC VNX/Celerra Event Publishing Agent (CEPA) is a mechanism that enables Data Insight to register with the EMC VNX or Celerra filer to receive event notifications from the filer. You can specify filters for the event type, the CIFS server, and the shares that you want to monitor during registration with

the CEPA facility in the CEE framework. CEPA then sends notifications regarding the registered events to Data Insight.

Preparing the EMC filer for CEPA

The Symantec Data Insight server registers with the EMC Celerra or the EMC VNX filer through the EMC Common Event Enabler (CEE) framework. Data Insight uses this framework to receive notifications of file access events from the filer.

See "About EMC Common Event Enabler (CEE)" on page 104.

To configure the EMC Celerra or EMC VNX filer to send event information to Symantec Data Insight

1   Create a cepp.conf file on the EMC filer. The following is a sample of the code that the cepp.conf file must contain:

```
surveytime=90

pool name=matrixpool \

servers=<IP Address/Hostname of Windows server running the EMC CAVA
service> \

postevents=* \

option=ignore \

reqtimeout=500 \

retrytimeout=50
```

Note: If the server pool contains more than one server, each of the `server` entry should be separated by a "|".

2   Copy the cepp.conf file to the root directory of the Data Mover. Run the following command: `server_file <datamover_name> -put cepp.conf cepp.conf`

For example, `server_file server_2 -put /tmp/CEPA/cepp.conf cepp.conf`

3   Start the CEPP service on the filer. Run the following command:

`server_cepp <datamover_name> -service -start`

Ensure that the service has started by running the following command:

`server_cepp name of data mover -service -status`

Note: For detailed information about configuring CEPA, refer to the EMC documentation.

Preparing Symantec Data Insight to receive event notification

The EMC Common Event Enabler (CEE) can be installed on the same Windows server as the Data Insight Collector node or on a remote server in the same Active Directory domain.

You must perform the following steps to route events from the Windows server on which the EMC CEE is installed to the Collector node.

To prepare Data Insight to receive event notification

1   Provision a Windows 2003 or 2008 server to run the EMC CEE or framework in the same Active Directory domain as the filers you wish to monitor.

2   Open Windows' Registry Editor (**Start** > **Run** > **regedit**).

3   In Registry Editor, navigate to `HKEY_LOCAL_MACHINE > SOFTWARE > EMC > Celerra Event Enabler/VNX Event Enabler >CEPP > Audit >Configuration`.

4   Double-click **Endpoint**.

5   Modify the registry entry for the EMC CAVA service to allow access to the Data Insight Collector node. Depending on the type of your Data Insight deployment, there can be the following different scenarios:

   ■   The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is only being used by Data Insight. In this case, add the Data Insight key, `SymantecDataConnector`, to the **Endpoint** option.

   ■   The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key, `SymantecDataConnector`, to the **Endpoint** option. Each entry must be separated by a semi-colon.

   Note: The above-mentioned scenarios are automatically configured at the time adding filers.

   ■   The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is being used only by Data Insight. In this case, add the Data Insight key in the format,

> `SymantecDataConnector@<IP address of the Collector>`, to the
> **Endpoint** option.

- The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key in the format,`SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.

If the EMC CAVA service is installed on multiple machines, modify the registry entries on each of these machines.

6   Double-click the **Enabled** setting, and set its value to 1.

7   To start the EMC CAVA service, run the following command on the EMC filer to check the service status. For example,

```
Server_cepp  server_2 -pool -info
```

8   Install Data Insight Collector node.

9   Login to the Data Insight Management Console.

10   Navigate to **Settings** > **Data Insight Servers** to open the Data Insight Servers details page for the Collector.

11   Navigate to the service configuration section on the filer, and click **Enable** to start the DataInsightCelerra service on the Collector node.

12   Under Credentials, enter the credentials that the service needs to run as. The specified credentials must be that of a domain user.

13   Click **Configure** to apply these settings to the server and start the EMC CAVA service.

Credentials required for configuring EMC Celerra filers

lists the set of credentials that are required by Symantec Data Insight during system configuration.

Table 7-1   Credentials for configuring EMC filers

Credential	Details
Credentials required to configure DataInsightCelerra service. The DataInsightCelerra service runs on the Collector and processes events sent by the CAVA services using the Windows RPC interface. This service must be configured on each Collector node that is used to connect to the EMC filers.	Required by the DataInsightCelerra service to run and authenticate itself with the EMC CAVA service provided by EMC, which runs on the Data Insight Collector node. The credential should belong to the user in the domain of which the Data Insight Collector node and the EMC filer are part.
Credentials required during filer configuration through the Symantec Data Insight Management Console.	Required to discover shares for EMC filer. This credential belongs to the EMC filer Control Station user who has administrative rights including XMLAPI v2 privilege (for example, nasadmin). See "Preparing Symantec Data Insight to receive event notification" on page 106.
Credentials required for scanning of shares.	Required for scanning of shares from the EMC filer. This credential belongs to the user in the domain of which the EMC filer is a part. Additionally, to be able to obtain share-level ACLs, the credentials must belong to the Domain Administrators group on the file server. You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system heirarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions See "Configuring EMC filers" on page 104.

Configuring EMC Isilon monitoring

This chapter includes the following topics:

- Configuring EMC Isilon Filers
- Prerequisites for configuration of Isilon file server monitoring
- Credentials required for configuring an EMC Isilon cluster
- Configuring audit settings on EMC Isilon cluster using OneFS GUI console
- Configuring audit settings on EMC Isilon cluster using the OneFS CLI
- Configuring Isilon audit settings for performance improvement
- Preparing Symantec Data Insight to receive event notifications from an EMC Isilon cluster
- Creating a non-administrator user for an EMC Isilon cluster

Configuring EMC Isilon Filers

Symantec Data Insight supports the monitoring of EMC Isilon OneFS file servers, version 7.1 or later. Isilon OneFS is a cluster that has multiple hosts and a unified namespace-based file system across those hosts. Data Insight uses the CEE framework that is provided by EMC to fetch access events from EMC Isilon file servers. Ensure that CEE version 6.1 is used to enable Data Insight to fetch access event information from the filer.

Note: Symantec Data Insight also supports scanning and event monitoring of user's home directories on the EMC Isilon storage system.

Prerequisites for configuration of Isilon file server monitoring

Before you start configuring Isilon file servers, ensure the following pre-requisites are complete:

- The EMC Common Event Enabler (CEE) version 6.1 or later is installed. You can install CEE either on the same Windows server as the Data Insight Collector or on a remote server in the same directory service domain.

   Note: Symantec recommends you to use the latest version of CEE.

- The DataInsightCelerra service is installed on the Data Insight Collector.

- Microsoft .Net Framework version 3 or 3.5 is installed on your Collector node.

- A local or a domain user is configured on Isilon for use by Data Insight.

- The audit settings are configured on the Isilon cluster.
   See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 112.

- Note the port number from the URL used to access the Isilon OneFS Management Console. This port number is used by Data Insight for discovery purposes.

- The Windows firewall is disabled on the Data Insight Collector node and also on the server where EMC CEE is installed.

Credentials required for configuring an EMC Isilon cluster

Table 8-1 lists the set of credentials that are required by Symantec Data Insight during system configuration.

Table 8-1

Credential	Details
Credentials required to configure DataInsightCelerra service.	Required by the DataInsightCelerra service to run and authenticate itself with the Isilon cluster. The service runs on the Data Insight Collector node. The credential should belong to the user in the domain of which the Data Insight Collector node and the Isilon cluster are part.
Credentials required during filer configuration through the Data Insight Management Console.	Required to discover shares for the Isilon cluster. See "Preparing Symantec Data Insight to receive event notifications from an EMC Isilon cluster" on page 116.
Credentials required for scanning of shares	Required for scanning of shares from the Isilon cluster. This credential belongs to the user in the domain of which the Isilon is a part. Additionally, to be able to obtain share-level ACLs, the credentials must belong to the Domain Administrators group on the file server. You do not need this privilege if you do not want to get the share-level ACLs. In this case you will only need privileges to mount the share and scan the file system hierarchy. You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions See "Configuring EMC Isilon Filers" on page 109.

Configuring audit settings on EMC Isilon cluster using OneFS GUI console

To enable Data Insight to receive access event information from the Isilon cluster, you must first configure auditing on the Isilon cluster. You can configure the audit settings either from the Isilon OneFS web-based GUI console or from the Isilon command line interface.

See "Configuring audit settings on EMC Isilon cluster using the OneFS CLI" on page 113.

To configure audit settings on Isilon using OneFS console

1   On the OneFs WebUI, navigate to the **CLUSTER MANAGEMENT** > **Auditing**.

2   Under the **Edit Auditing Settings** section, select **Enable Protocol Access Auditing**.

3   Under the **Audited Zones** section, add the access zone that you want to audit. To enable auditing for the entire Isilon cluster, you can select the default System zone. For more information about access zones, see the EMC Isilon documentation.

4   Under the **Event Forwarding** section, enter the uniform resource identifier for the server where the Common Event Enabler is installed. The format of the entry is: http://*<IP of the CEE server>*:port/cee.

    For example: http://10.209.302.152:12228/cee.

    Note that 12228 is the default CEE HTTP listen port. You must choose a port number that is same as the one configured in the registry on the computer where CEE is installed.

    See "Preparing Symantec Data Insight to receive event notifications from an EMC Isilon cluster" on page 116.

5   Under the section, **Event Forwarding**, add the host name of the storage cluster. You can either use the EMC Isilon SmartConnect cluster name or the DNS resolvable host name of one of the nodes of the cluster. Do not leave this field blank. This host name is later used to add the EMC Isilon cluster to the Symantec Data Insight configuration.

See "Configuring Isilon audit settings for performance improvement" on page 115.

Configuring audit settings on EMC Isilon cluster using the OneFS CLI

You can configure the audit settings on an EMC Isilon cluster using the command-line interface(CLI). You must be a root user of the Isilon cluster perform the configuration steps.

To configure and view audit settings on Isilon using the OneFS CLI:

1   Log on to the Isilon OneFS cluster using the command line interface.

2   Issue the following commands:

- To enable auditing:

```
di-isilon-1# isi audit settings
modify--protocol-auditing-enabled on
```

- To disable auditing:

```
di-isilon-1# isi audit settings
modify--protocol-auditing-enabled off
```

- To configure the audit settings:

```
di-isilon-1# isi audit settings modify
--add-cee-server-uris=http://cee.example.com:12228/cee
--audited-zones=system --hostname=di-isilon.example.com
--config-auditing-enabled=no
```

- To view the audit settings:

```
di-isilon-1# isi audit settings view
```

Depending upon the configured audit settings, you may see the following response:

```
Protocol Auditing Enabled: Yes
            Audited Zones: System
          CEE Server URIs: http://cee.example.com:12228/cee
                 Hostname: di-isilon.example.com
  Config Auditing Enabled: No
```

After you have enabled audit settings on the Isilon cluster you must also perform the configuration for the Access Zones on the cluster.

To configure Access Zones using the OneFS CLI:

1 Log on to the Isilon OneFS cluster using the command line interface.

2 Issue the following command:

```
isi zone zones modify <zone>

--audit-success {close | create | delete | get_ security | logoff
| logon | read | rename | set_security | tree_ connect | write |
all} | --clear-audit-success

--add-audit-success {close | create | delete | get_security |
logoff | logon | read | rename | set_security | tree_ connect |
write | all}

--remove-audit-success <string>

--audit-failure {close | create | delete | get_security | logoff
| logon | read | rename | set_security | tree_ connect | write |
all} | --clear-audit-failure

--add-audit-failure {close | create | delete | get_security |
logoff | logon | read | rename
```

Using the command line interface you can also enable specific audit events.

To enable specific audit events using the OneFS CLI:

1 Log on to the Isilon OneFS cluster using the command line interface.

2 Issue the following command:

```
isi zone zones modify system --audit-success
create,delete,rename,set_security
```

Depending on your configuration, you may see the following response:

```
di-isilon-1# isi zone zones list -v
               Name: System
         Cache Size: 4.77M
       Map Untrusted:
          SMB Shares: -
       Auth Providers: -
       Local Provider: Yes
         NetBIOS Name:
       All SMB Shares: Yes
   All Auth Providers: Yes
    User Mapping Rules: -
Home Directory Umask: 0077
   Skeleton Directory: /usr/share/skel
        Audit Success: create, delete, rename, set_security
        Audit Failure: -
              Zone ID: 1
```

See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 112.

See "Configuring Isilon audit settings for performance improvement" on page 115.

Configuring Isilon audit settings for performance improvement

Configuring audit setting by using the Isilon OneFS web-based GUI console by default enables auditing of the following event types:

- audit-success

- audit-failure

- create

- delete

- rename

- set_security

The number of different types of events that Data Insight monitors affects the system performance. To reduce performance overhead, you may disable auditing of audit-failure events. You must use the OneFS CLI to disable auditing

To disable auditing for audit-failure events:

1   Log on to the Isilon OneFS cluster using the command line interface.

2   Issue the command:

```
isi zone zones modify system --remove-audit-failure all
```

Preparing Symantec Data Insight to receive event notifications from an EMC Isilon cluster

EMC Common Event Enabler (CEE) version 6.1 or later is required to receive events from EMC Isilon OneFS cluster. You can install CEE either on the same Windows server as the Data Insight Collector node or on a remote server in the same directory service domain.

To prepare Data Insight to receive event notification

1   Install EMC CEE framework on a Windows server so that it runs in the same directory service domain as the filers you want to monitor.

2   Open Windows Registry Editor (**Start** > **Run** > **regedit**).

3   Modify the registry entry for the EMC CAVA service to allow access to the Data Insight Collector node. Depending on the type of your Data Insight deployment, there can be the following different scenarios:

   - The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is only being used by Data Insight. In this case, add the Data Insight key, SymantecDataConnector, to the **Endpoint** option.

   - The EMC CAVA service and the Collector node are running on the same machine, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key, SymantecDataConnector, to the **Endpoint** option. Each entry must be separated by a semi-colon.

   Note: The above-mentioned scenarios are automatically configured at the time adding filers.

- ■ The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is being used only by Data Insight. In this case, add the Data Insight key in the format, `SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.

- ■ The EMC CAVA service and the Collector node are running on separate machines, and the EMC CAVA service is also being used by applications other than Data Insight. In this case, append the Data Insight key in the format,`SymantecDataConnector@<IP address of the Collector>`, to the **Endpoint** option.

  If the EMC CAVA service is installed on multiple machines, modify the registry entries on each of these machines.

4   Configure the following registry values:

- ■ Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **Configuration**. Note the key called **HttpPort**. The default key value is specified as `12228`. You can modify this key value but you must ensure that this value matches the port in the CEE URL that is specified during Isilon audit configuration.
  See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 112.

- ■ Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **CEPP** > **Audit** > **Configurtion**. Set the value of key **Enabled** to `1`.

- ■ Navigate to the location: **HKEY_LOCAL_MACHINE** > **SOFTWARE** > **EMC** > **CEE** > **CEPP** > **Audit** > **Configuration**. Set the value of key **EndPointi** to `SymantecDataConnector@<IP of Data Insight Collector machine>`

5   Log on to the Data Insight Management Console.

   Navigate to **Settings** > **Data Insight Servers** to open the Data Insight Servers details page for the Collector.

6   Navigate to the service configuration section on the filer, and click **Enable** to start the DataInsightCelerra service on the Collector node.

7   Under **Credentials**, enter the credentials that the service needs to run as. The specified credentials must be that of a domain user.

8   Click **Configure** to apply these settings to the server.

9   Restart the EMC CAVA and EMC CEE Monitor service on the Collector machine or on the remote CAVA server if the services are remotely installed.

Creating a non-administrator user for an EMC Isilon cluster

Data Insight requires a user account on Isilon to perform automatic discovery of CIFS shares and to list all local groups, group memberships, and local users. Data Insight can use a non-administrator account for this purpose and the account can be a local Isilon OneFS account or a domain account.

To configure a domain user for discovery and scanning of CIFS shares

1   Log on to the Isilon cluster CLI using SSH or Telnet as an Isilon administrator.

2   Run the following commands:

   ▪   To create a role named dirole:

```
isi auth roles create --name dirole --description "Read-only
role for DataInsight"
```

   ▪   To give the user, the privileges to log on to the REST API platform framework, to get a list of CIFS shares and to list users and groups:

```
isi auth roles modify dirole --add-user=username@domain
--add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI
--add-priv-ro=ISI_PRIV_AUTH
```

   ▪   To add the user to the Backup Operators group on Isilon, which enables Data Insight to scan all the CIFS shares:

```
isi auth groups modify "Backup Operators" --add-user
username@domain
```

To configure a local user for discovery of CIFS shares

1   Log on to the Isilon cluster CLI using SSH or Telnet as an Isilon administrator.

2   Run the following commands:

   ▪   To create a new local user called diuser:

```
isi auth users create diuser --enabled yes --password xxxxxx
```

   ▪   To create a role named dirole:

```
isi auth roles create --name dirole --description "Read-only
role for DataInsight"
```

   ▪   To grant the user, the privileges to log on to the REST API platform framework, to get a list of CIFS shares and to list users and groups:

```
isi auth roles modify dirole --add-user=diuser
--add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI
--add-priv-ro=ISI_PRIV_AUTH
```

Configuring Windows File Server monitoring

This chapter includes the following topics:

- Configuring Windows File Servers
- Credentials required for configuring Windows File Servers
- Upgrading the Windows File Server agent

Configuring Windows File Servers

Data Insight uses an agent to collect access events from the Windows File Server. The agent resides on the file server. Before you can configure a Windows File Server, you must install the Data Insight agent on the filer. The Data Insight agent consists of a filter driver that monitors the file system and records events that are relevant for Data Insight. It also consists of the Data InsightWinNAS service, which receives the event information from the filter driver and transfers it to the collector node that is configured for that filer.

If you do not want Data Insight to access events for a Windows File Server, it is possible to configure Windows File Server without an agent. In this case Data Insight scans shares of the filer from the Collector.

You can choose to install the agent on the Windows File Server automatically when adding the filer, or manually. Before you can install the agent automatically, ensure that the communication service port 8383 on the Collector node is accessible from the Windows File Server.

For detailed information about installing the agent manually, see the *Symantec Data Insight Installation Guide*.

Note: All Data Insight worker nodes must be at the same level of major version as the Management Server. Windows File Server agents can be one level lower than the Management Server version. Thus, Management Server 4.5 is compatible with both 3.0RU1 (3.0.1) version as well as 4.5 of Windows File Server agents. This gives you enough time to plan out the upgrade of your Windows File Server agents.

You can either add a Windows File Server to Data Insight through the Management Console, or if you want to add multiple filers together, you can use the `installcli.exe` utility. The `installcli.exe` utility uses a `.csv` file with the following details as input:

- The host name or IP address of the Windows File Servers that you want Data Insight to monitor.

- The host name, IP address, or ID of the Collector node that is configured to scan the filer.

- The host name, IP address, or ID of the Indexer node that is configured for the filer.

- The credentials that Data Insight should use to install the agent on the Windows File Server. The credential should be in the format user@domain. `installcli.exe` also accepts LocalSystem credentials as value _LOCAL_. The same credentials must be added to Data Insight as a saved credential previously.

- True or false value indicating if the filer is clustered.

- The IP addresses of the agents. Separate multiple IP addresses with a semi-colon. If you do not want to use an agent to monitor the filer, indicate this with a hyphen (-).

- The credentials required to scan the filer. The credential should be in the format user@domain. The same credentials should be added to Data Insight as a saved credential previously.

- True or false value indicating whether the scan should be enabled according to the specified schedule.

- In case of a Windows File Server agent upgrade, RP or Full value indicating the type of upgrade you want to perform. This parameter is optional.
  Optionally, the name of the installer. If the name of the installer is not specified, an appropriate installer is picked from `installers` folder on the Collector.

- True or false value indicating whether event monitoring should be enabled. For example, *winnas.company.com,collector.company.com,indexer.company.com, Administrator@DOMAIN,FALSE,winnas.company.com, Administrator@DOMAIN,TRUE,TRUE,RP, Symantec_DataInsight_windows_winnas_4_0_0_3002_x64.exe*

See "Credentials required for configuring Windows File Servers" on page 121.

To add multiple Windows File Servers

1 Log in to the Data Insight Management Server.

2 Open a Windows command prompt and change to the `INSTALL_DIR\bin` directory, where `install_dir\bin` is the installation path for Symantec Data Insight.

3 Type the following command:

```
installcli -w <Path to .csv file with Windows File Server
specifications>
```

For detailed information on `installcli.exe`, see the *Command File Reference*.

You can also add a clustered Windows File Server to Data Insight. Data Insight supports only a Microsoft Cluster Server (MSCS) configuration.

In this release, the following limitations exists for Windows File Servers:

■ Windows filter driver does not capture IP address from which accesses are made.

See "Configuring a DFS target " on page 166. for details about configuring a DFS target.

See "Using the Upload Manager utility" on page 218.

See "Adding filers" on page 134.

See "Add/Edit Windows File Server options " on page 147.

Credentials required for configuring Windows File Servers

Table 9-1 lists the set of credentials that are required by Symantec Data Insight during system configuration.

Table 9-1   Credentials for configuring Windows File Servers

Credential	Details
Credentials required to install agent on the Windows File Server.	This credential belongs to a user in the Administrators group on the Windows File Server.

The credential is also used to discover shares and obtain storage utilization information from the filer. |

Table 9-1   Credentials for configuring Windows File Servers *(continued)*

Credential	Details
Credentials required to discover shares and obtain storage utilization information on the filer.	Required for monitoring shares or when configuring a Windows File Server cluster. This credential belongs to a user in the Administrators group on the file server. If your configuration is not a Windows cluster or you do not want to collect storage utilization information for the filer, a credential with the privilege to list shares on the filer is sufficient.
Credentials required for scanning shares on the Windows File Server.	Required to scan a share. This credential belongs to a user with necessary share-level permissions and file system ACLs on a Windows File Server share. To be able to obtain share-level ACLs, the credentials must belong to the Power Users or Administrators group on the Windows File Server. You do not need this privilege if you do not want to get the share-level ACLs. To be able to scan a Windows File Server share successfully, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs: ■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions **Note:** To enable Data Insight to successfully scan the shares on a clustered Windows File Server, you must ensure that the scanning user has domain level permissions of **Allow logon locally**.

Note: If you neither want Data Insight to install an agent automatically, nor do you want Data Insight to discover shares on the cluster or get storage utilization information, specifying the filer credentials is optional.

Upgrading the Windows File Server agent

You can upgrade the Windows File Server agent automatically from the Data Insight Management Console.

Note: The option to upgrade the agent automatically appears only if you have configured the Windows File Server to allow Data Insight to automatically install the agent.

To upgrade Windows File Server agent automatically

1   Log on to the Management Console as Administrator.

2   Use the Uploader Manager utility to upload the agent packages on Collector worker nodes corresponding to the Windows File Server agent.

    See "Using the Upload Manager utility" on page 218..

3   Select **Settings** > **Filers** to view the list of configured Windows File Servers.

4   Click the server on which you want the upgrade the agent.

5   On the configuration details page, click **Upgrade Agent**

6   Windows File Server agent upgrade window appears and displays a progress bar while upgrading.

7   Click **Finish** to exit setup.

Note: To upgrade the Windows File Server agent manually, see the *Symantec Data Insight Installation Guide*. You can upgrade multiple Windows File Server agents using the `installcli` utility. See "Configuring Windows File Servers " on page 119.

Configuring Veritas File System (VxFS) file server monitoring

This chapter includes the following topics:

- About configuring Veritas File System (VxFS) file servers

- Credentials required for configuring Veritas File System (VxFS) servers

- Enabling export of UNIX/Linux NFS shares on VxFS filers

About configuring Veritas File System (VxFS) file servers

A Data Insight agent plugin, `vxdiplugind`, is used to monitor access events on the VxFs file servers. The plugin is part of the VxFS package and is automatically installed on the file server when Veritas Storage Foundation is installed. The plug-in captures events from the VxFS filer that Data Insight is monitoring, and saves it to a temporary database.The event data is then pulled by Data Insight, which fetches the access event information through Veritas Operations Manager (VOM) to gain vital insight into the user activity on the filer.

Data Insight uses NFS to scan all or a portion of VxFS shares remotely from the Collector node. Data Insight only monitors the access events on the VxFS devices exported by NFS.

Before you start configuring VxFS filers, verify the following:

- The file server must be installed with Storage Foundation 6.0.1

- The file server must be installed with Veritas Operations Manager (VOM) 4.1.

- NFS version 3.0 is configured on the VxFS filer.

- The LDAP or NIS domains that your users are part of must be configured in Data Insight.

- The Collector node for the VxFs filer must be a Windows 2008 Enterprise server. Ensure that the Collector node monitoring the VxFS filer has services for NFS enabled as file server roles. You can install a role on Windows 2008 Enterprise server through the **Server Manager** > **Add roles** option.

- The filer is accessible from the Collector node using the host name or IP address you plan to use when adding the filer.

You can also add a clustered VxFS file server to Data Insight. Data Insight supports only a Veritas Cluster Server (VCS) configuration for VxFS file servers configured in failover mode. Parallel Clustered File System is not supported in this release.

See "Adding filers" on page 134.

See "Enabling export of UNIX/Linux NFS shares on VxFS filers" on page 127.

See "Add/Edit Veritas File System server options" on page 151.

Credentials required for configuring Veritas File System (VxFS) servers

Table 10-1 lists the set of credentials that are required by Symantec Data Insight during system configuration.

Table 10-1   Credentials required for configuring VxFS filers

Credentials	Details
Credentials required during filer configuration through the Symantec Data Insight Management Console.	Required to discover shares on the VxFs filer. This credentials belongs to a user on the UNIX server who has administrative rights on the VxFS filer (for example, root). The credential should belong to a root user on the VxFS filer.
	Optionally, this credential can also belong to a local user who has access to the Data Insight namespace in the Veritas Operations Manager (VOM) agent installed on the VxFS filer.
	To configure a user other than the root user, you must create or use an existing user account, which you can use to add the filer into the Data Insight namespace. To add a local user account under VOM:
	1 Log in as root on the VxFs filer.
	2 Change directory to `/opt/VRTSsfmh/di/web/admin.`
	3 Create a `.xprtlaccess` file, and add the user to that file. For example, add *vomuser@unixpwd:user*, where vomuser is the name of the local user account.

Table 10-1   Credentials required for configuring VxFS filers *(continued)*

Credentials	Details
Credentials required for scanning on VxFS filer server	Required for scanning of shares from the VxFS filer.
	For scanning NFS shares, Data Insight needs a Unix account with at least read and execute permissions on all folders, alongwith at least read permission on all files. By default, Data Insight uses the User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Data Insight Servers** > **Advanced Settings** section of the Collector node.
	See "Configuring advanced settings" on page 203.
	Additionally, you must also have share-level READ permissions on the NFS export.

Enabling export of UNIX/Linux NFS shares on VxFS filers

These instructions are for Red Hat Enterprise Linux operation system which has standalone Storage Foundation 6.0 installed and a file system created using VxFS. The steps will change depending upon other operating system flavors.

To enable export of NFS shares on VxFS filers

1 Login as root on the VxFS filer and open the `/etc/exports` file.

2 Specify the name of the share that you would like to monitor. For example, `/demoshare`, where the VxFS file system is mounted.

Ensure that the device entries are added in `/etc/fstab` to automatically mount NFS file systems after reboot.

Data Insight uses /etc/exports and /etc/fstab for NFS share discovery. Sample entries are shown below:

```
root@RHEL5-VxFS ~]# cat /etc/fstab | grep vxfs

/dev/vx/dsk/openldapdg/vol01 /openldaphome vxfs defaults,_netdev 0 0

/dev/vx/dsk/openldapdg/vol02 /data vxfs defaults,_netdev 0 0

/dev/vx/dsk/openldapdg/vol03 /didata vxfs defaults,_netdev 0 0

[root@RHEL5-VxFS ~]# cat /etc/exports

/openldaphome 192.168.0.10(ro,sync,no_root_squash)  192.168.0.11
(rw,syc) /data/exportshare *(rw,sync,no_root_squash)

/didata *(rw,sync,no_root_squash)
```

3 Specify the root access and read only access to Data Insight Collector node. For example,

```
/demoshare <Collector node IP> (ro,sync,no_root_squash)

ro:read only

no_root_squash: root access.
```

You can specify read `+write`, `root_squash`, `anonuid`, `anongid` or other settings, as required.

4 Run the following command to start the NFS daemon

```
#service nfs start
```

See "Adding filers" on page 134.

Configuring monitoring of a generic device

This chapter includes the following topics:

- About configuring a generic device
- Credentials required for scanning a generic device

About configuring a generic device

Data Insight supports scanning and auditing of access events on storage devices with varied file systems. You can configure Data Insight to monitor generic device filers in addtion to NetApp, EMC Celerra, Veritas File System (VxFS), Windows File Server.

Data Insight uses a web API to collect access event files from the generic device filer. The web API enables web clients to push events from the generic device filers to the DataInsightGenericCollector web service. The DataInsightGenericCollector web service collects incoming events and copies them to a specific location on the Collector worker node configured for the storage device. The events are later processed by the corresponding Indexer.

The DataInsightGenericCollector service uses one-way Secure Sockets Layer (SSL) to secure communication between the generic device filer and the Collector node. The client connects to the DataInsightGenericCollector service through a specific port, which is configurable.

You must develop a customized client using the API specification provided by Data Insight to add a filer as a generic device to Data Insight, and to start monitoring, auditing, and scanning events on the filer. For more information on the web API specification for the generic Collector service, refer to the *Data Insight Programmer's Reference Guide*.

Credentials required for scanning a generic device

Table 11-1 lists the set of credentials that are required by Symantec Data Insight to scan a generic storage device.

Table 11-1   Credentials for scanning a generic device

Credential	Details
Credentials required for scanning of shares.	Required for scanning of shares from the filer.
	When scanning CIFS shares, this credential belongs to the user in the domain of which the filer a part. While the exact set of permissions depends on the generic device being scanned, this user must generally belong to the Administrator's group on the device. If the credential is not part of Administrator's group, the scanner might not be able to get share-level ACLs for shares of this device.
	Typically, to scan a CIFS sharre, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs enabled for the scan credential:
	■ Traverse Folder/Execute File ■ List Folder/Read Data ■ Read Attributes ■ Read Extended Attributes ■ Read Permissions
	For scanning NFS shares, Data Insight needs a Unix account with at least read and execute permissions on all folders, alongwith at least read permission on all files. By default, Data Insight uses User ID or Group ID 0 to scan NFS shares. You can configure an alternate User ID or Group ID from the **Settings** > **Advanced Settings** section of the Collector node.
	See "Configuring advanced settings" on page 203.
	When monitoring only NFS shares, you can specify **Use Local System account** from the scanning credentials drop-down, else you can specify credentials required to scan CIFS shares.

Managing file servers

This chapter includes the following topics:

About configuring filers

Symantec Data Insight collects and stores access events from NAS devices to service queries on user activity and data accesses. Before Data Insight can start collecting events, you must ensure that auditing is configured properly on the storage device. Data Insight collects access events using asynchronous APIs, namely, Fpolicy for NetApp filers, the CEE framework for EMC Celerra filers, and file system filter drivers for Windows File Server. Additionally, other generic devices can also publish events to Data Insight using the web API.

See "Managing Data Insight product servers" on page 195.

See "About configuring a clustered NetApp file server" on page 96.

See "Configuring EMC Isilon Filers" on page 109.

Viewing configured filers

In the Management Console, you can view all the filers that Data Insight is configured to monitor.

Use the provided dynamic search filter to search for configured filers based on various pre-defined criteria, for example, the type of the filer. You can also use the **Filter** field at the top of the content pane to filter the list of filers based on the IP address or hostname of the filer in addition to the pre-defined filter criteria. The displayed list of filers changes automatically when you select the check box for a filter criteria. For instance, when you select NetApp in the **Device Type** category, the application displays a list of configured NetApp devices. Similarly, when you select a Collector node in the **By Collector** filter, Data Insight displays a list of filers associated with the selected Collector node.

To view configured filers

1   In the Console, click **Settings** > **Filers**.

The screen displays the list of configured filers

2   Review the following information about the filers:

  ■   The object ID of the filer. This numerical value is used to identify the filer when troubleshooting issues with the filer. This column is hidden by default. To view this column, click on the column header and select **Columns** > **ID**.

  ■   The name of the filer.

  ■   The number of shares monitored by the filer.

  ■   The health of the filer.

- The type of filer -NetApp, EMC Celerra, Windows File Server, Veritas File System (VxFS) server, or a generic device.

- Whether file system event monitoring is enabled.

- The Collector node for the filer.

- The Indexer node for the filer.

- The scanning schedule for the filer. This column is hidden by default.

To review filer details

1   In the Console, click **Settings** > **Filers**.

2   Click the filer that you want to review, or click the **Select Action** drop-down and select **View**.

    The filer details screen appears.

To view filer events

1   In the Management Console, click **Settings** > **Filers**.

2   Click the **Select Action** drop-down for the corresponding server in the filers listing table, and select **Event Log**.

    The event log for that filer appears.

3   To download Data Insight logs for the filer for troubleshooting purposes, click the **Select Action** drop-down for the corresponding filer, and select **Download Logs**.

    Data Insight downloads a compressed folder containing the logs related to this filer from all relevant Data Insight servers.

    See "Downloading Data Insight logs" on page 282.

Adding filers

You must add filers that you want Symantec Data Insight to monitor.

To add filers

1 In the Console, click **Settings** > **Filers**.

The Filers page displays the list of available filers.

2 On the Filers page, click the **Add New Filer** drop-down, and select the type of filer you want to add.

3 On the New Filer screen, enter the filer properties, and click **Add New Filer**.

If you are adding a Windows File Server, Data Insight can automatically install an agent on the filer. This agent enables Data Insight to receive event notifications from the filer.

For detailed information about installing the agent manually, see the *Symantec Data Insight Installation Guide*.

See "About configuring filers" on page 133.

Add/Edit NetApp filer options

Use this dialog box to add a new NetApp filer to Symantec Data Insight or to edit the configuration of an existing filer.

Table 12-1   Add/Edit NetApp filer options

Field	Description
Filer host name or IP address	Enter the host name or IP address of the filer that you want Data Insight to monitor.
	Note: The host name or IP address should be the same as the filer name is entered in Symantec Data Loss Prevention targets.
Collector	From the drop-down, select the Collector worker node that is configured to scan the filer.
	Data Insight connects to the filer from this server. Symantec recommends that the Collector worker node share a fast network with the filer.
	Note: When monitoring NFS shares, ensure that the Collector node monitoring the filer must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through the **Server Manager** > **Add roles** option.

Table 12-1   Add/Edit NetApp filer options *(continued)*

Field	Description
Indexer	From the drop-down, select the Indexer worker node that is configured for the filer. Events and metadata collected from the filer is processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Filer administrator credentials	See "Credentials required for configuring NetApp filers" on page 79. Specifying the filer administrator credentials is optional, if you choose to not monitor events on the filer, nor enable share discovery.
Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of specified credentials. Symantec recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.
Filer is vFiler	Select the check box to indicate that this filer is a NetApp virtual file server.
Physical filer for vFiler	The host name or IP address of the physical NetApp file server that is associated with the virtual file server. If the Data Insight FPolicy safeguard is not enabled for the virtual file server, the field is not editable. See "Configuring scanning and event monitoring " on page 35.
Enable CIFS monitoring	Select this check box to enable monitoring of CIFS shares.
Enable NFS monitoring	Select this check box to enable monitoring of NFS shares.

Table 12-1   Add/Edit NetApp filer options *(continued)*

Field	Description
Select domain	From the drop-down, select the domain to which the NetApp filer belongs.
	This option is enabled when you enable the monitoring of NFS shares.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to allow Data Insight automatically discover shares of the filer and add them configuration.
	Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M.
	You can also choose to add shares manually.
	See "Adding shares" on page 160.
Exclude shares from discovery	Enter the details of shares which should not be included during discovery.
	This option is available if you select **Automatically discover all shares on this filer.** Specify comma-separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$`, and others.
Enable storage utilization analytics	Select the check box to allow Data Insight to gather storage utilization information from the filer. This information is used when you generate Filer Utilization and Filer Growth Trend reports.
	The DataInsightFpolicy service running on the Collector node gathers information about storage utilization on the filer.
Enable file system event monitoring	Select to enable event monitoring on the filer.

Table 12-1 Add/Edit NetApp filer options *(continued)*

Field	Description
Enable FPolicy automatically	Select to automatically enable FPolicy on the filer.
	If you clear this check box, you must manually enable FPolicy on the filer.
	See "Preparing the NetApp filer for Fpolicy" on page 85.
Register for explicit Read events	Select the option to register for explicit Read events.
	When this option is not selected, OPEN events are treated as READ events.
	Note: NFSv3 does not support OPEN events. This means that you will not see READ events for NFS shares when this check box is cleared.
	Symantec recommends that you do not register for explicit Read events. This can increase the load on the filer during peak traffic from third party applications such as backups over CIFS.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer:
	■ Use Collector's default scanning schedule.
	■ Use custom schedule.
	See "Custom schedule options" on page 157.
	Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.
Scanner credentials	See "Credentials required for configuring NetApp filers" on page 79.

Table 12-1   Add/Edit NetApp filer options *(continued)*

Field	Description
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.

See "Enabling export of NFS shares on a NetApp file server" on page 93.

Add/Edit NetApp cluster file server options

Use this dialog box to add a new NetApp filer in cluster mode to Symantec Data Insight or to edit the configuration of an existing filer.

Table 12-2   Add/Edit NetApp cluster file server options

Field	Description
Cluster Management Host	Enter the host name or IP address NetApp Cluster Management host interface that is used to manage the nodes in the cluster.
Collector	From the drop-down, select the Collector worker node that is configured to scan the CIFS servers that are configured in the cluster.
	Data Insight connects to the NetApp Cluster Management host from this server. Symantec recommends that the Collector worker node share a fast network with the cluster management host.
Indexer	From the drop-down, select the Indexer worker node that is configured for the filer.
	Events and metadata that are collected from the cluster are processed and stored on the Indexer node.
	You can also migrate the file server to another Indexer.

Table 12-2   Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Cluster Management Interface credentials	Data Insight uses the credentials that you specify to discover the following: ■ The Vservers in the cluster and the CIFS server underneath each Vserver. ■ The CIFS shares for each of the CIFS servers. To enable Data Insight to successfully discover shares, ensure that HTTP port 80 is not blocked between Data Insight Collector and NetApp cluster management host. Specifying the filer administrator credentials is optional, if you choose to not monitor events on the filer, nor enable share discovery. See "Credentials required for configuring a clustered NetApp file server" on page 99.
Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of specified credentials. Symantec recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the cluster management host. By default, Data Insight does not test credentials for the following HOMEDIR shares: ■ ~ ■ ~CIFS.HOMEDIR ■ CIFS.HOMEDIR ■ %w ■ %d

Table 12-2   Add/Edit NetApp cluster file server options *(continued)*

Field	Description
CIFS server	Every Vserver node in the cluster has a CIFS server configured on it. The CIFS server represents a file server for Data Insight. Data Insight automatically discovers all the CIFS servers that are configured in the cluster. From the drop-down, select the CIFS server that you want Data Insight to monitor. Ensure that you can resolve the CIFS server host name from the Collector node.
Enable CIFS monitoring	Select this check box to enable monitoring of CIFS shares.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to allowData Insight to automatically discover shares of the filer and add them configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M. You can also choose to add shares manually. See "Adding shares" on page 160.
Exclude shares from discovery	Enter the details of shares which should not be included during discovery. This option is available if you select **Automatically discover all shares on this filer.** Specify comma-separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$`, and others.

Table 12-2   Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Enable storage utilization analytics	Select the check box to allow Data Insight to gather storage utilization information from the filer. This information is used when you generate Filer Utilization and Filer Growth Trend reports.
	The DataInsightFpolicyCmod service running on the Collector node gathers information about storage utilization on the filer.
	See "About configuring FPolicy in Cluster-Mode" on page 97.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable FPolicy automatically	Select to automatically enable FPolicy on the filer.
	Once you enable FPolicy, the ONTAP Vserver node server initiates a TCP connection to the Data Insight FPolicy server. Ensure that the communication port in the firewall between the FPolicy server and the cluster management host is open.
	If you clear this check box, you must manually enable FPolicy on the filer.
	See "Preparing the NetApp filer for Fpolicy" on page 85.
Register for permission change events	Select if you want Data Insight to monitor the changes to permissions in your storage environment.
	By default, this option is not selected because it has a significant impact on the performace of the file server.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.

Table 12-2   Add/Edit NetApp cluster file server options *(continued)*

Field	Description
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 157. Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 P.M. on the last Friday of each month.
Scanner credentials	See "Credentials required for configuring a clustered NetApp file server" on page 99.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.

Add/Edit EMC Celerra filer options

Use this dialog box to add a new EMC Celerra filer to Symantec Data Insight or to edit the configuration of an existing filer.

Table 12-3   Add/Edit EMC Celerra filer options

Field	Description
CIFS Server Name	Enter the host name of the CIFS server that is exported by the filer. Entering the IP address of the CIFS server is not permitted
Control Station Hostname/IP address	Enter the IP address of the filer's Control Station.

Table 12-3   Add/Edit EMC Celerra filer options *(continued)*

Field	Description
Collector	From the drop-down, select the collector worker node that is configured to scan the filer. Data Insight connects to the filer from this server. Symantec recommends that the Collector worker node share a fast network with the filer You can also migrate the file server to another Indexer.
Indexer	From the drop-down, select the Indexer worker node that is configured for the filer.
Control Station Credentials	Enter the credentials for the filer's Control Station. These credentials are used to discover shares on the filer and add them to the configuration. You can specify non-administrative credentials, however, Test Connection will fail. In this case, you can continue to add the filer, but you must add shares manually.
Virtual Data Mover	Select the check box if the filer is running a virtual data mover. This field is used to handle physical paths that are returned for virtual data movers.
Test credentials	Click to test the availability of network connection between the Collector worker node and the control station and the validity of the specified credentials. Symantec recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer. **Note:** You must skip testing of credentials if you choose to use a non-administrator user. When you click Test Credentials, Data Insight by default tries to discover shares on the filer using the Control Station credentials.

Table 12-3   Add/Edit EMC Celerra filer options *(continued)*

Field	Description
Monitoring details	Select **Automatically discover and monitor shares on this filer** to enable Data Insight to automatically discover shares of the filer and add them configuration.
	You can also choose to add shares manually.
	Clear the check box if you use Control Station credentials with insufficient privileges for share discovery. If you choose to use credentials that do not have administrator rights and XML v2 privilege, you must manually add shares to the configuration.
	Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 A.M. and 2:00 P.M.
	See "Adding shares" on page 160.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer:
	▪ **Use the Collector's scanning schedule**
	▪ **Use custom schedule** See "Custom schedule options" on page 157.
	From the drop-down, select the appropriate frequency option. Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares on the last Friday of each month.
	Note: You can also customize the schedule per share using the Add/Edit Share dialog box.
Scanner credentials	See "Credentials required for configuring EMC Celerra filers" on page 107.

Table 12-3   Add/Edit EMC Celerra filer options *(continued)*

Field	Description
Scan new share immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule.

Scanning will still run only when scanning is permitted on the Collector node. |

Add/Edit EMC Isilon file server options

Table 12-4

Field	Description
Cluster Management Host	Enter the host name for the Isilon cluster. It can be EMC Isilon SmartConnect Cluster name or it can be the DNS resolvable host name of one of the hosts of the cluster.

Note: The Cluster Management Host name is the same host name which is entered during the configuration of audit settings on the Isilon cluster. See "Configuring audit settings on EMC Isilon cluster using OneFS GUI console" on page 112. |
| Cluster Management Port | Enter the port number on Isilon to be used for discovery.

Note: This is the port which is part of the URL used to access EMC Isilon OneFS web-based management console. |
| Collector | From the drop-down, select the collector worker node that is configured to scan the filer.

Data Insight connects to the filer from this server. Symantec recommends that the Collector worker node share a fast network with the filer |
| Indexer | From the drop-down, select the Indexer worker node that is configured for the filer.

You can also migrate the file server to another Indexer. |

Table 12-4   *(continued)*

Field	Description
Cluster Management Host Credentials	Select the saved credentials from the drop-down list to access the Cluster Management Host.
Test credentials	Click to test the availability of network connection between the Collector worker node and the Isilon cluster.
Monitoring Details	Select **Automatically discover and monitor shares on this filer** to enable Data Insight to discover automatically, the shares of the filer and add them configuration.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning Schedule(Full Scan)	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's default scanning schedule** ■ **Use custom schedule** From the drop-down, select the appropriate frequency option.
Scanner credentials	See "Credentials required for configuring an EMC Isilon cluster" on page 110.
Scan newly added shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Note that a scan can run only when scanning is permitted on the Collector node.

Add/Edit Windows File Server options

Use this dialog box to add a new Windows File Server to Symantec Data Insight or to edit the configuration of an existing filer.

Table 12-5 Add/Edit Windows File Server options

Field	Description
Is a MSCS clustered file server	Select the check box if the Windows File Server is part of a Microsoft Cluster Server configuration.
Windows server name/Cluster name	Enter the host name or IP address of the filer that you want Data Insight to monitor.
	In case of a clustered Windows File Server, enter the host name or IP address of the cluster.
	Note: The hostname or IP address should be same as the filer name entered in Symantec Data Loss Prevention Discover targets.
Select Collector node for this filer	From the drop-down, select the collector worker node configured to scan the filer.
	Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer.
Select Indexer node for this filer	From the drop-down, select the Indexer worker node configured for the filer.
Monitor mode	Select one of the following monitoring options:
	■ **Monitor this filer using an agent** If you select this option, Data Insight is able to monitor all file system events on the filer and scan file system metadata.
	■ **Monitor this filer without an agent** If you select this option, Data Insight scans the filer using CIFS to discover shares and obtain file metadata. However, in this case, Data Insight will not be able to monitor file system events.

Table 12-5   Add/Edit Windows File Server options *(continued)*

Field	Description
Agent names for this filer	This option is visible when adding a clustered file server that is monitored using an agent, but where the agent is installed manually.
	Select one or more agent nodes from the list that belong to this cluster.
	This option is also visible when editing a clustered file server.
Let Data Insight install the agent automatically	Select to allow Data Insight to install or upgrade the agent on the Windows File Server.
	Data Insight automatically installs the Windows File Server agent on the filer using the WMI interface and also registers the filer with the Management Server.
Node names to install agent	This option is only visible if you have selected **Is a MSCS clustered file server**.
	In the text box, enter comma-separated IP addresses or hostnames of the Windows File Server nodes, on which you want to install the agent.
Filer Administrator Credentials	Enter the credentials that Data Insightshould use to install the agent on the Windows File Server.
	See "Credentials required for configuring Windows File Servers" on page 121.
Test Connection	Click to test the availability of network connection between the Collector worker node and the filer, and the validity of the specified credentials.
	Symantec recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.

Table 12-5   Add/Edit Windows File Server options *(continued)*

Field	Description
Automatically discover and monitor all shares on this filer	Use this option to have Data Insight automatically discover shares of the filer and add them configuration. You can choose to exclude certain shares using the **Exclude shares** field. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 a.m. and 2:00 p.m.
Exclude following shares from discovery	Enter the details of shares which should not be included in share discovery. This option is available if you select **Automatically discover all shares on this filer**. Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, `tmp*` ignores shares `tmp_A`, `tmp_abc`, `*$` ignores shares `C$`, `EXT$` and others.
Collect storage utilization information for the filer	Select to enable Data Insight to collect storage utilization information from the filer. This information is used to create Filer utilization and Filer Growth Trend reports.
Enable file system event monitoring	Select to enable event monitoring on the filer.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ▪ **Use the Collector's scanning schedule** ▪ **Define custom schedule** Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares on the last Friday of each month. **Note:** You can also customize the schedule per share using the **Add/Edit Share** dialog box.

Table 12-5 Add/Edit Windows File Server options *(continued)*

Field	Description
Scanner credentials	Select one of the following: ■ **Use LOCAL SERVICE credentials** Select to use the LOCAL SERVICE account to scan shares of the filer. This option is available only for the filers monitored using an agent. If you select this option, ensure that the LOCAL SYSTEM account has appropriate privileges to scan the shares. If the account does not have adequate privileges, the scans for such shares will fail if performed using this account. ■ **Use saved credentials** Select the saved credentials from the drop-down or specify new credentials. See "Credentials required for configuring Windows File Servers" on page 121.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning will still take place during the hours when scanning is permitted on the Collector node.

Add/Edit Veritas File System server options

Use this dialog box to add a new Veritas File System (VxFS) filer to Symantec Data Insight or to edit the configuration of an existing filer.

Table 12-6 Add/Edit Veritas File System (VxFS) filer options

Field	Description
This is a VCS clustered file server	Select the check box if the Veritas File System server is part of a Veritas Cluster Server (VCS) configuration.

Table 12-6   Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
VCS cluster name	Enter the logical name of the VCS cluster. This field is avalable only if you select the **This is a VCS clustered file server** check box.
Cluster Node IP addresses	Enter the comma-separated list of the host names or IP addresses of the physical nodes in the VCS cluster.
Filer hostname or IP address	Enter the hostname or IP address of the filer that you want Data Insight to monitor. **Note:** The hostname or IP address should be the same as the filer name entered in Symantec Data Loss Prevention Discover targets.
Collector	From the drop down, select the Collector worker node configured to scan the filer. Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer. **Note:** Ensure that the Collector node monitoring the NFS must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through **Server Manager** > **Add roles** option.
Indexer	From the drop down, select the Indexer worker node configured for the filer. Events and meta-data collected from the filer is processed and stored on the Indexer node. You can also migrate the file server to another Indexer.
Login credentials	See "Credentials required for configuring Veritas File System (VxFS) servers" on page 125. Specifying filer administrator credentials is optional, if you choose not to monitor events on the filer, nor enable share discovery.

Table 12-6   Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
Test credentials	Click to test the availability of network connection between the Collector worker node and the filer, and to test the validity of the specified credentials. Symantec recommends that you test the connection before proceeding to ensure that Data Insight is able to connect to the filer.
Monitoring details	Select **Automatically discover and monitor shares on this filer** to have Data Insight automatically discover shares of the filer and add them to the configuration. Discovery of shares takes place as soon as you add a new filer and then twice each day at 2:00 a.m. and 2:00 p.m. You can also choose to add shares manually. See "Adding shares" on page 160.
Exclude shares from discovery	Enter the details of shares which should not be included during discovery. This option is available if you select **Automatically discover all shares on this filer.** Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, tmp* ignores `tmp_A`, `tmp_abc`, *$ ignores shares `C$`, `EXT$` and others.

Table 12-6   Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
Audit details	Select to enable file system monitoring on the filer.
	Enter the following details:
	■ Select the **Enable file system event monitoring** check box to enable event monitoring on the VxFS filer.
	■ Time to live - The value indicates the time for which the VxFS plugin will try to communicate with Data Insight. If communication fails after the specified time, the plugin will terminate and stop capturing events from the VxFS filer. The default TTL value is 24 hours.
	■ Records per file - The number of records after which the events are flushed to the Collector node. You can also enable an advanced setting to flush the records to the Collector node every 10 minutes, irrespective of the number of records specified. By default, the limit is set to 100000 records per file. See "Configuring advanced settings" on page 203.
	■ Domain: The name of the LDAP or NIS domain that the filer is a part of. The VxFS filer that you want to add should not be part of two domains at the same time.
Enable filer scanning	Select the checkbox to enable filer scanning according to the specified schedule.

Table 12-6   Add/Edit Veritas File System (VxFS) filer options *(continued)*

Field	Description
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ **Use the Collector's scanning schedule** ■ **Define custom schedule** Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 p.m. on the last Friday of each month. **Note:** You can customize the schedule per share using the Add/Edit Share dialog box.
Scanner credentials	See "Credentials required for configuring Veritas File System (VxFS) servers" on page 125.
Scan newly added shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule.

See "Enabling export of UNIX/Linux NFS shares on VxFS filers" on page 127.

Add/Edit a generic storage device options

Use this dialog box to add a new generic storage device to Symantec Data Insight or to edit the configuration of an existing device.

Table 12-7   Add/Edit generic device options

Field	Description
Filer hostname or IP address	Enter the hostname or IP address of the device that you want Data Insight to monitor.

Table 12-7   Add/Edit generic device options *(continued)*

Field	Description
Collector	From the drop-down, select the Collector worker node configured to scan the filer. Data Insight connects to the filer from this server. It is recommended that the Collector worker node share a fast network with the filer. **Note:** When monitoring NFS shares, ensure that the Collector node monitoring the filer must have services for NFS enabled as file server roles. You can install the role on Windows 2008 through the **Server Manager** > **Add roles** option.
Indexer	From the drop-down, select the Indexer worker node configured for the filer. Events and meta-data collected from the filer is processed and stored on the Indexer node.
Domain	From the drop-down, select the domain to which the device belongs. This option is enabled when monitoring NFS shares.
Enable filer scanning	Select the check box to enable filer scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for shares of this filer: ■ Use Collector's default scanning schedule. ■ Use custom schedule. See "Custom schedule options" on page 157. Symantec Data Insight periodically scans shares of the filer to obtain file metadata and security descriptors. Each Collector worker node by default initiates a full scan of shares at 7:00 p.m. on the last Friday of each month.

Table 12-7   Add/Edit generic device options *(continued)*

Field	Description
Test credentials	Do the following: 1 Click to test whether the selected credentials have the required pemission to scan shares on the device. 2 Enter the name of a share on the the device, and click **OK**.
Scanner credentials	See "Credentials required for scanning a generic device" on page 130.
Scan new shares immediately	Select this option to scan newly added shares immediately, instead of waiting for the normal scan schedule. Scanning proceeds only when scanning is permitted on the Collector node.

Custom schedule options

Table 12-8 describes the options that you can use to define the frequency of the scans.

Table 12-8   Custom schedule options

Option	Description
Never	Runs the scan as and when required.
Once	Runs the scan once at the specified time and date.
Daily	Runs the scan once every day. You must specify the time when the scan should be run.
Weekly	Runs the scan once every week. You can choose to run it on every weekday, or on specific weekdays. Also, you must specify the time when the scan should be run.
Monthly	Runs the scan on the specified days of a month. You must specify the days of the month and the time when the scan should be run. Separate multiple days with a comma. For example, 2,5.

Table 12-8   Custom schedule options *(continued)*

Option	Description
Custom Cron	Runs the scan according to a defined cron schedule. You can build strings in the cron format to specify custom schedules such as every Friday at noon, or weekdays at 10:30 a.m., or every 5 minutes between 9:00 a.m and 10 a.m. on Wednesdays and Fridays. More information about defining a cron schedule, see *http://quartz-scheduler.org/api/2.1.7*.

Editing filer configuration

After you add a filer to Data Insight, you can edit the filer's configuration. For example, you might need to edit any of the following:

- The IP address or hostname of the filer.

- The username and password of the user authorized to log in to the filer.

- The IP address or hostname of the Collector worker node configured to scan the filer.

- The Indexer node associated with the filer.

- The scanning schedule.

- The scanner credentials.

- Whether all shares are to be monitored.

- Whether new shares are to be scanned immediately.

To edit filer configuration

1   In the Console, click **Settings** > **Filers**.

     This displays the list of available filers.

2   Do one of the following:

3   On the Edit Filer screen, do the following:

- Make the necessary configuration changes.

- To migrate a filer to another Indexer, click **Migrate**.

- On the **Select Indexer** panel, select the new Indexer that you want to associate with the file server.

- Click **Select**.

- Click **Save** to save the changes.

4   To view the status of migration to another Indexer, navigate to the Filers list page.

5   Click the **Select Action** drop-down corresponding to the filer being migrated, and select **View Migration Status**.

On the **Migration Status** panel, you can view the status of each share on the filer that is being migrated. The status can either be Success or Pending.

See "About migrating storage devices across indexers" on page 219.

Deleting filers

You can delete a configured filer.

To delete a filer

1   In the Console, click **Settings** > **Filers** to display the configured filers.

2   Do one of the following:

- In the filer summary table, click the **Select Action** drop-down, and select **Delete**.

- Click the filer you want to delete, and on the filer details page, click **Delete**.

3   Click **OK** on the confirmation message.

Viewing performance statistics for file servers

You can view the performance graphs for the configured filers.

To view file server statistics

1   Click **Settings** > **Filers**.

2   Select the filer for which you want to view the performance statistics.

3   On the filer details page, click the **Statistics** sub-tab.

4   From the **Period** drop-down, select the duration for which you want to view the data.

5   The **Statistics** sub-tab displays the following high-level data:

- The time (in milliseconds) required to perform one CIFS I/O operation for a filer. The graph displays both average and maximum latency statistics.

- The average and maximum rate of incoming events per second for a filer.

■ The count of all files and folders across all shares for a filer. The graph displays the number of paths that are scanned across all shares on a filer

Note: For EMC Celerra, VxFS, and Windows File Servers, you can only view the count of files and folders across all the shares on the filer.

Adding shares

After you add a filer, you can add shares present on the filer that you want Data Insight to monitor. You need to perform this operation if you have selected **Shares will be added manually** option when adding a filer.

To add a share

1 In the Console, click **Settings** > **Filers** to expand the Filer node.

2 Click the filer for which you want to add a share.

3 On the Filer Detail screen, click **Monitored Shares**.

4 On the Monitored Shares screen, click **Add New Share**.

5 On the **New Share** screen, enter the share properties, and click **Add New Share**.

See "Add New Share/Edit Share options " on page 160.

Add New Share/Edit Share options

Use this dialog box to add a new share to Symantec Data Insightor to edit the configuration of an existing share.

Table 12-9   Field Descriptions

Field	Descriptions
Share name	Enter the name of the share you want to add. For example, *share1*.
	If this share belongs to a clustered filer, enter share name as *fileserver@share*, where, *fileserver* is the name of the file server within the cluster that hosts the share.
Physical path on filer	Enter the physical path of the share on the filer. For example, `F:\<Share name>`.

Table 12-9   Field Descriptions *(continued)*

Field	Descriptions
Scanning schedule	Select one of the following to define a scanning schedule: ■ **Use filer's scanning schedule** ■ **Define custom schedule**
Enable legal hold for this share	Select to preserve the activity information on the share. Selecting this option disables the deletion or archiving of access event information on the share. See "About archiving data" on page 47. See "About purging data" on page 48.

About disabled shares

Data Insight automatically discovers new shares on file servers and adds them to Data Insight configuration, if automatic discovery of shares is enabled for the file server. Share discovery takes place when a new filer is added to Data Insight or when the scheduled share discovery process runs on the filer.

During the process of share discovery, Data Insight also discovers shares that have been deleted from the filer but are still present in the Data Insight configuration. Such shares are marked as disabled and Data Insight stops scanning these shares.

Managing shares

On the Monitored Shares details page you can view the detailed information about configured shares and run a customized scan on the configured shares.

Use the provided dynamic search filter to search for configured shares based on the name of the share.

To view configured shares

1   In the Console, click **Settings** > **Filers**.

2   Click the filer on which the share resides.

3   On the Filer Detail screen, click **Monitored Shares**.

    Review the following information about the shares:

- ID of the share. The ID is required during troubleshooting. This column is hidden by default.

- The name of the share.
  If this share belongs to a clustered filer, then the name should appear as *fileserver@share*, where, *fileserver* is the name of the file server within the cluster that hosts the share.

- Type of this share, CIFS or NFS.

- Enabled status of this share. This column is hidden by default.

- Legal hold status for this share. This column is hidden by default.

- The scanning schedule for the share. This column is hidden by default.

- The date and time of the last full scan of the share.

- The date and time of the last incremental scan.
  Incremental scans are scans of the file system that includes only those paths that have changed since the last full scan. Incremental scans are much faster than full scans and they take place once every night at 7:00 p.m. You can configure incremental scans on the **Settings** > **Data Insight Servers** > **Advanced Settings** page.

- The time this share's index was last updated with scan information.
  After every scan, the index is updated with information about the changes to the folder hierarchy on a share. This indicates whether the last update was successful or failed. It also indicates the number of scan files pending for this share on the Indexer and the number of files that failed to be indexed. Such files are present in the `$data/indexer/err` folder on the Indexer. If there are failed files on the Indexer, you can move them from the `err` folder to `$data/inbox` folder and attempt a full scan of the share. If the information fails to be indexed again, contact Symantec Support.

- The time this share's index was last updated with access event information.
  As new access events come in, the index for the share is periodically updated with information about the new access events. This indicates whether the last update was successful or had failed. It also indicates the number of audit files pending for this share on the Indexer and the number of files that failed to be indexed. Such files are present in the `$data/indexer/err` folder on the Indexer. If there are failed files on the Indexer, you can move them to the `$data/inbox` folder on the Indexer. If they fail to be indexed again, contact Symantec Support.

- The status of event monitoring for the share, whether enabled or disabled.

- Whether a legal hold is being enforced for the share. You can choose to prevent access information for a share from being archived or deleted by

putting a legal hold on the share. Data Insight preserves the access events information for such shares indefinitely.

See "Add New Share/Edit Share options " on page 160.

4   Click the Export icon at the bottom of the page to save the data on the Monitored Shares panel to a `.csv` file.

You can also add a new share, edit the share's configuration, delete the share, start an unscheduled scan for a share, view the scan history of a share, and download Data Insight logs from this page.

To view the scan history of a share

1   In the Console, click **Settings** > **Filers**.

2   Click the filer on which the share resides.

3   On the Filer Detail screen, click **Monitored Shares**.

4   Click the **Action** drop-down for the corresponding share, and select **Scan History**.

The scan history for the share appears. You can view the details in a tabular format or in a Timeline view. The tabular view displays the following details of a scan:

■   The start and end time of the scan.

■   The time taken for the scan.

■   The type of scan, whether full or incremental.

■   The Collector node associated with the share.

■   The details of the scan. For example, if a scan has failed, the **Details** column indicates the exit code for the error message.

■   The user account that initiated the scan.

The Timeline view displays an hourly and daily overview of the scans run on the share, including information about the status of the scan, whether failed, successful, partially successful, or aborted.

You can also view the scan history for a share from the **Scan History** sub-tab of the **Scanning** dashboard.

To view events pertaining to a share

1   In the Console, click **Settings** > **Filers**.

2   Click the filer on which the share resides.

3   On the filer details screen, click **Monitored Shares**.

4   Click the **Action** drop-down for the corresponding share, and select **Event Log**.

     The event log for that share appears.

5   To download the Data Insight logs for the share, click the **Select Action** drop-down for the corresponding share, and select **Download Log**.

     Data Insight downloads a compressed folder containing logs for this share from all relevant Data Insight servers.

     See "Downloading Data Insight logs" on page 282.

To scan one or more shares in a batch

1   On the **Monitored Shares** tab, select a share, and select Scan from the **Select Action** drop down corresponding to a share.

     Note: The **Scan** option is not available for shares that have been disabled.

     ---

2   To scan multiple share, select one or more shares using the check boxes.

3   Click **Scan**, and select **Scan Selected Records**.

     Optionally, filter shares as needed using the filters available on the page. Click **Scan**, and select **Scan Filtered Records**.

Note: You can use a command line utility, `scancli.exe`, to further customize the scan, view the scan jobs running on a specified node, or display the scan status for specified shares. See scancli.exe on page 302. You can also use the Scanning dashboard view to scan shares and site collections based on more granular criteria.

To enable or disable shares

1   In the Management Console, click **Settings** > **Filers**.

2   Click the filer on which the share resides.

3   On the filer detail screen, click **Monitored Shares**.

4   Click the **Action** drop-down corresponding to a share, and select **Enable** to enable a share that has been disabled.

5   Click the **Action** drop-down corresponding to a share, and select **Disable** to disable a share.

Editing share configuration

After you add a share to Data Insight, you can edit the share's configuration. For example, you might need to edit the scanning schedule.

To edit share configuration

1   In the Console, click **Settings** > **Filers** to expand the Filer node.

     This displays the list of available filers. Click the appropriate filer to open the Filer details page.

2   Select the share whose configuration you want to edit, click the **Select Action** drop-down and select **Edit**.

3   On the Edit Share screen, make the necessary configuration changes.

4   Click **Save**.

Deleting shares

You can delete a configured share.

To delete a share

1   In the Console, click **Settings** > **Filers** to display the configured filers.

2   Click the filer, on which the share that you want to delete exists.

3   On the filer details page, under **Monitored Shares**, select the share that you want to delete.

4   Click the **Select Action** drop-down and select **Delete**.

5   Click **OK** on the confirmation message.

About configuring a DFS target

Symantec Data Insight supports Microsoft Distributed File System (DFS) targets.

DFS simplifies access to and management of shares by mapping a single logical namespace to shared folders on multiple filers. You can create folders within a DFS to create an additional level of hierarchy. For example, if you have a NetApp filer, NETAPP01, which has a shared folder called `NetAppShare1`. You can link a target, `HQ\Sales\Test`, present on a DFS server, DFSSvr01, to the subfolder named `Finance` within `NetAppShare1`.

You must first import the DFS mappings to physical shares in to Data Insight before you can view data using DFS hierarchy.

Configuring a DFS target

Before you can configure a DFS target you must configure all file servers which map to your DFS targets.

To set up a DFS target

1   Log in to the Management Console.

2   Create a .csv file containing the following information:

- The name of the DFS server.

- The DFS target.

- The name of the filer that contains the share that you want to map to the DFS target.

- The share on the filer.

- Path under the physical share, if the DFS folder is mapped to a folder under physical share, else this value can be blank.

    For example, *DFSSvr01,HQ\Sales\Test,NETAPP01,NetAppShare1,\Finance*.

3   Click the **Settings** tab.

4   Click **Filers**, and select **Import DFS Mappings**.

5   In the Add new DFS mappings dialog, browse to the location of the .csv file, and click **Upload**.

6   Alternatively, open a Windows command prompt, and change to the `installdir\bin` directory, where `installdir\bin` is the installation path for Symantec Data Insight.

7   Type the following command:

    `configdb -H  <name of the .csv file>`

About the DFS utility

The DFS utility, `mxdfsscan.exe`, maps root level DFS paths to actual storage server or share paths. It is used to export the DFS components (roots, root targets, links, and link targets) for all Windows DFS namespaces. The utility finds out physical level storage/filer link for all Domain DFS paths. It takes DFS root UNC path as input, for example `\\<DFS domain>\root`. This utility only enumerates online links and skips all offline links. It generates the output in .csv format.

The `mxdfsscan.exe` is a command line utility.

mxdfsscan lets you specify only a single namespace at a time. You can run the utility twice to get information from two different DFS namespaces and create two different files from the output, for example, `test1.csv` and `test2.csv`. You can then import settings from `test1.csv` and `test2.csv` from the Data Insight Management Console

When you import a new DFS mapping file to Data Insight, the old mappings are maintained in Data Insight. For example, if you import mappings from `test1.csv` and then from `test2.csv`, the mappings from both files are displayed in Data Insight. However, if there are some duplicate mappings (the same DFS link appears twice – whether mapped to the same physical path or a different path), these mappings are not imported. A message is displayed indicating that there are duplicate mappings and hence one or more mappings cannot be imported.

Running the DFS utility

Ensure that the DFS servers are accessible from the machine you use to run the DFS utility.

To run the DFS utility

1   From the Windows Start menu, select **Run** and type `cmd` in the dialog box to open a command prompt window.

2   Change to the `installdir\bin` directory, where `installdir\bin` is the installation path for Symantec Data Insight.

3   Type the following command:

    `mxdfsscan -n \\<DFS domain>\root -f dfsmap.csv`

    where,

    `-n` is the name of the DFS root

    `-f` is the file name to which the DFS mappings have to be exported.

    `-e` is the option to exclude domain DFS paths. For example,

    `mxdfsscan.exe -n \\MSAD\newroot -f dfsmap.csv -e exclude.txt`

    An exclude list can have max 128 exclude entries. For example,

    `\\DFS\root\AP\NUY`

    `\\DFS\root\users`

    `-c` is the option to traverse a specified number of intermediate DFS servers to find a physical storage path. If the -c option is not specified then the utility takes the default value 5. This option helps avoid circular links in a DFS environment. If there are more hops then it logs all such links into `dfs_log_links.txt`.

Importing DFS mapping

You can import DFS mappings to Data Insight from the Management Console. To import DFS mappings, complete the following steps.

To import DFS mappings

1   Create a .csv file that contains information about the DFS mappings.

    See "Running the DFS utility" on page 167.

2   In the Console, click **Settings** > **Filers** to display the list of available filers.

3   Click **Import DFS mappings**.

4   On the Import DFS mappings window, browse to the location of the `.csv` file that contains information about the mapped DFS namespaces.

5   Click **OK**.

Configuring SharePoint monitoring

This chapter includes the following topics:

- About SharePoint server monitoring

- Credentials required for configuring SharePoint servers

- Configuring a web application policy

- About the Data Insight Web service for SharePoint

- Viewing configured SharePoint web applications

- Adding web applications

- Editing web applications

- Deleting web applications

- Adding site collections

- Managing site collections

- Removing a configured web application

About SharePoint server monitoring

You can use Symantec Data Insight to monitor unstructured data residing on servers running any of the following:

- Microsoft SharePoint™ 2013

- Microsoft SharePoint™ 2010

- Microsoft Office SharePoint™ Server 2007 (MOSS 2007)

Data Insight monitors accesses to the data in the following SharePoint library types:

- Document library - Stores documents in the .pdf, .doc, .xls, .txt and other such file extensions.

- Picture library - Stores images.

Before you use Data Insight to scan a SharePoint server, you must complete the following tasks:

- To be able to configure Data Insight, you must know the URLs of the target SharePoint Web applications.

- Ensure that .NET Framework 3.0 or 3.5 is installed on the Collector node that is responsible for the discovering site collections and collecting audit logs.

- Configure a policy for each Web application.
  See "Configuring a web application policy" on page 171.

- Install and configure the Data Insight Web service on the SharePoint server.

- Enable auditing on the SharePoint server. You can enable auditing from the Management Console when you add Web applications, or directly from the SharePoint server.
  See " Add/Edit Web application options" on page 175.

See "Supported file servers and platforms" on page 18.

Credentials required for configuring SharePoint servers

Table 13-1   Credentials required for configuring SharePoint servers

Credential	Details
Credentials required to install the Data Insight Web service on the SharePoint Server.	This credential belongs to a user in the Administrators group on the SharePoint server.

Table 13-1   Credentials required for configuring SharePoint servers *(continued)*

Credential	Details
Credentials required to discover Web applications or site collections, and to collect scan information and audit data	This credential belongs to a site collection administrator for the configured sites and it must be in the same domain as the SharePoint server. It must have full control permissions not only on the configured Web applications, but also on the Web applications that are added to SharePoint subsequently. See "Configuring a web application policy" on page 171. for details on adding such a user credential.

Configuring a web application policy

When configuring SharePoint from the Data Insight console, you must specify an account for monitoring the configured site collections. This account must be a site collection administrator for the configured sites and it must be in the same domain as the SharePoint server. It must have full control permissions not only on the configured web applications, but also on the web applications that are added to SharePoint subsequently. The account should have the necessary privileges to set the appropriate audit flags, gather metadata about site collection content, and gather audit data from SQL Server databases for SharePoint

To enable Data Insight to gather audit and metadata from multiple site collections using a single user account, you must configure a policy for each web Application from the SharePoint Central Administration Console.

To configure a policy for web Application in SharePoint 2007

1   In the Central Administration website, click **Application Management**.

2   Under the Application Security section, click **Policy for Web application**.

3   Click **Add Users**.

4   In the **Web Application** drop-down list, select the web application that contains the site collections that you want Data Insight to monitor.

5   Select the appropriate zone. You can select **(All Zones)** if you want the user to be given permissions on all zones for the web application.

6   Click **Next**.

7   Choose the user account that has Full Control permissions.

     In the **Choose Permissions** section, select **Full Control - Has full control**

8   Specify whether this account operates as SharePoint System account. If you select the **Account operates as System** check box, all accesses made by this user account are recorded with the user name, *SharePoint System*.

9   Click **Finish**.

To configure a policy for web Application in SharePoint 2010 and SharePoint 2013

1   In the Central Administration Console, click **Application Management**.

2   Under the web Applications section, click **Manage Web Applications**.

3   In the table displaying web application details, select the appropriate web application.

4   Click **User Policy**.

5   In the Policy for web Application popup, click **Add Users**.

6   Select the appropriate zone. You can select **(All Zones)** if you want the user to be given permissions on all zones for the web application.

7   Click **Next**.

8   Choose the user account that will have Full Control.

     In the **Choose Permissions** section, select **Full Control - Has full control**

9   Specify whether this account operates as SharePoint System account. If you select the **Account operates as System** check box, all accesses made by this user account are recorded with the user name, *SharePoint System*.

10   Click **Finish**.

About the Data Insight Web service for SharePoint

Before you can configure SharePoint monitoring in Data Insight, you must install the Data Insight Web service on the SharePoint server. The Data Insight Web service performs the following functions:

■   Enables or disables auditing on the SharePoint server.
     You can enable or disable auditing of access events at the site collection level, either manually or from the Data Insight Management Console when adding Web applications.

■   Discovers site collections within a Web application.

■   Discovers all Web sites and lists or libraries within a site collection

■   Retrieves access event data from the SharePoint server. Data Insight uses this data to service queries on user activity and data access.

■   Deletes audit logs from a site collection.

Installing the Data Insight web service for SharePoint

To enable Data Insight to collect access events, you must install the Data Insight web service, on a SharePoint server. When installing in a SharePoint farm, you must ensure that the web service is configured on all front-end web servers in the farm. Perform the following steps on any one front-end web server in your SharePoint farm. The web service installer automatically deploys the web service to all front-end web servers in the farm.

To install the Data Insight web service

1   Log on to the SharePoint server with an account that has SharePoint administrator privileges.

2   Load the Data Insight media on your SharePoint server.

3   Navigate to the folder where you have extracted or copied the installers.

4   To start the installation, double-click `Symantec_DataInsight_sharepoint_4.5_N.exe`, where, **N** is the build number.

5   Work through the installation wizard.

6   Click **Finish** to complete the installation process.

7   Verify whether the web service is deployed as expected.

After installing the Data Insight web service, you must verify whether it is successfully deployed on all front-end web servers in the SharePoint farm.

To verify the deployment of the web service in SharePoint 2007

1   In the Central Administration console, click the **Operations** tab.

2   Under **Global Configurations** section, click **Solution Management**.

3   Verify that the status for Data Insight solution for SharePoint is set to **Deployed**.

4   Click the link for the solution. Verify that the solution is deployed to all the front-end web servers in the farm by checking the value of **Deployed To** field.

To verify the deployment of the web service in SharePoint 2010 and SharePoint 2013

1   In the Central Administration console, click the **System Settings**.

2   Under the Farm Management section, click **Manage Farm Solutions**.

3   Verify that the status for Data Insight solution for SharePoint, *sdispwebsvc.wsp*, is set to **Deployed**.

4   Click the link for the solution. Verify that the solution is deployed to all the front-end web servers in the farm by checking the value of **Deployed To** field.

Viewing configured SharePoint web applications

In the Management Console, you can view all the SharePoint web applications that Data Insight is configured to monitor.

Use the provided dynamic search filter to search for configured web applications based on various predefined criteria, for example, the state of event monitoring for the web application. You can also use the **Filter** field at the top of the content pane to filter the list of web applications based on the URL in addition to the predefined filter criteria. The displayed list of web applications changes automatically as you type the term in the **Filter** text box or when you select the check box for a filter criteria. For instance, when you select a Collector node in the **By Collector** filter, Data Insight displays a list of web applications that are associated with the selected Collector node.

To view configured web applications

1   In the Console, click **Settings** > **SharePoint Web Applications**.

    The screen displays the list of configured web applications.

2   Review the following information about the web applications:

    ■   The URL of the web application.

    ■   The status of the web application — whether scanning and event monitoring are enabled for this web application.

    ■   The Collector node for the web application.

    ■   The Indexer node for the web application.

    ■   Click on a configured web application to view its detailed information, or click the **Select Actions** drop-down and select **View**.
        The web application details page appears.

Adding web applications

You must install the Data Insight web service on the SharePoint server, before you can add the web applications that you want Data Insight to monitor. In case the web service is not installed, Data Insight prompts you to install it before you can proceed with adding web applications.

To add web applications

1   In the Console, click **Settings** > **SharePoint Web Applications**.

    The SharePoint page displays the list of configured web applications.

2   Click **Add SharePoint Web Application**.

3   On the Add web Application screen, enter the URL of the web application you want to add and enter the properties.

4   Click **Save**.

Add/Edit Web application options

Use this dialog box to add a new Web application to Symantec Data Insight or to edit the configuration of an existing web application.

Table 13-2   Add/Edit Web application options

Field	Description
Web application URL	Enter the URL of the web application that you want Data Insight to monitor.
Collector for this Web application	From the drop-down, select the Collector worker node configured to scan the SharePoint server.
	Data Insight connects to the SharePoint server from Collector node. It is recommended that the Collector worker node share a fast network with the SharePoint server.
Indexer for this Web application	From the drop-down, select the Indexer worker node configured to scan the SharePoint server.
Default Site Collection Administrator	Enter the credentials that Data Insight should use to provide authenticated access to the Data Insight Web service on the SharePoint server.
	See "Configuring a web application policy" on page 171.
Verify credential	Click to test the availability of network connection between the Collector worker node and the SharePoint server, and to test the validity of specified credentials. You must first ensure that the Data Insight Web service is already installed on the SharePoint server.
	Symantec recommends that you verify the credentials before proceeding to ensure that Data Insight is able to connect to the SharePoint server.

Table 13-2   Add/Edit Web application options *(continued)*

Field	Description
Automatically discover and add all site collections in the selected Web applications to Data Insight	This checkbox is selected by default. This option allows you to automatically include all site collections in the selected Web application for the purpose of monitoring. Clear the check box to add site collections manually. You can do this from the Web Application details page.
Exclude following site collections from discovery	Enter the details of the site collections which should not be included during discovery. This option is available if you select **Automatically discover and add site collections in the added SharePoint Web Applications**. Specify comma separated patterns that you want to ignore. Patterns can have 0 or more wildcard * characters. For example, https://webapp1/sites/test* ignores site collections https://webapp1/sites/testsite1 and https://webapp1/sites/testsite2.
Monitor access for this Web application	Select to enable monitoring of access events for the Web application.
Automatically enable auditing for site collections of this Web application	Select to automatically enable event monitoring for all site collections of this Web application. You can also choose to manually enable auditing by logging in to the SharePoint server. For this purpose, you must have site collection administrator privileges on the SharePoint server.

Table 13-2   Add/Edit Web application options *(continued)*

Field	Description
Delete audit logs from SharePoint database after importing in Data Insight.	Select to delete audit logs from SharePoint to prevent the Web application database from growing too large. By default, Data Insight deletes audit logs that are older than two days from the SharePoint server once every 12 hours. You can configure this interval from the **Advanced Settings** page for the corresponding Collector. You can choose to customize how often Data Insight should delete old audit logs from the Data Insight Servers node on the Management Console. See "Configuring advanced settings" on page 203.
Enable scanning for this Web application	Select the checkbox to enable SharePoint scanning according to the specified schedule.
Scanning schedule for full scans	Select one of the following to define a scanning schedule for the SharePoint servers in this farm: ■ **Use the Collector's scanning schedule.** ■ **Define custom schedule for farm.** From the drop-down, select the appropriate frequency. See "Custom schedule options" on page 157. Symantec Data Insight periodically scans site collections to obtain file metadata. Each Collector worker node by default initiates a full scan the SharePoint servers at 11:00 p.m. each night. **Note:** You can also customize the schedule for each site collection using the Add/Edit Site Collection dialog box.
Scan newly added site collections immediately	Select this option to scan newly added site collections immediately, instead of waiting for the normal scan schedule. Scanning will still proceed only when scanning is permitted on the Collector node.

Editing web applications

After you add a web application to Data Insight, you can edit its configuration. For example, you might need to edit any of the following:

- The user who is authorized to log in to the SharePoint server.

- The Collector worker node that is configured to scan the SharePoint server.

- The Indexer node that is associated with the web application.

- Enable or disable web application scanning or audit.

- The scanning schedule.

To edit web applications

1   In the Console, click **Settings** > **SharePoint web applications**.

2   Do one of the following:

- In the web application summary table, click the **Select Actions** drop-down and select **Edit**.

- Click the web application whose configuration you want to edit. On the web application detail screen, click **Edit**.

3   On the Edit Web Application screen, do the following:

- Make the necessary configuration changes.

- To migrate a filer to another Indexer, click **Migrate**.

- On the **Select Indexer** panel, select the new Indexer that you want to associate with the web application.

- Click **Select**.

- Click **Save** to save the changes.

4   To view the status of migration to another Indexer, navigate to the Web Application list page.

5   Click the **Select Action** drop-down corresponding to the web application being migrated, and select **View Migration Status**.

On the **Migration Status** panel, you can view the status of each site collection of the web application that is being migrated. The status can either be Success or Pending.

See "About migrating storage devices across indexers" on page 219.

See " Add/Edit Web application options" on page 175.

Deleting web applications

You can delete a configured web application.

To delete a web application

1   In the Console, click **Settings** > **SharePoint Web Applications**.

2   Do one of the following:

   ■   In the web application summary table, click the **Select Actions** drop-down and select **Delete**.

   ■   Click the web application that you want to delete. On the web application detail screen, click **Delete**.

3   Select the check box to disable auditing on the web application.

4   Click **Yes** on the confirmation dialog box.

Adding site collections

You can configure Data Insight to scan one or more site collections within a Web application.

To add site collections

1   In the Console, click **Settings** > **SharePoint Web Applications**.

2   In the Web application summary table, click the Web application that the site collections are a part of.

3   Click the **Monitored Site Collections** tab.

   The screen displays a list of all configured site collections.

4   Click **Add Site Collection**.

5   On the Add New Site Collection dialog, enter the site collection properties, and click **Add New Site Collection**.

Data Insight monitors access events on the SharePoint servers and maps all SharePoint access types such as CheckOut, CheckIn, ChildDelete, Copy, Update, Delete, Move, SecurityChange, and Undelete to Data Insight meta access types - Read, Write, Delete, and Rename. Data Insight automatically configures the audit settings for a site collection for these events after you enable auditing on the SharePoint server.

You can verify the audit events that Data Insight is monitoring for every site collection.

To verify the audit events monitored by Data Insight.

1 Use a web browser to open a site collection URL.

2 Click **Site Actions** > **Site Settings** > **Site Collection Audit Settings**.

3 Ensure that the appropriate events that you want to audit are selected for **Documents and Items** and for **Lists, Libraries, and Sites**.

See "Add/Edit site collection options" on page 180.

Add/Edit site collection options

Use this dialog box to add a new site collection to a configured Web application or to edit the configuration of an existing site collection.

Table 13-3   Add/Edit site collection options

Field	Description
Site Collection URL	Enter the URL of the site collection that you want to add.
Site Collection Title	Enter a logical name for the site collection.
Scanning schedule	Select one of the following to define a scanning schedule for the site collection: ■ **Use Collector's scanning schedule** ■ **Define custom scanning schedule** From the drop-down, select the appropriate frequency. ■ See "Custom schedule options" on page 157. Symantec Data Insight periodically scans site collections to obtain metadata. Each Collector worker node by default initiates a full scan of the SharePoint servers at 11:00 p.m. each night.
Enable legal hold for this site collection	Select to preserve the activity information on the site collection. Selecting this option disables the deletion or archiving of access event information on the site collection. See "About archiving data" on page 47. See "About purging data" on page 48.

Managing site collections

On the site collections details page, you can view detailed information about a site collection, and edit or delete the site collection. You can also run a scan on the site collection from this page.

Use the provided dynamic search filter to search for configured sites based on the name of the site.

To view configured site collections

1   In the Console, click **Settings** > **SharePoint Web Application**.

2   In the web application summary table, click the web application that the site collections are a part of.

The screen displays a list of all configured site collections.

3   On the web application details page, click **Monitored Site Collections**.

4   On the Site Collections listing page, review the following information about the configured site collections.

- The name of the site collection.

- The URL of the site collection

- The scanning schedule for the site collection.

- The date and time of the last full scan of the site collection.

- The time this site collection's index was last updated with scan information. After every scan, the index is updated with information about the changes to the folder hierarchy on a site collection. This column indicates whether the last update was successful or has failed. It also indicates number of scan files pending for this site collection on the Indexer and the number of files that failed to be indexed. The files that failed to be indexed, are present in the `$data/indexer/err` folder on the Indexer. If you do have failed files on the indexer, you can move them from the `err` folder to the `$data/inbox` folder and attempt a full scan of the site collection.
  If the scan information again fails to be indexed, contact Symantec support.

- The time this site collection's index was last updated with access event information.
  As new access events come in, the index for the site collection is periodically updated with information about these events. This indicates whether the last update was successful or has failed. It also indicates number of audit files pending for this site collection at the Indexer and the number of files that failed to be indexed. Audit files are present in the `$data/indexer/err`

folder on the Indexer. If you do have failed files on the indexer, you can try moving them back to `$data/inbox` folder on the Indexer.

If the new audit information again fails to be indexed, contact Symantec support.

- The status of event monitoring for the site collection, whether enabled or disabled.

- Whether a legal hold is being enforced for the site collection. You can choose to prevent access information for a site collection from being archived or deleted by putting a legal hold on the site collection. Data Insight preserves the access events information for such site collections indefinitely.

  See "Add/Edit site collection options" on page 180.

5   Click the Export icon at the bottom of the page to save the data on the **Monitored Site Collections** panel to a .csv file.

You can also edit the properties of the site collection, start an unscheduled scan of the site collection, delete the site collection, view the event log or scan history of the site collection, or download logs for troubleshooting purposes.

To edit a site collection

1   On the web application details page, click **Monitored Site Collections**.

2   Select the site collection that you want to edit, and from the **Select Action** drop-down, select **Edit**.

3   On the Edit site collection screen, make the necessary configuration changes.

4   Click **Save**.

To delete a site collection

1   On the web application details page, click **Monitored Site Collections**.

2   Select the site collection that you want to delete, and from the **Select Action** drop-down, select **Delete**.

3   Click **OK** on the confirmation message.

To view the scan history of a site collection

1   On the web application details page, click **Monitored Site Collections**.

2   Select the site collection for which you want to view the scan history, and from the **Select Action** drop-down, select **Scan History**.

The scan history for the site collection appears. You can view the details in a tabular format or in a Timeline view. The tabular view displays the following details of a scan:

- The start and end time of the scan.

- The time taken for the scan.

- The type of scan, whether full or incremental.

- The Collector node associated with the site collection.

- The details of the scan. For example, if a scan has failed, the **Details** column indicates the exit code for the error message.

- The user account that initiated the scan.

The Timeline view displays an hourly and daily overview of the scans run on the site collection, including information about the status of the scan, whether failed, successful, partially successful or aborted.

You can also view the scan history of a site collection from the **Scan History** sub-tab of the **Scanning** dashboard.

To view events pertaining to a site collection

1   In the Console, click **Settings** > **SharePoint Web Applications**.

2   On the web application details screen, click **Monitored Site Collections**.

3   Click the **Select Action** drop-down for the corresponding site collection, and select **Event Log**.

    The event log for that site collection appears.

4   To download the logs for the site collection, click the **Select Action** drop-down for the corresponding site collection, and select **Download Logs**.

    Data Insight downloads a compressed folder containing the logs for this site collection from all relevant Data Insight servers.

    See "Downloading Data Insight logs" on page 282.

To scan site collections in a batch

1   On the Monitored site collections tab, click the **Scan** button.

    Note: The **Scan** option is not available for site collections that have been disabled.

    ---

2   On the Scan Site Collections pop-up, select one of the following:

- **Scan all** - To scan all the configured site collections immediately.

- **Scan with last scan status** - To scan site collections based on the following criteria:

    - Site collections on which the last scan has failed completely.

- Site collections that have never been scanned before.

- Site collections on which the last scan has failed on certain paths.

3   Select one or more of the following conditions:

- Scan the site collections that have not been scanned for *n* number of days. Enter the interval in the field.

- Include site collections matching specified patterns. You can enter multiple patterns separated by a comma. You can also specify one or more wildcards in the pattern. For example vol*, *$.

- Exclude site collections matching specified patterns. You can enter multiple patterns separated by a comma. You can also specify one or more wildcards in the pattern. For example vol*, *$.

- Select **Add to the top of scan queue** to add the scans to the top of the scan queue.

4   Select **Start scanning**.

Note: You can use a command line utility, scancli.exe, to further customize the scan, view the scan jobs running on a specified node, or display the scan status for specified site collections. For details, See scancli.exe on page 302.

To enable or disable site collections

1   In the Management Console, click **Settings** > **SharePoint Web applications**.

2   Click the web application that the site collections are a part of.

3   On the filer detail screen, click **Monitored Site Collections**.

4   Click the **Action** drop-down corresponding to a share, and select **Enable** to enable a site collection that has been disabled.

5   Click the **Action** drop-down corresponding to a share, and select **Disable** to disable a site collection.

See "Adding site collections" on page 179.

See "Add/Edit site collection options" on page 180.

Removing a configured web application

If you want to remove an existing SharePoint web application from Data Insight you must complete the steps in the correct order.

To remove a web application from Data Insight

1   Delete the configured web applications from the Data Insight console.

    Deleting the web application enables you to disable auditing for the monitored SharePoint web applications

    See "Deleting web applications" on page 179.

2   On the SharePoint server, disable auditing of the web applications that are deleted from Data Insight.

3   Uninstall the Data InsightWeb service from the SharePoint server.

Configuring containers

This chapter includes the following topics:

- About containers
- Managing containers
- Adding containers

About containers

A container can consist of similar entities such as filers, shares, Web applications, site collections, or DFS paths. Grouping the entities under a single container allows you to easily define the scope of a role assigned to a user.

For example, User1 is assigned the Product Administrator role. You can further define the scope of the role by selecting a container that contains only the filers that you want User1 to access.

Managing containers

You can add containers to Data Insight, view details of the configured containers and delete one or more containers on the Containers listing page.

To manage containers

1   In the Console, click **Settings** > **Containers** to display the Containers details page.

2   The list of configured containers appears.

Adding containers

You must add containers to Data Insight that group the filers, Web applications, shares, site collections or DFS paths, as required.

To add a new container

1   In the Console, click **Settings** > **Container**.

2   On the Containers page, click **Add new container**.

3   On the Add new container screen, enter the container properties, and click **Add new container**.

4   Enter

Add new container/Edit container options

Use this dialog box to add a container to Symantec Data Insight or to edit the configuration of an existing container.

Table 14-1   Add new container/ Edit container options

Field	Description
Container Name	Enter a logical name for the container.
Container Type	From the drop-down, select **Filer/Web Application**, **Shares/Site Collection**, or **DFS paths**.
	Based on the selection, Data Insight filters the list of entities.
	Do the following to select the resources:
	1 Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint Web applications.
	Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain.
	2 The selected data set is listed in the **Selected resources** pane.

Configuring Data Insight product users

This chapter includes the following topics:

- About Data Insight users and roles
- Reviewing current users and privileges
- Adding user
- Editing users
- Deleting users
- Configuring authorization for Symantec Data Loss Prevention users

About Data Insight users and roles

Before a user can log in to Symantec Data Insight, you must add an account for that user. The user can then use that account to log in to the Console. The user account can be any account that is valid on the Management Server system. This includes local system accounts as well as users belonging to the domain which the Management Server is a part of.

When you create an user account, a role (set of access privileges) is associated with the account. Roles specify access privileges to the Symantec Data Insight system. For example, a role might let users view access and permissions data, but prevent them from adding or deleting filers. Data Insight role-based access control governs access to product features and functionality. Roles consist of the user privileges that determine what a user can see and do in the Management Console.

The Data Insight administrator (a user mapped to the predefined Server Administrator role) assigns roles to users. Users can be mapped to one role only. Data Insight ships with predefined roles that you can assign to user accounts.

Table 15-1 summarizes the various Data Insight roles.

Table 15-1   Symantec Data Insight roles

Role name	Description
Server Administrator	Allows the user to perform all actions in the product GUI that includes setting up all infrastructure (including filers, users, and others) and view all the access and permissions data.
Product Administrator	Allows the users to manage filer settings and optionally to view all the access and permissions data for the given filers. Product administrator role, configured for a select set of filers/Web applications, is not allowed to add new filers or delete configured filers.
User	Allows the users to view all the product access and permissions data. Users in this role do not have access to any settings tasks.
Storage User	Allows the users to view storage-related data in the **Workspace** tab, but does not allow them to view permissions data or audit logs. Users in this role do not have access to the **Settings** tab.

Reviewing current users and privileges

You can review the current Data Insight users and the roles assigned to them on the Product Users listing page. On this page you can also review the filers and Web applications that these users are allowed to monitor.

To review current users and privileges

1   In the Console, double-click **Settings** > **Product Users** to display the Product Users listing page.

2   Click the Export icon at the bottom of the page to save the data to a .csv file.

Adding user

This section describes how to add users to Symantec Data Insight.

To add new a Data Insight user

1   In the Console, click **Settings** > **Data Insight Users** to display the Product Users listing page.

2   Click **Add New Data Insight User**.

3   On the Configure new product user page, enter the user properties, click **Add New Data Insight User**.

See "Configure new Data Insight user /Edit Data Insight user options " on page 190.

Configure new Data Insight user /Edit Data Insight user options

Use this dialog box to add a new user to Data Insight, or edit the properties of an existing user.

Table 15-2   Add/Edit Data Insight user options

Field	Description
Username	Enter the username for the user.
Domain name	Enter the name of the domain to which the user belongs.
Role	From the drop-down, select the role you want to assign the user. See Table 15-1 on page 189.
Select view options	From the drop-down, select **Allowed** or **Denied**. Setting this option to **Allowed** enables the user to view the screens on the **Workspace** and **Reports** tabs. This option is only available if the user is assigned the Product Administrator role.
Allow access to Workspace data	Select the check box to enable the user to view the screens on the **Workspace** and the **Reports** tabs This option is only available if the user is assigned the Product Administrator role.

Table 15-2   Add/Edit Data Insight user options *(continued)*

Field	Description
Resources/Containers to grant access to	Select one of the following: ■ **All filers/Web applications (Includes the ones added in the future)** ■ **Selected Filers/Web applications** ■ **Selected Shares/Site Collections** ■ **Selected DFS paths** ■ **Containers** If you select **Selected filers/Web Applications** , **Selected Shares/Site Collections**, **Selected DFS paths**, or **Containers** , the system displays a list of the appropriate configured entity. Use the arrows to select the entities you want the user to monitor. **Note:** A user, assigned the Server Administrator role, has the scope set to **All Filers/Web Applications**, by default. The scope by DFS paths is applicable only for User and Storage User roles.

Editing users

After you add a user to Data Insight, you can edit the user properties. For example, you might need to edit any of the following:

■ The role assigned to the user

■ The view option for the user

■ The filers and/or Web applications that the user is allowed to monitor

To edit the properties of a user

1   In the Console, double-click **Settings** > **Data Insight Users** to display the Product Users listing page.

2   Click the **Edit** button for the corresponding user.

3   On the Edit Data Insight user page, make changes. as necessary, and click **Save**.

Deleting users

You can delete Data Insight users.

To delete an user

1   In the Console, double-click **Settings** > **Data Insight Users** to display the Product Users listing page.

2   Select the user, and click **Delete**.

3   Click **OK** on the confirmation message.

Configuring authorization for Symantec Data Loss Prevention users

Symantec Data Loss Preventions makes Web Services calls into Data Insight to obtain ownership information for sensitive files and folders. However, you must first provision a Data Insight account for Symantec Data Loss Prevention in Data Insight.

You can provision a Active Directory service account OR a local system account and assign it the Server Administrator privilege. Symantec Data Loss Prevention can use this account to access Data Insight data.

Configuring Data Insight product servers

This chapter includes the following topics:

About Data Insight product servers

A Data Insight product server is any server which has Symantec Data Insight software installed. This includes the Management Server, zero or more Collectors, zero or more Indexers, and zero or more Windows File Server agents. You can view information about configured product servers, check the status of running scans, and change advanced settings from the **Settings** tab of the Management Console.

Adding a new Data Insight server

You can add a new Data Insight from the Management Console.

To add a Data Insight server

1   In the Console, click **Settings** > **Data Insight Servers** to display list of configured product servers.

2   Click **Add new server**.

3   On the Add New Server page, enter the following details:

   ■   The host name or IP address of the server.

   ■   The Communication Service port. By default, the Communication Service connects through service port 8383.

   ■   The Configuration Service port. By default, the Configuration Service connects through service port 8282.
       See "Configuring Data Insight services" on page 201.

   ■   Select a credential from the **Select Saved Credential** drop-down. The credential must belong to a user with Administrator privileges on the node where the software needs to be installed.
       See "Managing saved credentials " on page 46.

   ■   The directory in which you want Data Insight to be installed. By default, the destination directory is `C:\Program Files\Symantec\DataInsight`.

   ■   The location where you want to store the product data. Select a location with enough free space and high-performance disks. By default the data directory is located at **C:\DataInsight\data**.

   ■   Depending on your deployment scenario, select **Collector** only or **Indexer and Collector** as the installation option.

4   Click **Install**.

You can view the progress and status of the installation on the Installation Status page.

Managing Data Insight product servers

On the Data Insight servers listing page you can do following tasks:

- View detailed information about all configured servers.

- Add a new Data Insight server.
  See "Adding a new Data Insight server" on page 194.

- Get a list of currently running scans.

- Edit the server's configuration.

- Apply a node template.

- View and install recommended patches for a server node.
  See "Viewing and installing recommended upgrades and patches" on page 217.

- Remotely upgrade and push-install rolling patches on Collector and Windows file server agent nodes.
  See "Deploying upgrades and patches remotely" on page 218.

- Run a specific Data Insight process.
  See "Monitoring Data Insight jobs" on page 214.

- Delete the server.

Use the filter on the left to filter the list of servers based on the health of the server or role of the server. You can also search for a server by entering the name or IP address of the server in the **Filter** text box. The displayed list of servers changes automatically as you enter the term in the **Filter** text box or when you select the check boxes. For example, you can search for all Healthy servers that are assigned the Collector role. Data Insight displays all the Collectors in your environment that are in the Healthy state.

The following parameters determine the health of a Data Insight server:

- Whether a node is online.

- Whether all the required services are running on a node.

- The disk space on the server.

- Whether the node has error files in the `err` folder.

- CPU and Memory usage on the server

The health of a server is not known for five minutes after starting the DataInsightWatchdog service, or if the DataInsightWatchdog service is stopped.

The Servers pie-chart on the **System Overview** dashboard gives a high-level overview of the number of Data Insight severs in Faulted, at Risk, and Healthy state.

See "Viewing the system health overview " on page 25.

To view configured product servers

1   In the Console, click **Settings** > **Data Insight Servers** to display list of configured product servers.

2   Use the provided dynamic search filter to search for configured servers based on the name of the server.

3   Review the following information about the servers:

   ■   The ID of the server.

   ■   The name of the server.

   ■   The role of the server.

   ■   The status of the server — whether the server is online or offline.

   ■   A list of matching templates for this server.

   ■   The version of the Data Insight software that is installed on the server.

   ■   The credentials used to install the server.

   ■   The links to recommended patches and hot fixes.

Data Insight also displays the recommendation to upgrade to the latest version, if available.

To view server events

1   In the Console, click **Settings** > **Data Insight Servers**.

2   Click the **Select Action** drop-down for the corresponding server in the servers listing table, and select **Event Log**.

   Or, on the details page for a server, click **Event Log**.

   The event log for that server appears.

You can create a node template to change one or few settings on multiple nodes. Using node templates is useful when multiple nodes need to inherit the same settings, for example, more number of indexer threads for all indexers in your environment.

See "Managing node templates" on page 221.

To apply a template to a Data Insight node

1   In the Console, click **Settings** > **Data Insight Servers**.

2   Select the server to which you want to apply a template, and from the **Apply a node template** drop-down, select a configured template.

3   Click **Yes** to confirm.

To delete a server

1   In the Console, click **Settings** > **Data Insight Servers**.

2   Click the **Select Action** drop-down for the corresponding server in the servers listing table, and select **Delete**.

Note: Data Insight does not allow you to delete a server, if it is associated with a storage device.

See "About automated alerts for patches and upgrades" on page 217.

See "Configuring advanced settings" on page 203.

Viewing Data Insight server details

You can view server details on the **Overview** tab and perform various tasks, such as, changing th eadvanced settings, monitor in-progress scans, view server statistics, and check event logs for the server.

To review server details

1   In the Console, click **Settings** > **Data Insight Servers**

2   Click the server that you want to review, or click **Select Action** and select **View**.

The **Overview** tab of the server details screen appears. It displays the following information:

- Name
  This is the address of the server configured when the server was added. Remote product servers use this address when communicating with this server. At this time, Data Insight does not support changing the address of the server.

- Roles
  This indicates the roles that the server plays. Possible server role values are Management Server, Indexer, Collector, and Windows File Server Agent.

- Filer Name
  If the server is a Windows File Server Agent, the name of the associated
  file server is displayed here.

- Data Insight version
  Indicates the version of Data Insight installed on this server.

- Operating System
  Indicates the operating system installed on this server.

- CPUs
  Indicates the number of CPUs available on this server.

- Memory
  Indicates the RAM installed on this server in MBs.

- Associated Windows File Server
  This detail is available only if you select the server with the Windows Filer
  Server agent role. Indicates the host name or IP address of the Windows
  File Server that the agent is monitoring.

- Server Health
  Indicates the current state of the server - whether a server is in a healthy
  state, is faulted, or at risk. You can also view the reasons for the current
  health state of the server.

- Composition of Data Directory
  The pie chart shows the disk space utilized by various folders under the
  `data` directory. These folders include `collector`, `indexer`, `console`,
  `workflow`, `attic`, `inbox`, `outbox` and `others`

- Product Updates
  The suggestion for upgrades if a newer version of Data Insight is available.

See "Configuring advanced settings" on page 203.

Monitoring the performance of Data Insight servers

Data Insight enables you to review the performance of the remote Collector and
Indexer nodes from the Management Console. You can analyze various performance
parameters, such as the disk and CPU utilization and backlog of files accumulating
on individual server nodes.

Processes running on the Collector node collect information about file and folder
hierarchy and access events on storage devices, This information is stored in
database files with appropriate timestamps, and is periodically processed and sent
to the Management Server for the storage device.

Performance alerts are triggered when the resource utilization exceeds the threshold values for multiple scan cycles, and the backlog of files being processed on a node starts accumulating. The view enables you to decide whether to reassign or add a Collector, or to make appropriate configuration changes on the **Advanced Settings** page for the server.

See "Configuring advanced settings" on page 203.

To view the performance of the Data Insight servers

1   On the console, click **Settings** > **Performance**.

2   The overview list page displays the configured Data Insight servers.

3   From the **Processing Backlog** drop-down, select the node or the folder for which you want to view the performance statistics.

4   Select the period for which you want to view the data.

     Data Insight displays the following statistics:

     ■   The host name of the server node.

     ■   The size and number of files accumulated on the server.

     ■   The CPU and memory consumption on the server.

Table 16-1 lists the type of files and their location on server nodes.

Table 16-1   Files and their location on Data Insight servers

File type	Server	Location	Cause
Event files	Collector	`<datadir>\collector`	Check the throughput rate of the `collector.exe` process and the number of collector threads.

Table 16-1   Files and their location on Data Insight servers *(continued)*

File type	Server	Location	Cause
Snapshot and audit files	Collector, Windows File Server	`outbox` folder on the Collector node or the Windows File Server	`FileTransferJob` is lagging. The `FileTransferJob` transfers the files from the `outbox`of folder on the Collector to the `inbox`o folder on the Indexer. Check the throughput statistics of `FileTransferJob`.
Snapshot and audit files	Indexer	`inbox` folder on the Indexer node	Check the throughput rate of `idxwriter.exe` process and the number of indexer threads. The throughput of the `idxwriter.exe` process in turn depends on the size of the batch being processed.

See "About the Health Audit report" on page 281.

Viewing in-progress scans

You can view a list of currently running scans on the **In-progress scans** tab.

To review the in-progress scans

1   On the product server details page, click **In progress** scans.

2   Review the following information from the in-progress scans table:

   ■   **Object Name** - Name of the object being scanned.

   ■   **Object Type** - The type of object being scanned. This can be a share, site collection or Active Directory.

- **Task Name** - Indicates the type of the scan.

- **Task State** -Whether the task is RUNNING or IN QUEUE.
  If none of the tasks are in RUNNING state, it usually means a scan pause
  window is in effect. You can configure the pause interval for the server from
  Advanced Settings page for the server. To override the pause schedule
  for a share or site collection and start the scan immediately, from the Action
  drop-down, select **Override pause schedule**.
  See "Configuring advanced settings" on page 203.

- **Start Time** - Time the scan started if it's in RUNNING state.

- **Time elapsed** - Indicates how long the scan has been running.

- **Task Statistics**- Indicates the statistics of the in-progress scans. It shows
  the number of folders scanned and the files or folders scanned per minute.

3   Click **Cancel** to cancel a particular scan or click **Cancel All** to cancel all scans.

Note: To view the in-progress scans across all nodes in your environment, navigate
to the **Settings** > **Scanning** > **In-progress Scans** tab.

Configuring Data Insight services

You can view the status of all Data Insight processes from the **Services** tab on the
Management Console. You can manage the following services from this page:

- Data Insight Web server service - The process runs on the Management Server.
  It starts by default when you start your Management Server.

- Data Insight Communication service - The process runs on each node in a Data
  Insight deployment. This service is responsible for all inter-node communication.

- Data Insight Configuration service - The process that provides interface to
  configuration and other product data that is stored on the local system.

- DataInsightFpolicy service - The process runs on the Collector node or the
  Management Server. This service is responsible for registering the Data Insight
  server with the NetApp filer and enables Data Insight to receive access events
  from the filer.

- DataInsightFpolicyCmod service - The Windows service runs on the Collector
  node. It is responsible for interacting with the NetApp Cluster Management host
  to receive access events from the nodes in the cluster.

- DataInsightCelerra service - The process runs on the Collector Worker node or
  the Management Server. This service is responsible for registering the Data

Insight server with the EMC Celerra filer and enables Data Insight to receive access events from the filer.

■ DataInsightWatchdog service - The process runs on all nodes in a Data Insight deployment and monitors the health of the node. The service also monitors the disk usage on the Windows File Server node and prevents it from running out of disk space by implementing safeguards.

■ DataInsightWinNAS service - The process runs on the Windows File Server. The service receives event information from the Windows File Server filter driver and transfers it to the Collector node that is configured for the filer.

■ DataInsightGenericCollector service - The service runs on the Collector associated with a generic file server. The service collects all incoming events from generic file servers and web API clients, and copies them to a specific folder on the Collector.

■ DataInsightWorkflow service - The service runs only on the Management Server . This service is responsible for managing the lifecycle of various actions initiated from the Management Server.

For detailed information about the Data Insight services, see the *Symantec Data Insight Installation Guide*.

Depending on the type of the filers managed by the Collector, you can enable the FPolicy, EMC Celerra, or Genericcollector service on the server from this page.

To enable or reconfigure the FPolicy or EMC Celerra service

1   On the **Services** tab, click the service that you want to enable on the server.

2   From the **Select saved credential** drop-down, select the credential that the service uses to run. Ensure that the user used to configure the FPolicy service is added to to the Group Policy object with the **Log on as a service** privilege in your Active Directory Domain Controller.

     Note: In case of a NetApp file server, if the file server belongs to a different untrusted domain, select the Local System account to run the DataInsightFpolicy service.

3   If configuring the DataInsightFpolicy service, enter the name of the policy.

4   If configuring the DataInsightCelerra service, select one of the following to specify the location of the server on which the EMC CAVA service is installed:

     ■ EMC CAVA Service is installed locally on this server

- Remote EMC CAVA Server Pool will publish events to this server

5 Click **Configure**.

See "Configuring advanced settings" on page 203.

See "Credentials required for configuring NetApp filers" on page 79.

See "Credentials required for configuring EMC Celerra filers" on page 107.

Configuring advanced settings

You can edit various settings of the Data Insight servers from the **Settings** > **Data Insight Servers** > **Advanced Settings** page.

The advanced settings are divided into the following categories:

- File system settings - Configures how the server scans file systems. Data Insight performs two types of scans on the configured shares:
  - Full scans
    During a full scan, Data Insight scans the complete share. These scans can run for several hours, if the share is very big. Typically, a full scan should be run once for a newly added share. After the first full scan, you can perform full scans less frequently based on your preference. Ordinarily, you need to run a full scan only to scan those paths which might have been modified while event monitoring was not running for any reason. In all other cases, the incremental scan is sufficient to keep information about the file system metadata up-to-date.
    See Table 16-2 on page 205.
  - Incremental scans
    During an incremental scan, Data Insight re-scans only those paths of a share that have been modified since the last full scan. It does so by monitoring incoming access events to see which paths had a create event or write event on it since the last scan.
    See Table 16-3 on page 206.

- Indexer settings - Configures how the indexes are updated with new information. This setting is applicable only for Indexers.
  See Table 16-5 on page 208.

- Audit events preprocessor settings - Configures how often raw access events coming from file servers must be processed before they are sent to the Indexer.
  See Table 16-6 on page 209.

- High availability settings - Configures how this server is monitored.

Each server periodically monitors its CPU, memory, state of essential services, number of files in its inbox, outbox, and err folders. Events are published if these numbers cross the configured thresholds. Also, each worker node periodically heartbeats with the Management Server. The Management Server publishes events if it does not receive a heartbeat from a node in the configured interval.
See Table 16-7 on page 209.

- Report settings - Configures settings for reports.
  See Table 16-8 on page 210.

- Windows File Server Agent settings - Configures the behavior of the Windows File Server filter driver. This setting is applicable only for the Windows File Server Agent server.
  See Table 16-9 on page 211.

- Veritas File System server (VxFS) settings - Configures how Data Insight scans the VxFS filer.
  See Table 16-10 on page 212.

- NFS settings - Configures how Data Insight scans NFS shares.
  See Table 16-11 on page 212.

- SharePoint settings - Configures the duration for which old audit logs are kept on the SharePoint server. Audit logs that are fetched from the SharePoint server are automatically deleted from the Data Insight database. You can disable this feature at the Web application level.
  See Table 16-12 on page 213.

- Troubleshooting settings - Configures settings that aid troubleshooting.
  See Table 16-13 on page 214.

You can configure the advanced settings per node or save commonly used settings as a templates. See "About node templates" on page 221.

To configure advanced settings

1   In the Console, click **Settings** > **Data Insight Servers**.

2   Click the server, for which you want to configure the advanced settings.

3   Click **Advanced settings**.

4   Click **Edit**.

5   Make necessary configuration changes, and click **Save**.

See "Managing node templates" on page 221.

Each of the categories for the advanced settings are described in detail below.

Table 16-2   File system settings - Full scan settings

Setting	Description
Total scanner threads	The Collector can perform multiple full scans in parallel. This setting configures how many full scans can run in parallel. The default value is two threads. Configure more threads if you want scans to finish faster.
Scan multiple shares of a filer in parallel	This setting indicates if the scanner can perform a full scan on multiple shares of the same filer in parallel. The setting is disabled by default.
Maximum shares per filer to scan in parallel	If multiple shares of a filer can be scanned in parallel, this setting puts a limit on the total number of shares of a filer that you can scan in parallel.
Default scan schedule	Specifies how often full scans need to be performed. You can override this setting at a filer or at a share level. By default, full scans are scheduled to repeat last Friday of each month.
Pause scanner for specific times	You can configure the hours of the day when scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the filer. The setting is enabled by default. Scans resume from the point they were at before they were paused.

Table 16-2   File system settings - Full scan settings *(continued)*

Setting	Description
Pause scanner schedule	Configures when scanning should not be allowed to run. By default, scanning is paused from 7a.m to 7p.m, Monday to Friday. You can specify multiple scanner pause schedules for different days of the week. For example, you can choose to pause scanning from 7:00 A.M. to 7:00 P.M. on weekdays and from 7:00 A.M. to 9:00 P.M. on Saturdays and Sundays. To add a scanning schedule: 1 Click **Add**. 2 On the Pausing schedule pop-up, select the time period and the days on which you want to pause scanning. 3 Click **Save**. You can also edit or delete existing scanning schedules.

Table 16-3   File system settings - Incremental scan settings

Setting	Description
Total scanner threads	The Collector can perform multiple incremental scans in parallel. This setting configures how many incremental scans can run in parallel. The default value is two threads. Configure more threads if you want scans to finish faster.
Scan multiple shares of a filer in parallel	The setting indicates whether the scanner can perform an incremental scan on multiple shares of the same filer in parallel. The setting is enabled by default.
Maximum shares per filer to scan in parallel	If multiple shares of a filer can be scanned in parallel, this setting puts a limit on total number of shares of a filer that can be scanned in parallel. The default value is 2.

Table 16-3   File system settings - Incremental scan settings *(continued)*

Setting	Description
Default scan schedule	Specifies how often incremental scans must be performed. By default, incremental scans are scheduled to run at 7:00 P.M. each night. Schedule incremental scans more or less frequently based on how up-to-date you need information in Data Insight to be.
Pause scanner for specific times	You can configure hours of the day when scanning should not be allowed. This setting ensures that Data Insight does not scan during peak loads on the filer. This setting is enabled by default. Scans resume from the point they were at before they were paused.
Pause scanner schedule	Configures when scanning should not be allowed to run. By default, scanning is paused from 7:00 A.M. to 7:00 P.M. Monday to Friday.

Table 16-4   File system settings - Common settings

Setting	Description
Scanner snapshot interval	Scanning a big share can take several hours. The scanner periodically saves information to a disk so that information is visible sooner without waiting for the entire scan to finish. You can configure how often information is saved to the disk by the scanner. By default, the scanner creates a snapshot of new information every 300 seconds (5 minutes). The minimum value you can set for this parameter is 300.

Table 16-5   Indexer settings

Setting	Description
Total indexer threads	The indexer processes incoming scan and access event information for various shares and updates the per-share database. This setting configures how many databases can be updated in parallel. By default 2 threads are configured.
	Specify a larger value for bigger setups where indexer is not able to keep up with incoming rate of information. This is indicated when you observe too many files in the inbox of the Indexer worker node. However, you must ensure that the Indexer has adequate CPU and memory when configuring a higher number of indexer threads. You need approximately 1 GB of RAM per indexer thread.
Limit maximum events processed in memory	By default, the indexer processes all new incoming events in memory before saving the events to the disk. If your are falling short of RAM on your Indexer, you can limit the maximum number of events that the indexer processes in memory before it saves them to the disk.
	Note that specifying a small number makes the indexing very slow.
Reconfirm deleted paths when reconciling full scan information	After Data Insight indexes full scan data, it computes the paths that no longer seem to be present on the file system. Set this option to true to have Data Insight re-confirm if those paths are indeed deleted using an incremental scan before removing them from the index.

Table 16-5 Indexer settings *(continued)*

Setting	Description
Indexer schedule	Specify how often an index should be updated with new information. By default, all new data is consumed once every four hours. Indexer gets better throughput if more information is given to it when indexing. However, if you configure a very high value, new information will not be visible in the Console for a much longer period.
Indexer integrity checking schedule	Data Insight checks the integrity of its databases once a week. If any errors are found in the database, an event is published. You can configure a different schedule if required.

Table 16-6 Audit events preprocessor settings

Setting	Description
Audit events preprocessor schedule	Incoming raw audit events from file servers must be pre-processed before sending them to the Indexer. At this stage, collector.exe applies various heuristics to the raw events and also removes transient events. By default, raw events are processed every 2 hours.
Batch size (MB)	The maximum size of the raw audit event files that a single Collector thread can process. The default batch size is 2 GB.
Total Collector threads	The Collector can run multiple pre-processors in parallel. This setting configures how many instances can run in parallel.

Table 16-7 High availability settings

Setting	Description
Ping timeout (in minutes)	If a worker node does not heartbeat in the specified interval, Management server will publish an event to that effect. This setting is only applicable for the Management Server.

Table 16-7   High availability settings *(continued)*

Setting	Description
Notify when CPU continuously over (percentage)	If CPU used on this server is consistently over the specified percentage, an event is published. (Default value: 90%)
Notify when memory continuously over (percentage)	If Memory used on this server is consistently over the specified percentage, an event is published. (Default value: 80%)
Notify when disk usage over (percentage)	If disk usage, either for the system drive or data drive, is over the specified threshold, an event is published. (Default value: 80%)
Notify when disk free size under (MB)	If the free disk space for the system drive or data drive is over the specified threshold in megabytes, an event is published. (Default value: 500 MB)
Notify when number of files in `err` folders over	If Data Insight is not able to process an incoming file for some reason, that file is moved to an `err` folder. Data Insight publishes an event if number of files in the `err` folder crosses the specified threshold. (Default value: 50)
Notify when number of files in `inbox` and `outbox` folder over	If Data Insight is not able to process incoming data fast enough, the number of files in the transient folders, `inbox` and `outbox`, goes on building up. Data Insight publishes an event if number of files crosses the configured threshold. (Default value: 5000)

Table 16-8   Reports settings

Setting	Description
Maximum memory when generating report output	Specifies the maximum memory that can be used for generating a report output. By default, it is 1024 MB on a 32 bit machine and 2048 MB on a 64 bit machine.
Total threads for generating report output	Configure the number of report outputs that can be generated in parallel. Default value is 2.

Table 16-8   Reports settings *(continued)*

Setting	Description
Total threads for generating report data	By default, Data Insight executes two reports in parallel. However, you can configure a higher value to run multiple reports in parallel.
Maximum reports that can run simultaneously	Specify the number of report instances that can run in parallel. This setting helps you speed up the process of report generation. For a particular Data Insight server, the thread count applies to all types of reports.

Table 16-9   Windows File Server agent settings

Setting	Description
Maximum kernel ring buffer size	The Windows File Server filter driver puts events in an in-memory buffer before the DataInsightWinnas service, consumes them. By default, it uses a 10MB buffer. You can use a bigger buffer. Data Insight publishes an event that indicates events are being dropped due to a high incoming rate. Note that this buffer is in kernel and is limited on a 32 bit operating system.
Ignore accesses made by Local System account	The Windows File Server filter driver ignores accesses made by processes running with Local System account. This setting ensures that Data Insight can ignore most events originating from the operating system processes or other services like antivirus and backup. Clear this check box to enable monitoring accesses made by LOCAL SYSTEM account. This is not recommended on a production file server.

Table 16-10   Veritas File System server settings

Setting	Description
Flush events on VxFS filer before audit	Set this option to true, if you want to force VxFS to flush its events to disk each time Data Insight requests for information. This option is useful in Proof-of-Concept (POC) setups and enables you to see events faster.
Maximum number of audit threads	This option determines how many filers to fetch audit information from in parallel.
Maximum kernel ring buffer size (Number of records)	The access event records are saved in a log file on the VxFS filer before Data Insight consumes them. By default, 50,000 records can be saved in the log file. You can also specify a larger number. Data Insight publishes an event that indicates that events are being dropped due to a high incoming rate.

Table 16-11   NFS settings

Setting	Description
Set default credentials for NFS scanner	Set this option to true if you want to allow Data Insight to use the specified User and Group ID to log in to scan NFS shares.
User ID	The ID of the NFS user that the Data Insight uses to scan the filer. You can set the value to 0 to allow root access from the Data Insight scan hosts.
Group ID	The ID of the group that the Data Insight uses to scan the filer. You can set the value to 0 to allow root access from the Data Insight scan hosts.

Table 16-12 SharePoint settings

Setting	Description
Automatically delete audit events from SharePoint server that are older than (days)	When configuring a SharePoint Web application, you can choose to let Data Insight delete audit logs that have already been fetched from SharePoint. By default, Data Insight deletes audit logs older than two days. Deletion of audit logs takes place every 12 hours.

You can change the interval using this setting. |
Schedule to fetch audit events from SharePoint server	Data Insight fetches new audit events from SharePoint periodically. By default, it does so every 2 hours. You can configure a different schedule.
Total scanner threads	The Collector can perform multiple full scans in parallel. This setting configures how many full scans can run in parallel. The default value is 2 parallel threads. Configure more threads if you want scans to finish faster.
Scan multiple site collections of a web application in parallel	This setting indicates if the scanner can perform a scan on multiple site collections of the same web application in parallel. The setting disabled by default.
Maximum site collections per web application to scan in parallel	If multiple site collections of a web application can be scanned in parallel, this setting puts a limit on the total number of site collections of a web application that you can scan in parallel
Default scan schedule	Specifies how often scans need to be performed. You can override this setting at a web application or site collection level. By default, scans are scheduled to repeat 11:00 p.m. each night.
Pause scanner for specific times	You can configure the hours of the day when scanning should not be allowed. This ensures that Data Insight does not scan during peak loads on the SharePoint servers. The setting is enabled by default. Scans resume from the point they were at before they were paused.

Table 16-12   SharePoint settings *(continued)*

Setting	Description
Pause scanner schedule	Configures when scanning should not be allowed to run. By default, scanning is paused from 7:00 a.m to 7:00 p.m, Monday to Friday.

Table 16-13   Troubleshooting settings

Setting	Setting
Preserve intermediate files	As new data comes into a Data Insight system, it moves between various modules. In this process the original files are deleted and a new processed file is generated for the next stage of processing.

To aid troubleshooting, select this check box to retain the intermediate data files. These files get stored in `attic` folder in the data directory. |
| Preserve raw audit event files | Events processed by the Audit Pre-processor stage are deleted once consumed. If this setting is enabled, raw audit event files will be preserved in the `attic` folder in the data directory. |

See "Managing Data Insight product servers" on page 195.

Monitoring Data Insight jobs

You can monitor and execute the jobs that run on remote Data Insight servers from the Management Console. The view enables you to search for the jobs are part of the communication Web service, the DataInsightWatchdog service, and the DataInsightWorkflow service. These jobs either run continuously or are scheduled.

For more information about the different Data Insight services, see the *Symantec Data Insight Installation Guide*.

Note: You can view the SPEnableAuditJob and the SPAuditJob only if the server is configured to be the Collector for a SharePoint site collection.

To view and execute jobs from the console

1   On the Management Console, click **Settings** > **Data Insight Servers**.

2   On the list page, click **Select Action** >**View** for the server on which you want to view or execute a process or a job.

    The overview page for the server opens by default.

3   Click the **Jobs** sub-tab.

    The page lists all jobs that are scheduled on the server. Review the following details about a job:

    ■   The status of the job, whether stopped or running.

    ■   The date and time when the job was previously run.

    ■   The date and time of the next scheduled run.

    ■   The schedule of the job. For example, every day, every hour, every one minute.

    ■   The service the job runs on, such as the Communication service or the DataInsightWatchdog service.

4   To execute a job, do one of the following :

    ■   On the Jobs sub-tab, click **Select Action** > **Run** to run the job without waiting for the next scheduled run.

    ■   On the **Data Insight Server** list page, click the **Run Job** drop-down, and select the job you want to execute.

See

Viewing Data Insight server statistics

You can view the line graphs that indicate the health of the Data Insight servers. The DataInsightWatchdog service collects the server statistics. You can use the information to know the performance trends for each Data Insight server. For more information about the DataInsightWatchdog service, see the *Symantec Data Insight Installation Guide*.

The line graphs display hourly, weekly, monthly, and yearly data.

To view server statistics

1   Click **Settings** > **Data Insight Servers** > *Server Name* > **Statistics**.

2   Select the following to filter the data that is displayed on the page:

    ■   The charts that you want to view.

- The duration for which you want to view the data.

- The type of statistics that you want to view - Average, Minimum, and Maximum.

3 The **Statistics** sub-tab displays the following high-level data:

- The number of files in the `inbox` and the `outbox` folders.

- The number of error files in the `scanner/err`, `collector/err`, and `indexer/err` folders.

- The CPU and memory usage on the Data Insight servers.

- The disk space utilization on the Data Insight servers on the system disk, Data Insight installation disk, and the Data Insight `data` directory.

- Incoming event rate for this server if it is a collector for one or more filers.

- Throughput for the event pre-processor.

- The raw event files being processed by the Collector, that is, the number of events that are processed by `collector.exe` per second.

- The scan and audit files being processed by the Indexer

- Internal event files being processed by the Management Server.

4 You can view each processing backlog chart from three perspectives. Click one of the following aspects of the backlog:

- The size of the backlog in MB or GB; the total size of all files that are yet to be processed.

- The count of files that are yet to be processed.

- The time lag in terms of hours or days; the time lag is the difference between the current time and the file with the oldest timestamp that is yet to be processed.

5 For the Collector and Indexer backlogs, click the drop-down to the right of the chart to view the Top 10 objects that are contributing to the backlog. In case of Collector backlog, objects are the filers or web application and case of Indexer backlog, objects mean the shares or site collections

The Top ten chart is a bar chart.

About automated alerts for patches and upgrades

Data Insight simplifies the task of installing upgrades and patches by providing you with automated alerts and suggestions. Data Insight fetches this information from Symantec Operations Readiness Tool (SORT) to help you keep track of the product updates applicable to your installed version.

For Data Insight to communicate with SORT, the Data Insight Management Server must have an active Internet connection to the web site, *https://sort.symantec.com/*.

Data Insight displays the following upgrade recommendations for each of its product servers:

- Rolling Patches (RPs) that are available for the installed version. Only the latest rolling patch is displayed.

- Recommendation for a new product version appears at the footer of the Data Insight servers page.

See "Viewing and installing recommended upgrades and patches" on page 217.

Viewing and installing recommended upgrades and patches

You can view the recommendations regarding patches and the product version upgrades from the **Data Insight Servers** page. Using the download links you can download the recommended patched and install them on your **Data Insight Servers**.

To view and install a patch for your product server:

1   In the Management Console, navigate to **Settings** > **Data Insight Servers** to display a list of configured product servers.

2   Data Insight displays the recommendations for the patches under the **Product Updates** column. The recommendation for upgrading the product version is displayed on the footer of the page.

> **Note:** The **Product Updates** column is displayed only when the Data Insight Management Server is able to connect to the SORT website. When there is no connection, an error message is displayed in the footer.

3   Click the link to the latest patch that Data Insight recommends. You will be redirected to the SORT website.

4   Download the patch from the **Downloads** page on the website.

5   You can refer to the README on the page for the installation instructions and to verify the problems that have been fixed in the patch.

Deploying upgrades and patches remotely

You can remotely deploy installers and patches on the Data Insight worker nodes and Windows file server agents from the Data Insight Management Console. This simplifies the task of upgrading and configuring numerous nodes individually in a large Data Insight deployment. You can do the following:

- Install rolling patches on the Indexer nodes, the Collector nodes, and the Windows file server agents.

- Deploy the product installer on the Collector nodes and the Windows file server agents.

Note: Remote upgrade of an Indexer node is not supported. You must upgrade an Indexer node manually.

You can either use existing saved credentials for upgrading a node or create new credentials.

You can also perform all the remote deployment actions using the `installcli.exe` utility from the Windows command prompt. For detailed information on `installcli.exe`, see the *Command File Reference*.

To remotely deploy upgrades and patches

1   In the Management Console, navigate to **Settings** > **Data Insight Servers** to display a list of configured product servers.

2   Select the server node for which you want to remotely deploy patch or upgrade.

3   Click the **Install** drop-down.

4   Do any of the following:

- Click **Install Rolling Patch** for installing a rolling patch on the selected node.

- Click **Upgrade** for installing a product version upgrade.

You can view the progress of the remote deployment operation from the **Installation Status** page.

See "Viewing the status of a remote installation" on page 220.

Using the Upload Manager utility

Use the Upload Manager to upload agent bundle zip files, var files, and patch installers on the worker nodes.

Before you can install the agent, ensure that the Windows File Server agent packages are uploaded on the relevant Collector nodes. You can use the Upload Manager utility to upload the agent packages to the Collector nodes in your Data Insight configuration.

To upload the agent packages

1   In the Console, click **Settings** > **Upload Manager**.

2   Browse to the location where the agent packages are saved.

3   Select the Collector nodes on which you want to upload the packages.

4   Click **Upload Bundle**.

The agent installation bundle is a zip file that contains the agent installer and various installation template files. There is one bundle for each processor architecture. You must upload the appropriate bundles to the Collector worker nodes based on the architecture of your file servers. The bundles are available along with the main install media and have the name, Symantec_DataInsight_windows_winnas_4.5_XXX_arch.zip. You can customize the agent installation by extracting the bundle in a temporary location, editing the installation templates as required, recreating the zip bundle, and then uploading the updated bundle to the appropriate Collector nodes using the Upload Manager utility.

To install a rolling patch, upload the rolling patch executable to the Management Server node. To upgrade Collector nodes, upload the newer version of the product installer to the Management Server.

Note: When a new version is installed on the Management Server, the installer automatically copies itself to the `installers` folder of the Management Server. In such cases, you do not need to separately upload the package to the Management Server for upgrading other worker nodes (except the Windows File Server agents).

Note: Remote install for Linux indexers is not supported.

About migrating storage devices across indexers

Every storage device that is configured in Data Insight has exactly one Indexer node associated with it. The Indexer node processes the audit data from the Collector to service queries from the Management Server.

You may consider migrating a storage device to another Indexer in the following situations:

- If the existing indexer for the storage device displays high resource utilization and the backlog of files being processed starts accumulating on the node. Migrating to another indexer helps balance the load on the existing Indexer. See "Monitoring the performance of Data Insight servers" on page 198.

- If the existing indexer is being decommissioned.

- If you want to migrate from a Windows server to a Linux server.

Before you decide to migrate the storage device to another Indexer, ensure the following:

- Both the source and the destination indexers are in good health and can communicate with each other using the fully qualified domain name.

- The destination Indexer has enough disk space, CPU, and memory capacity to take the additional load.

- The Collector nodes that are associated with the storage device have enough disk space, because the files may accumulate on the Collector nodes during the migration.

- The migration should take place during quiet hours, so that additional files for consumption are not flowing in rapidly.

See "Editing filer configuration" on page 158.

See "Editing web applications" on page 178.

Viewing the status of a remote installation

You can view the progress of a remote installation of a Data Insight server or a Windows File Server agent, whether you have initiated the installation from the Management server or by using the command line utility.

To view the status of a remote installation

1   From the Data Insight Management Console, navigate to **Settings** > **Installation Status**.

2   The progress of the ongoing operations is displayed along with the following information:

- The node for which the installation was initiated.

- The status of the installation.

- Time when the installation was initiated.

3   Click **View Progress** to view a more detailed status of the install operation.

Configuring node templates

This chapter includes the following topics:

- About node templates
- Managing node templates
- Adding or editing node templates

About node templates

A node template consists of a set of pre-defined node settings. To ensure consistency of configuration across all Data Insight server nodes, you can create templates with the required settings and apply them when configuring the servers.

Data Insight comes bundled with three standard templates which include frequently used settings to help you set up nodes for well known configurations, for example, templates for POC configuration, large Indexer, and large Collector. The standard node templates are stored in the `INSTALL_DIR\conf\node_template` folder. You cannot edit or delete the standard node templates from the console.

You can apply multiple templates when configuring a server. However, when you apply multiple templates to a node, Data Insight applies each template serially when evaluating the configuration for that node. You can also use a node template to change a single setting on all nodes. For example, if you want to change the **Ping timeout** setting for all nodes, you can create a template with the required timeout setting and apply to all Data Insight nodes in your environment.

Managing node templates

You can view configured templates, create new node templates, edit existing templates, and delete a template from the **Node Templates** list page.

To view configured node templates

1   In the Console, click **Settings** > **Data Insight Servers**.

2   From the **Node Templates** drop-down, select **Manage Node Templates**.

3   The list page displays all standard and configured node templates.

You can choose to delete configured node templates.

To delete a node template

1   In the Console, click **Settings** > **Data Insight Servers**.

2   From the **Node Templates** drop-down, select **Manage Node Templates**.

3   On the node templates list page, select the template that you want to edit, and from the **Select Action** drop-down, select **Delete**.

4   Click **OK** on the confirmation dialog to delete the node template.

See "Adding or editing node templates" on page 222.

Adding or editing node templates

You can create new templates that you can apply to multiple nodes simultaneously or edit existing templates to change the configuration settings.

To create a node template

1   In the Console, click **Settings** > **Data Insight Servers** .

The Data InsightData Insight Servers list page displays.

2   From the **Node Templates** drop-down, select **Manage Node Templates**.

3   Click **Add New Node Template**.

4   In the node template name field , enter a unique name for the template.

5   For each configuration category, click **Edit**,.

Select the check box for the settings that you want in the template, and configure appropriate values for the settings, as required.

See "Configuring advanced settings" on page 203.

6   Click **Close Configuration** to save the settings as a node template.

To edit a node template

1   In the Console, click **Settings** > **Data Insight Servers**.

2   From the **Node Templates** drop-down, select **Manage Node Templates**.

3 On the node templates list page, select the template that you want to edit, and from the **Select Action** drop-down, select **Edit**.

4 On the **Edit Node Templates** page, change the required settings, and click **Close Configuration** to save the changes.

When you edit a node template, the changes in configuration do not automatically reflect on the Data Insight servers on which the template is applied. You must apply the modified template to the Data Insight server again for the configuration changes to take effect on the server.

Configuring remediation settings

This chapter includes the following topics:

About configuring permission remediation

Data Insight provides permission recommendations on paths based on activity on the paths. To fine tune these recommendations and take action on the recommendations, the Data Insight administrator must enable and configure permission remediation. Depending on your organization process, you can configure Data Insight in any of the following two ways:

- Raise a remediation ticket.
  Configure when you have a ticketing system, which can assign tickets to relevant stakeholders. The recipients are responsible for the actual implementation of the recommendation. You can create remediation tickets by using one of the following two ways:

  - By sending an email to the ticketing system.
    Use this option if you have a ticketing system which can create a ticket by reading an email.

  - By executing a custom script.

Use this option if you have an alternate mechanism of creating a ticket.

The scripts can be created in the *.exe*, *.bat*, *.pl*, or *.vbs* formats.

■ Apply changes by invoking custom scripts.

Configure custom scripts to enable Data Insight to directly apply the recommended changes.

You can use the following types as custom scripts,.*exe*, *.bat*, *.pl*, or *.vbs*.

You can view the status of remediation actions on the **Settings** > **Action Status** tab of the Data Insight Management Console.

For information about custom scripts, see the *Symantec Data Insight Programmer's Reference Guide*.

See

Managing and configuring permission remediation

You can configure Data Insight to handle the way it implements the recommended permission changes. You must have a Data Insight Server Administrator role to be able to configure the permission settings.

To enable permission remediation

1   From the Data Insight Management Console, click **Settings** > **Permissions**. The **Remediation** sub-tab opens by default.

2   Click **Edit**. The page expands to display the configuration for permission remediation.

3   Select **Enable Permission Remediation**.

To configure remediation for raising a ticket

1   From the Data Insight Management Console, click **Settings** > **Permissions**. The Remediation sub-tab opens by default.

2   Click **Edit**. The page expands to display the configuration for permission remediation.

Select **Enable Permission Remediation** if it is not already enabled.

3   Select **Raise a ticket**. The panel expands to display the configuration details.

4   Select either of the two options:

■   **Send email** - Select to configure settings for an email-based ticketing system.

■   **Use custom action** - Select to configure settings for a non-email based ticketing system.

5   If you selected the **Send email** option, provide the relevant information in the email template:

- The email ID of the sender

- The email IDs of the recipients

- The email IDs of other recipients

- The subject line

- The header line showing priority and the queue status. The headers can be set to have custom information sent to the ticketing/request systems. For example, you can set priority=high, assign_to=permission_queue etc.

- The body of the email. You can use the default variables to enter relevant text. The variables are evaluated during run-time and are replaced by their corresponding values. Currently Data Insight provides the following dynamic variables:

  - Recomendation_text
    The value of this variable is the recommendations generated by Data Insight.
    For information about reviewing permission recommendations, see the Symantec Data Insight User's Guide.

  - Requester_name
    The value of this variable is the user who accepted the Data Insight recommended changes.

  - Action_id
    Data Insight generates the value of this variable. It is a unique identifier for the operation.

  - Todays_date
    The value of this variable is the system date.

6   If you selected **Use custom action**, do the following:

- Create a custom script by following the guidelines documented in the *Symantec Data Insight Programmer's Reference Guide*.

- Save the script at the location:
  `$datadir\conf\workflow\steps\permission_remediation\ticketing.`

7   In the **Enter the command to be executed** field, provide the file name of the saved script.

8   Select the relevant saved credential if your system needs to run the script using the specified credentials. The script runs with the Local System account credentials, however network calls made by the script will impersonate the specified user credential.

9   Click **Save**.

To configure the process of applying recommendations

1   Write the relevant scripts to handle changes to the following:

   ■   The Active Directory.

   ■   CIFS permissions.

   For more information about the custom scripts refer to the *Symantec Data Insight Programmer's Reference Guide*.

2   Save the scripts in the following locations:

   ■   For changes to Active Directory -
       `$DATADIR\conf\workflow\steps\permission_remediation\AD`

   ■   For changes to CIFS permissions -
       `$DATADIR\conf\workflow\steps\permission_remediation\CIFS`

3   From the Data Insight Management Console, click **Settings** > **Permissions**. The **Remediation** sub-tab opens by default.

4   Click **Edit**. The page expands to display the configuration for permission remediation.

5   Select **Enable Permission Remediation** if it is not already enabled.

6   Select **Remediate using custom scripts**. The panel expands to show you the configuration details.

7   In the **Enter the command to be executed** field, specify the file name of the custom script(s) that you have created in step 1

8   Click **Save**.

   The saved scripts are used to handle the permission remediation actions after you accept the permissions recommendations displayed on the **Workspace** tab.

For information on reviewing recommendations and initiating the process of applying them, see the *Data Insight User's Guide*.

Configuring exclusions for permission recommendation

You can specify the users and groups that you want to exclude from the purview of the permission recommendation. Once you exclude a user or group, Data Insight

does not consider that group for presenting a recommendation for permission changes. You can exclude a groups by directly selecting it. You can also exclude large groups which have more than the specified number of users.

To exclude a user or group from remediation

1   From the Data Insight Management Console, click **Settings** > **>Permissions**. The **Remediation** tab opens by default.

2   Click **Recommendation**.

3   Do the following:

   ■   To exclude groups with more than a certain number of users, specify the value in the space provided.
       To exclude specific groups, in the **Exclude following groups from recommendations** pane, click the group's name which you want to exclude.

   ■   To exclude a specific user, in the **Available Members** pane, select the user.

   You can use the name filters and domain filters to view and sort the available user groups.

   The users and groups that you select are displayed in the **Exclusion List** pane.

4   Click **Save**.

About managing data

The storage devices in your environment may accumulate data that is orphan or has not been accessed for a long time. A large amount of such data on your storage devices can consume valuable storage space. Data Insight enables you to reclaim storage space occupied by inactive data. You can manage the inactive data directly from the Data Insight Management Console in the following ways:

■   Archive data using Symantec Enterprise Vault.

■   Manage data by invoking custom scripts.
    See "Using custom scripts to manage data" on page 233.

About configuring archive options for Enterprise Vault

You can handle archiving and maintenance of the data that is stored on network shares using Symantec Enterprise Vault™ . This feature enables you to identify old and inactive data residing on storage devices and to archive it directly from the Data Insight Management Console. To perform an archive operation on a file, you must have relevant read-write permissions on that file.

To ensure proper functioning of the archive operations, you must have Microsoft .NET Framework 3.5 running on the Data Insight Management Server. Ensure that you restart the Management Server after you have finished installing the Microsoft .NET Framework 3.5.

Enterprise Vault (EV) must be successfully configured before you can archive data from the Data Insight Management Console. To configure Enterprise Vault complete the following tasks:

- Run the Enterprise Vault Configuration Wizard to create Enterprise Vault Directory Database and Enterprise Vault site.

- Add indexing locations.

- Add Vault Store Groups and Vault Store Partitions to store archived items. When configuring a Vault Store, if it makes sense for your organization, set the **Remove safety copies** option to **Immediately after archive**. Selecting this option ensures that the post-processing actions, such as creation of a placeholder or deletion of file are performed immediately after the file is archived. This is especially useful during Testing and Proof of Concept deployments.

- Modify the File System Archiving (FSA) task properties, if necessary.

- Add the necessary retention categories.

- Create archiving policies or modify the Default FSA Volume Policy and Default FSA Folder Policy, if necessary.

- Add archiving targets for NetApp filers and Windows file servers. When adding a Windows File Server as a target, select the option to install the placeholder service on the file server.

- Add volumes and folders for the targets, and create the necessary archive points.

For detailed instructions on completing the Enterprise Vault configuration tasks, see the Symantec Enterprise Vault™ documentation.

Note: If there are multiple Enterprise Vault servers in an Enterprise Vault site, then only one of the EV servers in the site must be added to the Data Insight configuration. If any of the Enterprise Vault servers is down, archiving of files from the shares that use vault stores that are managed by that Enterprise Vault server fails. For information about Enterprise Vault sites and vault stores, see the *Symantec Enterprise Vault™ Introduction and Planning Guide*.

For instructions on initiating archive requests, see the *Symantec Data Insight User's Guide*.

Adding new Enterprise Vault servers

You can configure Enterprise Vault servers to archive inactive data directly from the Data Insight Management Console.

To add a new Enterprise Vault server

1   In the Data Insight Management Console, click **Settings** > **Data Management**.

The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of servers which are already configured.

2   Click **Add New EV Server**.

3   In the **New EV Server** window, provide the following information:

  ■ The IP address of the Enterprise Vault server.

  ■ The port at which the Enterprise Vault server runs.

  ■ The relevant login credentials for the Enterprise Vault server. From the **Login Credentials** drop-down, select the credentials which are saved in the system. See "Managing saved credentials " on page 46.

4   Click **Test Credentials** to verify that Data Insight can connect to the server using the saved credentials.

5   Click **Save**.

Managing Enterprise Vault servers

You can do the following tasks on the **Data Management** page:

■ Review the configured Enterprise Vault servers

■ Add a new Enterprise Vault server.
   See "Adding new Enterprise Vault servers" on page 230.

■ Edit the configuration of an Enterprise Vault server.

■ Configure the archive options.

■ Configure the schedule to pause an archive operation.

To view and manage the existing Enterprise Vault servers

1   From the Data Insight Management Console, click **Settings** > **Data Management**.

The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of configured servers.

2   Select the server for which you want to edit the configuration, and from the **Actions** drop-down, select **Edit**.

3 In the **Edit EV Server** dialog, make necessary changes.

4 Click **Save**.

You can configure additional options, such as the total size and number of the files and folders that can be archived in one archive request.

To configure archive options for Enterprise Vault servers

1 From the Data Insight Management Console, click **Settings** > **Data Management**.

The **Archiving (Enterprise Vault Configuration)** tab opens by default. It displays a list of configured servers.

2 On the **Archive Options** panel, specify the preferred batch size in MB(s). When archiving files to Enterprise Vault, the batch of files sent to Enterprise Vault in one call does not exceed the given size.

3 Enter the number of files that you want to archive in one operation. By default, you can archive 50 files in one archive request.

4 Click **Save**.

Note: The batch size has a higher priority than the file count for deciding the list of files in an archive operation. Thus, Data Insight limits files in the archive operation after the batch size limit is reached, even if the file count does not exceed the specified limit.

You can configure a pause window for the Enterprise Vault operations by scheduling Data Insight to pause all the archive activities during a specific duration of time. When the pause occurs, Data Insight submits no more new archive requests. It places all new requests in a queue and executes them after the pause window.

To configure a pause schedule for the Enterprise Vault operations

1 From the Data Insight Management Console, click **Settings** > **Data Management**.

2 On the **Pause Workflow Schedule** panel, select **Pause workflow for specific times**. Data Insight displays a list for all the previously configured pause schedules.

3 Do any of the following:

- To change an existing pause schedule, click the name of the schedule and click **Edit**

- To add a new pause schedule, click **Add**.

4 Provide the following information:

- The start time of the pause.

- The end time of the pause.

- The days of the week for which you want to schedule the pause.

5   Click **Save**.

Mapping file server host names

A filer that is assigned a specific host name in Data Insight could be assigned a different host name in an Enterprise Vault server. To resolve this conflict, you must map the host name assigned to a file server in Data Insight to the host name assigned by the Enterprise Vault server.

To map the host names of a file server

1   In the Data Insight Management Console, click **Settings** > **Data Management**.

The **Archiving (Enterprise Vault Configuration)** tab opens by default.

2   Click **Add Filer Mappings**. The **Filer Mappings** page displays a detailed list of all the file servers that are configured in Data Insight. The list also displays some already mapped file servers.

For each filer configured in Data Insight, the **Enterprise Vault Filer** field displays the corresponding mapped filer in Enterprise Vault. For a filer that has no mapping, Data Insight attempts to automatically map it with the corresponding filer in Enterprise Vault using host name matching. If Data Insight does not find an automatic match, it does not display an entry. In this case, you must manually map the filers.

3   For each file server displayed in the **Filer** column, verify if its Data Insight host name is correctly mapped to its Enterprise Vault host name.

4   If an existing mapping is incorrect, then enter the correct value for the host name in the **Enterprise Vault Filer** field. Data Insight populates the **Enterprise Vault server** field based on the Enterprise Vault Filer selected.

5   Click **Save**.

Note: Clustered Windows File Servers are added to Data Insight using the name of the cluster. Enterprise Vault requires that virtual file servers configured in the cluster must be added to the Enterprise Vault configuration. Tto enable EV to archive paths on the virtual file servers, Data Insight automatically maps the virtual file servers in the Windows cluster to the configured Enterprise Vault server.

Using custom scripts to manage data

Data Insight enables you to archive inactive data on your storage devices using Symantec Enterprise Vault. However, if you use other archiving tools, or if you want to take actions such as copy your data to a cheaper storage, or delete orphan or inactive data, you can write custom scripts to manage the data.

You can initiate up to two custom actions directly from the Data Insight Management Console. You can apply the scripts to run on the following data:

- The files that are listed under **Workspace** > **Folder Activity** > **Inactive Subfolders** tab.

- The files that are listed in the following types of reports:

  - Access Details reports

  - Access Summary reports

  - DQL reports

  - Data Lifecycle reports

- The paths displayed in the **Workspace** > **ContextMap** view.

Before you can configure Data Insight to run the custom action scripts, do the following:

- Define the action that you want to perform on the data. For example, you may choose to upload all inactive data on your storage devices to the cloud.

- Write a script to perform this action.
  For more information, see the *Symantec Data Insight Programmer's Reference Guide*.

- Place the script at a specified location on the Management Server. By default, all custom scripts must be placed at
  `$datadir\conf\workflow\steps\CUSTOMACTION\scripts`. Data Insight invokes the scripts form this location when you initiate an action from the Management Console.

To configure a custom action script

1   In the Management Console, click **Settings** > **Data Management**. The Archiving (Enterprise Vault Configuration) page displays by default.

2   Click **Custom Action 1** or **Custom Action 2**.

3   On the **Custom Action** page, enter the following details:

    ■ The name of the custom action.

    ■ The name of the script, for example, `copy.pl`.

    ■ The credentials of the user to run the script. You can either use Local System account credentials or the credentials of a user with privileges to perform the desired action on the data. The script continues to run with Local System account, however the specified credentials are used for any network calls made by the script.

4   Select **Do not expand paths** to apply the action defined in the script to the paths selected in the view or the report. The selected paths are passed as-is to the custom script.

5   Select **Expand paths** to apply the action defined in the script to all child folders under the selected folder recursively. If you select this option to invoke an action on the folder, Data Insight passes individual files present in that path's hierarchy to the script, instead of the parent folder.

6   Select the additional data that you want to pass to the script.

7   Click **OK** to save the settings.

Viewing and managing the status of an operation

You can track and manage the progress of the operations that are initiated from the Data Insight Management Console. You can perform the following tasks from the **Action Status** page.

■ View the progress of an operation.

■ Cancel an ongoing operation.

■ Re-run a completed or canceled operation.

■ Delete a completed or canceled operation.

You can view the progress of the following types of operations:

■ Archive operations using Enterprise Vault.

■ Permission remediation operations.

- Custom Action triggered operations.

The **Action Status** page displays the following information:

- The unique identification number of the operation.

- The system-generated name for the triggered operation indicating its origin.

- The type of the operation. For Enterprise Vault operations, the type is specified as *EV*. For permission remediation operations, the type is *PR*. For custom action operations, the type is *CUSTOM*.

- The time when the user triggered the operation from the Management Console.

- The user who triggered the operation.

- The time when the destination server starts processing the triggered request. The destination server is an external server which is responsible for the actual execution of an operation that is triggered from the Data Insight Management Console. For example, in the case of an archive operation, the destination server is the Enterprise Vault server.

- The time when the destination server completes processing the request.

- The time it takes to complete the operation.

- The status of the operation.

Note: Only some columns are displayed in the default view. You can view any other columns by selecting them from the column header drop-down.

The **Details for Action** panel shows you the step-by-step break-down of the selected operation.

To view the status of an operation

1 In the Management Console navigate to **Settings** > **Action Status**. The **Action Status** page displays the details of recently triggered operations.

2 Use the check box filter to display the operations based on their **Type** or **Status**. Additionally, you can use the search facility to display the operations based on their attributes such as **Origin** , **Type**, or **Status**.

3 Click the **Origin** of the selected operation to view granular details of an operation. Alternatively, click the **Select Action** drop-down, and select **View**.

4   The details of the selected operation are displayed in the **Details for Action**
     panel.

5   Use the check box filter to display the operations based on the attributes such
     as: **Status** or **Filer**. Additionally, you can use the search facility to display the
     details based on attributes such as **Path** or **Status**.

You can cancel an operation that is in progress. Cancelling an operation, pauses
all the activities of the operation. You can re-run a canceled operation later.

To cancel an ongoing operation

1   In the Management Console, navigate to **Settings** > **Action Status**. The **Action
     Status** page displays the details of recently triggered operations.

2   Use the check box filter to display the operations based on their **Type** or **Status**.
     Additionally, you can use the dynamic filter to display the operations based on
     their attributes such as **Origin** , **Type**, or **Status**.

3   Click **Select Action** for the operation you want to cancel.

4   Click **Cancel**.

You can re-run a canceled or a completed operation.

To re-run a canceled or a completed operation

1   In the Management Console navigate to **Setting** > **Action Status**. The **Action
     Status** page displays the details of recently triggered operations.

2   Use the check-box filter to display the operations based on their **Type** or **Status**.

3   Click **Select Action** for the operation you want to re-run.

4   Select **Run Again**.

5   Select any of the following:

     ■   **All** - To run all the sub-steps for the operation.

     ■   **Unsuccessful** - To run all the failed sub-steps for the operation.

     Note: For a permission remediation or custom action-related operation that is
     canceled, the option to run the unsuccessful steps again is not available.

You can delete an operation that is canceled or completed.

To delete a canceled or a completed operation

1   In the Management Console navigate to **Setting** > **Action Status**. The **Action
     Status** page displays the details of recently triggered operations.

2   Use the check-box filter to display the operations based on their **Type** or **Status**.

3 Click **Select Action** for the operation you want to delete.

4 Select **Delete**.

Configuring remediation workflows

This chapter includes the following topics:

- About remediation workflows
- Prerequisites for configuring remediation workflows
- Configuring Self-Service Portal settings
- About workflow templates
- Managing workflow templates
- Creating a workflow using a template
- Managing workflows
- Monitoring the progress of a workflow

About remediation workflows

In large storage environments, it can become difficult to assign the responsibility of remediating data resources to data owners and custodians. Security and storage administrators have to manually inform data owners about issues with the resources that they own, and it can be tedious to track remediation actions on such resources

Remediation workflows provide an easy way to fan out remediation tasks among configured custodians and data owners. The custodians are responsible for the data resources and can take a decision about the best way to remediate them. To understand how custodians are assigned in Data Insight, refer to the *Symantec Data Insight User's Guide*.

You can use workflows to define a process to distribute remediation tasks to custodians. You can create different workflows for different remediation tasks. Data Insight workflows enable custodians to address the following use cases:

- Review the user permissions on the folders that the custodians own and attest the permissions or suggest changes.

- View policy violations and take action on the files that violate Data Loss Prevention (DLP) policies. The actions are Smart Response rules defined by DLP administrators. DLP uses the Smart Response rules to remediate the resources that violate configured DLP policies.
  Data Insight uses two DLP Web services for incident remediation - the Response Rules Listing Service and the Response Rule Execution Service. The Response Rule Listing Service provides a list of available response rules, such as delete or quarantine, for a given incident. The Response Rule Execution Service takes the response rule requests submitted by users from the Self-Service Portal and executes them in DLP. By default, the Response Rule Execution Service is disabled. You must enable the service to allow the portal users to remediate incidents.

- Confirm the ownership of resources.

Once you submit a workflow from the Data Insight console, the custodians receive an email notification with a link to the Self-Service Portal. They can log in to the portal, choose the necessary remediation actions, and submit the same for execution by the DLP Enforce Server or the Data Insight Management Server, depending on the type of workflow.

Depending on the type of workflow, the custodian may perform the following actions:

Workflow	Actions
Entitlement Review	Review the user permissions on folders that the custodian owns and automatically trigger a permission remediation workflow to execute the changes.
	To trigger a permission remediation action, you must first configure the permission remediation settings.
	See "About configuring permission remediation" on page 224.
DLP Incident Remediation	Choose the configured remediation actions, and submit the same for execution by the DLP Enforce Server.

Workflow	Actions
Ownership Confirmation	Confirm the ownership of resources. Once the custodians confirm or deny the ownership, and the workflow is complete, the status summary is displayed in the Data Insight Management Console. A Data Insight administrator may review the status and take further actions based on it. See "Monitoring the progress of a workflow" on page 257.

See "About workflow templates" on page 242.

Prerequisites for configuring remediation workflows

Before you can submit a remediation workflow to the Self-Service Portal, verify that the following configuration tasks are complete:

- The Portal Server is installed and registered with the Management Server. You can verify the installation on the **Settings** > **Data Insight Servers** page.
  For more information about installing the Self-Service Portal, see the *Symantec Data Insight Installation Guide*.

- To create DLP Incident Remediation workflows, ensure that Data Loss Prevention (DLP) 12.5 is installed and the DLP settings are configured in Data Insight.
  See "Configuring Symantec Data Loss Prevention settings" on page 50.

- The directory service domains in your organization are configured in Data Insight, and the user and user group information is imported in Data Insight.

- The mail custom attribute is configured for the directory domain. Setting this attribute enables Data Insight to send the email alerts for submitted workflows to the custodian's email address.

- Custodians are assigned on paths that are configured in Data Insight. If some paths do not have any custodians assigned to them, you can assign custodians at the time of creating the workflow request.
  See "Create DLP Incident Remediation Workflow options" on page 252.

- The SMTP server settings are configured.

See the *Symantec Data Loss PreventionData Insight Implementation Guide* for information about using the Data Loss Prevention web services to remediate incidents.

Configuring Self-Service Portal settings

You can personalize the look and feel of the Self-Service portal to match the branding of your organization.

To configure the portal settings

1   In the Management Console, click **Settings** > **Self-Service Portal Settings**.

2   Edit all or any of the following settings:

Session timeout	Your login session on the Self-Service Portal times out after certain period of inactivity. The default time out period is 30 minutes. To change the default timeout period, enter the time in minutes.
Branding	Customize the header section of the portal by adding the following elements:

- To add the logo of your organization to the header, browse to the location where the image is saved, and select it.
  The image must be in the `.png`, `.gif` or `.jpg` formats only. It is recommended that the size of the image must be 32x32.
- Enter the text that you want to appear in the header section of the screen. For example, you can enter the name of your organization.
- Enter the hexadecimal values to define the font and the background colors to be used in the header area.

Login help text	Enter any information that the portal users may need to login to the portal. For example, the login credentials that are required for the portal. This information is optional.

| Support information | Enter information for the portal users to get assistance with the problems that they may encounter when using the portal. Support information can include an email address or the help desk number of the local support office. |

3 Click **Save**.

About workflow templates

Data Insight provides a way to create templates to help you quickly create remediation workflows for the resources that are monitored by Data Insight. Using workflow templates saves you time because you can use certain common values defined in the template to create multiple workflow instances of the same type. For resources that need remediation, the workflow template lets you define the attributes to be displayed on the Self-Service Portal and the actions that the Self-Service Portal users can take on these resources. For example, in case of a DLP Incident Remediation template, you can selectively choose the Data Loss Prevention (DLP) Smart Response rules that you want to present for action by the portal users.

Refer to the Data Loss Prevention documentation for details about DLP policies and Smart Response rules.

You can create multiple workflow instances from a template of the same type. You can also choose to edit a template to suit your requirement before you submit a workflow. For example, you can choose to change the frequency of email reminders, or customize the default email included in the template.

You can create a template for the following types of remediation workflows:

- Entitlement review
  See "Create/Edit Entitlement Review workflow template" on page 243.

- DLP Incident Remediation
  See "Create/Edit DLP Incident Remediation workflow template" on page 245.

- Ownership Confirmation
  See "Create/Edit Ownership Confirmation workflow template" on page 247.

See "Managing workflow templates" on page 243.

See "Creating a workflow using a template" on page 248.

Managing workflow templates

You can create multiple templates for each type of workflow. You can customize templates to define the different options that appear on the Self-Service Portal.

To create a workflow template

1   On the Management Console, click **Settings** > **Workflow Templates**.

2   On the list page, click **Add New Template**. Select the type of workflow template that you want to create. For example, DLP Incident Remediation.

3   Specify relevant values in each of the fields and click **Save**.

    You can use the template to create a remediation workflow of the same type.

    See "Creating a workflow using a template" on page 248.

You can edit, copy, or delete an existing template.

To manage existing templates

1   On the Management Console, click **Settings** > **Workflow Templates**.

2   Select a workflow template, and select the appropriate action:

    ■   To edit a template, click **Select Action** > **Edit**.
        Make necessary changes to the template, and click **Save**.

    ■   To copy a template, click **Select Action** > **Copy**.
        Enter the name of the new template. Data Insight creates a replica of the selected template with the new name.

    ■   To delete a template, click **Select Action** > **Delete**.
        Click **Yes** on the confirmation message.

Note: You cannot delete a template if it is being used for creating a workflow.

Create/Edit Entitlement Review workflow template

Use the dialog to create a template of type Entitlement Review.

Table 19-1   Entitlement Review template options

Option	Description
Template Type	Describes the type of workflow that can be created using the template.
Name	Enter a logical name for the template.

Table 19-1   Entitlement Review template options *(continued)*

Option	Description
Description	Enter a short description for the template. The description can state the kind of Entitlement Review workflow for which the template should be used.
Welcome Text	This text appears in a pop-up when the custodian first logs in to the Self-Service Portal. You can include the specific instructions for remediation in this field.
Portal Options	Select the check-boxes to display additional information on the Self-Service Portal. For example, select the **DLP Information** check box to display the number of sensitive files in a folder and the policies that they violate.
	To display custom attributes of a user in the portal, select the **Include custom attributes of user** check box. From the drop-down menu, select any of the custom attributes as per your requirements.
	To allow the reviewer to delegate the review task to another user, click **Allow delegation**.
Email Reminder	Select the frequency, day, time for sending email reminders to the custodians.
Customize Email	Do the following:
	1 Click to customize the email that is sent to custodians when a workflow is submitted for remediation.
	2 Insert the variable in the **To**, **From**, **CC**, and **Subject** fields.
	3 Add the **${workflow.link}** variable in the body of the email to include the link to the portal in the request.
	Note that the SMTP server settings must be configured to enable Data Insight to send an email.
	See "Configuring SMTP server settings " on page 35.

See "About workflow templates" on page 242.

See "Managing workflow templates" on page 243.

Create/Edit DLP Incident Remediation workflow template

Use the dialog to create a template of type Data Loss Prevention (DLP) Incident Remediation.

Table 19-2 DLP Incident Remediation template options

Option	Description
Template Type	Describes the type of workflow that can be created using the template.
Name	Enter a logical name for the template.
Description	Enter a short description for the template. The description can state the kind of DLP Incident Remediation workflow for which the template should be used.
Welcome Text	Select the check box to display a message to the portal users. Use the variables from the adjoining drop-down to create the message.
	This text appears in a pop-up when the custodian first logs in to the Self-Service Portal. You can include the specific instructions for remediation in this field.

Table 19-2   DLP Incident Remediation template options *(continued)*

Option	Description
Portal Options	Select the following: ■ **DLP Smart Response** - Data Insight fetches the configured Smart Response rules that are configured in DLP using the DLP Response Rule Listing Service API. The Response Rule Listing Service provides the available response rules for a a given incident. These rules define the actions that portal users are allowed to take on the paths that violate DLP polices, such as delete or quarantine. From the drop-down, select the configured Smart response rules. Click the **Refresh** icon to fetch the latest rules from DLP. ■ Select the check boxes for the file attributes that you want to display on the Self-Service Portal. The attributes displayed include information about the suggested owner and the DLP policy name. The file owner information may not reflect the responsible party. The responsible party or data owner can be a line manager in the business unit, the head of a department, or an information security officer. Data Insight provides information to tie the most active user of a file to a manager or responsible party for remediation steps. ■ Select **Allow delegation**, if you want to let the custodians delegate the workflow to other users being monitored by Data Insight.
Email Reminder	Select the frequency, day, time for sending email reminders to the custodians.
Customize Email	Do the following: 1 Click to customize the email that is sent to custodians when a workflow is submitted for remediation. 2 Insert the variable in the **To**, **From**, **CC**, and **Subject** fields. 3 Add the **${workflow.link}** variable in the body of the email to include the link to the portal in the request. Note that the SMTP server settings must be configured to enable Data Insight to send an email. See "Configuring SMTP server settings " on page 35.

See "About workflow templates" on page 242.

See "Managing workflow templates" on page 243.

Create/Edit Ownership Confirmation workflow template

Use the dialog to create a template of type Ownership Confirmation.

Table 19-3   Ownership Confirmation template options

Option	Description
Template Type	Describes the type of workflow that can be created using the template.
Name	Enter a logical name for the template.
Description	Enter a short description for the template. The description can state the kind of Ownership Confirmation workflow for which the template should be used.
Welcome Text	This text appears in a pop-up when the custodian first logs in to the Self-Service Portal. You can include specific instructions for the portal users in this field.
Portal Options	Select the check boxes for the file attributes that you want to display on the Self-Service Portal. ■ Select the **DLP Information** check box to display the number of sensitive files in a folder and the policies that they violate. ■ Select the **Show active user count** check box to display the number of active users for the data that is being remediated.
Email Reminder	Select the frequency, day, time for sending email reminders to the custodians.

Table 19-3   Ownership Confirmation template options *(continued)*

Option	Description
Customize Email	Do the following:
	1 Click to customize the email that is sent to custodians when a workflow is submitted for remediation.
	2 Insert the variable in the **To**, **From**, **CC**, and **Subject** fields.
	3 Add the **${workflow.link}** variable in the body of the email to include the link to the portal in the request.
	Note that the SMTP server settings must be configured to enable Data Insight to send an email.
	See "Configuring SMTP server settings " on page 35.

See "About workflow templates" on page 242.

See "Managing workflow templates" on page 243.

Creating a workflow using a template

You can create an instance of a workflow using an existing template or by creating a new template that precisely suits your needs.

To create and submit a workflow

1 In the Management Console, click **Settings** > **Workflow**.

2 On the list page, click **Create Workflow**, and click the type of workflow you want to create.

   On the workflow panel, enter the relevant information.

3 Click **Submit** to submit the workflow for further action by the custodians, or click **Save & Close** to save the workflow details.

See "Create Entitlement Review Workflow options" on page 249.

See "Create DLP Incident Remediation Workflow options" on page 252.

See "Create Ownership Confirmation Workflow options" on page 254.

Create Entitlement Review Workflow options

Use the dialog to create an instance of an Entitlement Review workflow. You can view the summary of the options you select in the right-hand panel of the page.

Table 19-4   Create Entitlement Review workflow

Option	Description
Workflow Information	Enter information in the following fields: ■ Name - Enter a logical name for the workflow. ■ Description - Enter a short description for the workflow. Workflow Type - Describes the type of workflow. Template - Select the template you want to use for creating the workflow. See "About workflow templates" on page 242. ■ Portal Node for Execution - From the drop-down, select the Self-Service Portal node to which you want to submit the workflow. ■ Click **Test portal connection** to test the availability of SMTP connection to the Self-Service Portal. Enter the email IDs of the recipients of the workflow request, and click **Test**. You will see a response from the SMTP server if the connection to the Portal node succeeds. ■ Action - Select **Apply configured permission remediation action automatically** to let Data Insight automatically take the configured actions by a remediation workflow. To avail this feature, you must first configure Data Insight for permission remediation. See "About configuring permission remediation" on page 224. ■ Schedule - Select the start and the end date for completing the workflow.

Table 19-4   Create Entitlement Review workflow *(continued)*

Option	Description
Data Selection	Do the following:
	1 Select the **Physical** radio button to view the configured file servers or SharePoint web applications. Or, select the **DFS** radio button to view the configured DFS paths in a domain.
	2 From the **Resource Selection** drop-down, select one of the following options: ■ Physical or DFS paths - Select the physical or DFS paths for which you want to review the user permissions. ■ Opens Shares - Select the open shares that need to be remediated. ■ Containers - Select configured containers. Data Insight presents the paths in the containers to remediate user permissions. ■ Enter paths manually - Enter the full path that you want to remediate. ■ Upload CSV - Browse to the location of the .csv file that contains the paths that you want to remediate. Only valid paths in the .csv file are displayed in the **Selected Resources** pane. ■ Select paths having custodians - Data Insight retrieves only the list of paths that have custodian assignments. Select the paths from the list.
	The selected data set is listed in the **Selected Resources** pane.
	Note: Data Insight does not support NFS and SharePoint paths for the Entitlement Review workflows. If you select a container which contains NFS paths, then those paths will not be sent to the custodian for review.

Table 19-4   Create Entitlement Review workflow *(continued)*

Option	Description

Resource -Custodian Selection	This panel displays the following:
	
	■ The paths that you select under the **Data Selection** tab.
	■ The paths for which custodians are already assigned and those paths for which custodians are not assigned.
	■ The email address of the custodian.
	Data Insight displays the email address only if you have added the email custom attribute, and have also marked the attribute as email alias when you add the directory service. See "Adding a directory service domain to Data Insight" on page 65.
	
	You can assign custodians on paths or remove already assigned custodians. For example:
	
	■ Click **Import Custodian** to assign custodian to a selected path. Select any of the following options:
	  ■ Upload a .csv file with custodian information.
	  ■ Select a user who is configured in Data Insight as the custodian.
	  ■ Select a Data Insight suggested data owner as the custodian.
	  ■ Select a custom attribute of a Data Insight suggested data owner and assign it as a custodian. For example, you can select the manager of a user who is a suggested data owner as the custodian.
	■ Click **Assign Custodian** to manually assign the custodian for a selected path. Use the domain filter to filter the users based on their directory domains.
	■ Click **Remove Custodian** to remove a custodian from a selected path.
	■ Click **Delete Paths** to remove the selected paths.

Exclusion List	Select the groups or users that you want to exclude from the scope of the review. Click the group or user to select it. The selected data set is listed in the **Selected Groups/Users** panel. Once you have excluded a user or a group, the activities of the user or the group on the paths will be ignored and thus will not be considered for the review.

Create DLP Incident Remediation Workflow options

Use the dialog to create an instance of a Data Loss Prevention (DLP) Incident Remediation workflow. You can view the summary of the options you select in the right-hand panel of the page.

Table 19-5   Create DLP Incident Remediation workflow

Option	Description
Workflow Information	Enter information in the following fields: ■ **Name** - Enter a logical name for the workflow. ■ **Description** - Enter a short description for the workflow. ■ **Workflow Type** - Describes the type of workflow. ■ **Template** - Select the template you want to use for creating the workflow. See "About workflow templates" on page 242. ■ **Portal Node for Execution** - From the drop-down, select the Self-Service Portal node to which you want to submit the workflow. Click **Test portal connection** to test the availability of network connection between the Data Insight Management Server and the Self-Service Portal. ■ Click **Test portal connection** to test the connection between the SMTP server and the DLP Enforce Server to the Self-Service Portal. Enter the email IDs of the recipients of the workflow request, and click **Test**. You will see a response from the SMTP server if the If the connection to the Portal node succeeds. ■ Select the start and the end date for completing the workflow.

Table 19-5   Create DLP Incident Remediation workflow *(continued)*

Option	Description
Data Selection	Do the following:

1   Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint web applications. Or, select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain.

2   From the Resource Selection drop-down, select one of the following options:

- Physical or DFS paths - Select the physical or DFS paths that violate DLP policies.
- Opens Shares - Select the open shares that need to be remediated.
- Containers - Select configured containers. Data Insight presents the paths in the containers that violate DLP policies.
- Policies - Data Insight displays the configured DLP policies. Select a policy to remediate the paths that violate the policy.
- Enter paths manually - Enter the full path that you want to remediate.
- Upload CSV - Browse to the location of the .csv file that contains the paths that you want to remediate. Only valid paths in the .csv file are displayed in the **Selected Resources** pane
- Select paths having custodians - Data Insight retrieves only the list of paths that have custodian assignments. Select paths from the list.

  You must you run the Custodian Summary report to fetch recent custodian assignments.

The selected data set is listed in the **Selected Resources** pane.

Table 19-5 Create DLP Incident Remediation workflow *(continued)*

Option	Description
Resource -Custodian Selection	This panel displays the following:
	■ The paths that you select under the **Data Selection** sub-tab.
	■ The paths for which custodians are already assigned and those paths for which custodians are not assigned.
	■ The email address of the custodian. Data Insight displays the email address only if you have added the email custom attribute and have also marked the attribute as email alias when you add the directory service. See "Adding a directory service domain to Data Insight" on page 65.
	For the paths that do not have custodians, you can assign custodians using the following methods:
	1 Click **Import Custodian**, and select one of the following options:
	■ Upload a .csv file with information about paths and corresponding custodians
	■ Select a user who is configured in Data Insight as the custodian.
	■ Select a Data Insight suggested data owner as the custodian.
	■ Select a custom attribute of a Data Insight suggested data owner and assign it as a custodian. For example, you can select the manager of a user who is a suggested data owner as the custodian.
	2 Click **Assign Custodian**, and select the custodian from the users list.
	You can remove custodians from selected paths or delete paths from the workflow. Do the following:
	1 Click **Remove Custodian** to remove a custodian from a selected path.
	2 Click **Delete Paths** to remove the selected paths from the workflow.

Create Ownership Confirmation Workflow options

Use the dialog to create an instance of an Ownership Confirmation workflow. You can view the summary of the options you select in the right-hand panel of the page.

Table 19-6   Create Ownership Confirmation workflow

Option	Description
Workflow Information	Enter information in the following fields: ■ **Name** - Enter a logical name for the workflow. ■ **Description** - Enter a short description for the workflow. ■ **Workflow Type** - Describes the type of workflow. ■ **Template** - Select the template you want to use for creating the workflow. ■ **Portal Node for Execution** - From the drop-down, select the Self-Service Portal node to which you want to submit the workflow. ■ Click **Test portal connection** to test the availability of SMTP connection to the Self-Service Portal. Enter the email IDs of the recipients of the workflow request, and click **Test**. You will see a response from the SMTP server if the the connection to the Portal node succeeds. ■ **Schedule** - Select the start and the end date for completing the workflow.
Data Selection	Do the following: 1 Click **Select All Resources** button to select all the paths which have custodians assigned. 2 Filter the list of displayed paths , based on custodian name or custodian attribute, and manually select the resources. 3 Click **Re-Generate** to regenerate the custodian map and refresh the listed paths in the panel. The selected data set is listed in the **Selected Resources** pane.
Resource -Custodian Selection	This panel displays the data set selected in the **Data Selection** tab. You can review the selected paths on the basis of criteria such as custodians and custodian email. You can remove a selected path from the list. Click **Delete Paths** to remove any paths from the selected resources.

Managing workflows

On the workflow details page, you can view detailed information about submitted workflows, the start and the end date of the workflow, their completion status, and the number of custodians who have been assigned the workflow.

You can also create new workflows from the workflow list page.

To view submitted workflows

1   In the console, click **Settings** > **Workflows** .

2   On the **Workflows** list page, review the following information about the
    submitted workflows:

    ■   The name of the workflow.

    ■   The status of the workflow.
        If there are multiple workflows listed on the page, use the **Filter** field to
        search for the workflow you are interested in. Or use the dynamic search
        option to search for workflows based on the type or their status.
        Hover the mouse over the progress bar to know the number of completed
        and pending requests in the workflow.

    ■   The date when the workflow was submitted.

    ■   The number of days within which the workflow request must be completed.

    ■   The number of custodians that have been assigned the workflow.

To extend the deadline of a submitted workflow

1   On the **Workflows** list page, click the **Select Action** drop-down corresponding
    to the workflow for which you want to extend the deadline.

2   Select **Extend Deadline**.

3   On the pop-up, select the new end date for the workflow.

4   Click **Extend**.

Note: Once you submit a workflow, you can modify only the deadline.

You can cancel a submitted or in-progress workflow. For example, you can cancel
a workflow, if the custodian who is required to take action has left the organization.
Custodians will stop receiving email reminders to complete the workflow when you
cancel it.

You can delete a workflow if the deadline for completing the tasks in the workflow
has expired, or if the workflow is complete or has been canceled.

To cancel or delete a workflow

1 On the **Workflows** listing page, click the **Select Action** drop-down corresponding to the workflow you want to delete or cancel.

2 Click **Delete** or **Cancel**, as appropriate.

See "About remediation workflows" on page 238.

Monitoring the progress of a workflow

On the Workflows listing page, you can view the progress of workflows that are submitted to the Self-Service Portal. You can also view the details of the actions that are taken on all paths that are part of a workflow.

To view the status of a workflow

1 On the console, click **Settings** > **Workflows**.

On the **Workflows** list page, you can view the status for each workflow. The following table describes the possible status for any workflow:

Status	Description
Draft	When the workflow is saved as a draft but is not submitted it to the portal server.
Submitted	When the workflow is submitted from the Management Server but is not picked up by the portal server for processing.
In-progress	When workflow is being processed by the portal server for processing .
Completed	A workflow is marked as complete if: ■ The end date of the workflow lapses and after a day of grace period from the end date. ■ An action is taken on all the paths by all the custodians and the portal server has processed the workflow
Canceled	If you have canceled the workflow.

Status	Description
Grace Period	After the due date of the workflow, an extra day is given as grace period. In this case, the state of the workflow is set to **Grace Period**. If actions are still not taken by the end of the grace period, the status changes to **Completed**, and the state of the paths will be shown as **Expired**.
Failed	If Data Insight fails to create a workflow database based on the input that is provided for the workflow.

2   On the workflow listing page, click **Select Action** > **View**, or click the workflow link to view details of a submitted, completed, or canceled workflow.

3   On the workflow summary page, you can view the list of paths that are submitted for custodians' actions on the Self-Service Portal. The page also displays the summary of the total paths in the workflow, the percentage of paths on which an action is submitted on the portal, and the time within which the workflow must be completed.

You can also view the following details:

- The list of paths that are part of the workflow.

- In case of a DLP Incident Remediation workflow, the Data Loss Prevention (DLP) policies that the paths violate, the severity of the incidents0, and the incident IDs that need to be remediated. The incident ID is associated with the available response rules for a given incident.

- The custodian(s) for whose action the workflow is submitted.

- The status for each path can be one of the following:

  - Pending - Indicates that the custodian has not taken any action on the assigned paths.

  - Success - Indicates that the custodian has submitted an action and the action has been registered with the Data Insight Management Server. In case of a DLP Incident Remediation workflow, it means that Data Insight has sent the response rule request for execution to the DLP Response Rule Execution Service.

  - Failed - Indicates that the action submitted by the portal user on the Self-Service Portal is not registered with the Data Insight Management Server for any reason.

- Expired - Indicates that the due date for completing the workflow has expired, and the portal users will not be able to take any action on the paths in that particular workflow,

- Depending on the type of workflow, you can also view the following information about the actions taken on the files assigned for remediation:

Workflow	Details
Entitlement Review	Click the name of a path to see the details of the user permissions on that path. For each of the users for the path you can view the following information: ■ The user name ■ The login ID of the user ■ The type of permission the user has on the path. For example, read, write etc. ■ Activity status of the user. ■ Whether the user is allowed access on the path or not.
DLP Incident Remediation	The actions are based on configured DLP Smart Response rules, for example, Quarantine, Mark for Deletion, or Archive. For information about Smart Response rules, see the *Symantec Data Loss Prevention Administration Guide*. A possible action can also be *Delegate* if the custodian delegates the incident remediation for certain paths to another user.
Ownership Confirmation	The possible actions for any path can be *Confirm* or *Decline* ownership.

Configuring policies

This chapter includes the following topics:

- About Data Insight policies
- Managing policies
- Managing alerts

About Data Insight policies

A policy is a set of conditions that you configure to monitor access events on files and folders stored on various repositories. Symantec Data Insight policies help you detect the sources of threat, access patterns on sensitive data, and anomalous user behavior. Data Insight receives information about sensitive files from Symantec Data Loss Protection (DLP).

Policies must include at least one condition that is configured to detect abnormal access patterns or user behavior. Data Insight generates an alert whenever it detects any violation of a condition in a configured policy.

Policies can be configured with three severities, namely, high, medium, and low. You can assign the severity level to a policy based on your organizational needs. For example, the Information Security team can define policies to monitor accesses on the share `\Finance`. For this purpose, they can configure a policy with a medium severity to monitor accesses on folders containing Finance policies and guidelines files. Whereas, they can configure a policy with a high severity to monitor accesses on files containing payroll information. When an alert is generated for a policy violation, the severity of the policy is associated with the alert.

Data Insight comes packaged with the following out-of-the-box policies that you can configure according to your needs:

- Data Activity Trigger policy

Use this policy to define the maximum cumulative count of the meta operations on the selected paths. For example, if you have defined the maximum accesses per day as 500 on the share `\\netapp1\finshare`, and the total access count by the active set of users exceeds 500, then Data Insight generates an alert.

- User Activity Deviation policy
  Use this policy to define the threshold of deviation from the baseline activity. The baseline activity on a file or folder is the average number of accesses that are considered normal based on past access counts. If the activity, by the selected users, on the selected data exceeds the specified threshold of the baseline (for example, three standard deviations above the baseline activity), or the maximum accesses allowed per day, Data Insight generates an alert. You can configure how many standard deviations a user is allowed to deviate from the defined baseline.

- Data Activity User Whitelist-based policy
  Use this policy to define a whitelist of users based on the Active Directory custom attributes, who can access selected shares or paths. Also, you can create such a policy with multiple conditions with multiple values for the same custom attributes .
  If users, other than those defined in the whitelist, access selected data, Data Insight generates an alert.

Managing policies

You can view, edit and delete configured policies, and add new policies to Data Insight from the **Policies** tab.

To manage policies

1 In the Console, click the **Policies** tab.

   The left pane displays the default policy groups.

2 Click a policy group.

   The policy listing page displays the configured policies for that policy group.

3 To edit an existing policy, from the Actions drop-down, click **Edit**.

4 To delete a policy, select the corresponding check box and click **Delete**.

To add a new policy

1 In the Console, click the **Policies** tab.

   The left pane displays the default policy groups.

2 Click the policy group that you want to base your policy on.

3   On the policy listing page, click **Add new policy**. Or in the tree-view panel, right-click the policy type, and select **Add**.

4   On the Add new policy page, select the options to create the policy.

Click the collapsed panels on the page to expand them.

5   Click **Save**.

By default, policies are evaluated at 12:00 A.M. every night. You can schedule policies to be evaluated more frequently for proof-of-concept (POC) setups. Note that a schedule that is too aggressive can put excessive load on the Indexer.

You can set a custom schedule to evaluate policies from the **Settings** tab. The schedule must be specified in the cron format.

To set a custom schedule for policies

1   Click **Settings** > **Data Insight Servers**.

2   Click the entry for the Management Server.

3   On the page for the Management Server node, click **Advanced Settings**.

4   Click **Edit**.

5   Scroll to bottom of the page and expand the **Set custom properties** section. Specify property name to be job.PolicyJob.cron and property value to be the new schedule. Schedule needs to be specified in cron format

6   In the Property name field, enter job.PolicyJob.cron.

7   In the Property value fields, enter the values as follows:

To evaluate values every N minutes, specify value as 0 0/N * * * ? *.	For example, to evaluate policies every 10 minutes, specify value as 0 0/10 * * * ? *.
To evaluate policies every N hours, specify value as 0 0 0/N * * ? *.	For example, to evaluate policies every two hours, specify value as 0 0 0/2 * * ? *.

See "Create Data Activity Trigger policy options" on page 262.

See "Create User Activity Deviation policy options" on page 265.

See "Create Data Activity User Whitelist-based policy options" on page 267.

Create Data Activity Trigger policy options

Use this dialog to create a new Data Activity Trigger policy. Click the collapsed panels on the page to expand them and enter the relevant information in each panel of the dialog. Options that are selected in the respective panels are displayed in the **Summary** panel on the right of the page.

Table 20-1   Create Data Activity Trigger policy options

Option	Description
Policy Information	Enter information in the following fields: ■ **Name** - The name of the policy. ■ **Description** - A short description of the policy. ■ **Policy Type** - Data Activity Trigger is selected by default. ■ **Severity** - The severity of the policy. From the drop-down, select High, Medium, or Low. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs. Select the **Enable Policy** check box to enforce the policy. The policy is not evaluated if the check box is not selected.
Configure Policy	Select the following conditions to configure the policy: ■ **Select Activity** - Select the type of accesses to be monitored on the selected data set. Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events. Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events. ■ **Additional Condition** - From the **Minimum accesses per day for alerts** drop-down, select the minimum number of accesses on the selected data set on that day that are required to trigger an alert.

Table 20-1   Create Data Activity Trigger policy options *(continued)*

Option	Description
Data Selection	Do the following to select the resources:
	1 Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint Web applications.
	Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain.
	2 Click a path to select it.
	Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**.
	3 To limit the scope of the files to be monitored, select **Select all files in folder** or **Select only sensitive files**.
	If you select the **Select all files in folder** option, accesses on all files in the folder are evaluated for determining any violation of the policy.
	If you select the **Select only sensitive files** option, accesses on only the sensitive files in the folder are evaluated for determining any violation of the policy.
	Note: The list of sensitive files is obtained from Symantec Data Loss Prevention (if configured).
	See "Configuring Symantec Data Loss Prevention settings" on page 50.
	4 The selected data set is listed in the **Selected resources** pane.

Table 20-1   Create Data Activity Trigger policy options *(continued)*

Option	Description
Notification	Enter one or more specific email addresses for people to whom you want to send alerts that are generated for the policy.
	Select the **Notify custodians** check box to send email alerts to the custodians of the selected paths along with other email addresses on the notification list.

Create User Activity Deviation policy options

Use this dialog to create a new User Activity Deviation policy. Click the collapsed panels on the page to expand them and enter the relevant information in each panel of the dialog. Options that are selected in the respective panels are displayed in the **Summary** panel on the right of the page.

Table 20-2   Create User Activity Deviation policy options

Option	Description
Policy Information	Enter the following information:
	■ **Name** - The name of the policy.
	■ **Description** - A short description of the policy.
	■ **Type** - User Activity Deviation is selected by default.
	■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs.
	Select the **Enable Policy** check box to enforce the policy.
	The policy is not evaluated if the check box is not selected.

Table 20-2   Create User Activity Deviation policy options *(continued)*

Option	Description
Configure Policy	Do the following: 1 From the drop-down, select the time range for the baseline activity. Baseline activity is then computed as the average access in that time range. From the **Threshold Configuration** drop-down, select the threshold of normal activity. The threshold is the acceptable number of standard deviations that a user is allowed to deviate. Accesses above the defined threshold trigger an alert. 2 **Additional Condition** - From the drop-down, select the minimum accesses per day per user. Alerts are raised only if the total accesses exceed the minimum value specified. This prevents Data Insight from raising too many alerts when baselines are very low.
Data Selection	Do the following to select the resources: 1 Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint Web applications. Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. 2 Click a path to select it. Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**. The selected data set is listed in the **Selected resources** pane.

Table 20-2   Create User Activity Deviation policy options *(continued)*

Option	Description
User Selection	Do the following: 1 Select the **Users** or **Group** radio button to view configured users or groups respectively. The list of users and groups is sorted alphabetically. Click the star icon to display all the configured users or groups. You can use the Domain filter search bar to filter users or groups according to domains. You can also filter the users according to their Active Directory custom attributes. 2 Click a user or group to select it.
Notification	Enter one or more specific email addresses for people to whom you want to send the alerts that are generated for the policy.

Create Data Activity User Whitelist-based policy options

Use this dialog to create a new Data Activity User Whitelist-based policy. Click the collapsed panels on the page to expand them and enter the relevant information in each panel of the dialog. Options selected in the respective panels are displayed in the **Summary** panel on the right of the page.

Table 20-3   Create Data Activity User Whitelist-based policy options

Option	Description
Policy Information	Enter the following information:
	■ **Name** - The name of the policy.
	■ **Description** - A short description of the policy.
	■ **Type** - Data Activity User Whitelist-based is selected by default.
	■ **Severity** - The severity of the policy. The severity level associated with the policy helps you decide the possible course of action when an event that matches the policy occurs.
	Select the **Enable Policy** check box to enforce the policy.
	The policy is not evaluated if the check box is not selected.
Configure Policy	**Select Activity** - Select the type of accesses to be monitored on the selected data set.
	Select the **Meta Access** radio button to monitor only the high-level access events that Data Insight maps from the detailed file system and SharePoint access events.
	Select the **Detailed Access** radio button to monitor specific file system and SharePoint access events.

Table 20-3   Create Data Activity User Whitelist-based policy options *(continued)*

Option	Description
Data Selection	Do the following to select the resources:
	1 Select the **Physical Hierarchy** radio button to view the configured file servers or SharePoint Web applications.
	Or select the **DFS Hierarchy** radio button to view the configured DFS paths in a domain. 2
	2 Click a path to select it.
	Or you can use a `.csv` file with information about the paths that you want to apply the policy to. Click **Browse** to navigate to the location of the `.csv` file, and click **Upload**.
	3 To limit the scope of the files to be monitored, select **Select all files in folder** or **Select only sensitive files**
	If you select the **Select all files in folder** option, accesses on all files in the folder are evaluated for determining any violation of the policy.
	If you select the **Select only sensitive files** option, accesses on only the sensitive files in the folder are evaluated for determining any violation of the policy.
	The selected data set is listed in the Selected resources pane.
	Note: Data Insight obtains information about sensitive files from Symantec Data Loss Prevention (DLP).
	See "Configuring Symantec Data Loss Prevention settings" on page 50.

Table 20-3   Create Data Activity User Whitelist-based policy options *(continued)*

Option	Description
Whitelist Conditions	Do the following: 1 Click **Add Condition**. 2 Select the criteria to build the condition. Use the expression builder to select users based on their Active Directory custom attributes. You can add multiple conditions to a Data Activity User Whitelist-based policy.
Notifications	Enter one or more specific email addresses for people to whom you want to send alerts that are generated for the policy. Select the **Notify custodians** check box to send email alerts to the custodians of the selected paths along with other email addresses on the notification list.

See "Configuring Symantec Data Loss Prevention settings" on page 50.

Managing alerts

An alert is a signal generated by a policy when the condition specified in the policy is violated.

You can view alerts on the **Alerts** tab on the Management Console.

To manage alerts

1   In the Console, click the **Policies** tab.

The **Alerts** tab display by default. On the tab, you can view all the alerts that were generated by Data Insight.

2   In the Alerts Summary, click the drop-down arrow on any column header and select Columns. Then, select the parameters you want to show or hide. You can sort by:

■   The name of the policy.

■   The severity of the alert.

■   The type of policy associated with the alert - Data Activity Trigger, User Activity Deviation, or Data Activity User Whitelist-based.

- The name of the user account that violated the policy.

- The date on which the alert was generated.

- The resolution, if any, taken in response to the alert.

3   To send alerts in email, select the alerts and click **Send Email**.

4   Enter the email addresses and click **Send**.

5   To enter the resolution for an alert, select the alert, click in the Resolution column for the alert and type in the resolution.

To update the resolution for multiple alerts, select the alerts and click **Update Resolution** at the top of the summary table.

To delete alerts

◆   To delete an alert, select an alert and click **Delete**.

To delete alerts by severity, click Delete and select the severity. This deletes all alerts that match the selected severity.

To delete alerts older than a certain date, click Delete and select the date at the top of the table.

Note: You can configure automatic deletion of alerts older than the specified interval on the Data Retention screen. However, you cannot restore the alerts once they are deleted. Alerts are also automatically published to the Windows event log.

See "Configuring data retention settings" on page 48.

Events and Notifications

This chapter includes the following topics:

- Configuring email notifications
- Enabling Windows event logging
- About high availability notifications
- Viewing events
- Viewing scan errors

Configuring email notifications

Data Insight provides email notifications for important events happening in the product. For example, CIFS scan failure or a directory scan failure. Notifications are sent out every 15 minutes, if new events are available. Email notifications are not enabled by default.

Note: Before you enable email notifications, you must enable configure the SMTP settings.

See "Configuring SMTP server settings " on page 35.

To configure email notifications

1 In the Management Console, click **Settings** > **Global Settings** > **Event Notifications**

2 On the Event Notifications page, select **Enable event notifications** checkbox.

3 In the Email recipients field, enter a comma separated list of email addresses to be notified.

4   Select the severity of events for which the email notifications must be sent.

5   Click **Save**.

Enabling Windows event logging

Symantec Data Insight can publish events to the Windows Event log. Events are published on the same machine where they originate. Event logging is enabled by default.

To configure Windows event logging

1   In the Management Console, click **Settings** > **Global Settings** > **Event Notifications**.

2   Select the **Enable Windows logging** checkbox.

3   Select the severity of events for which you want to enable Windows logging.

4   Click **Save**.

About high availability notifications

Data Insight raises events for various conditions that might result in a loss of availability of a Data Insight system or component. Events are raised for the following conditions:

■   Changes in the state of various essential services

■   Saturation of the data volume

■   Worker node misses heartbeat with the Management Server

■   Accumulation of excessive files on the worker node

■   Loss of connection between the filers and the Collector

■   Excessive usage of CPU, memory, or disk space for extended period

See "Configuring advanced settings" on page 203.

Viewing events

You can monitor Symantec Data Insight recent system events on the **Events** page. The report displays entries for all system events. These events include the following information about an event:

■   Time

■   Severity

- Event summary

- Symantec Data Insight server where the event originated

- The user if any performing the action

- The object for which the event originated

To view system events

1   A list of recent system events appears.

2   You can choose to filter the events further using one or all of the following
    criteria:

    - By time

    - By any text appearing in the event summary

    - By severity

    - By the product server on which the event originates
      Enter the filter criteria in the relevant fields and click **Go**.

3   Click the Export icon at the bottom of the page to save the data to a `.csv` file.

Viewing scan errors

You can view a list of all the paths on which a scan has failed. In the **Workspace**
tab tree-view panel, the folder icon displays a red cross mark for the paths on which
a scan has failed. The scan errors displayed are from the latest scan completed on
the share.

To view scan errors

1   Do one of the following to view the details of the scan errors on paths:

    - In the Management Console, click **Settings** > **Scanning**.
      The **Overview** sub-tab of the Scanning dashboard displays the paths on
      which scans have failed in the last 24 hours.

    - On the **Scan Status** sub-tab of the **Scanning** dashboard, click the **Select
      Action** drop-down corresponding to a path and select **Scan Errors**.

    - Navigate to a share or a site collection on the **Monitored Shares/Monitored
      Site Collections** sub-tab on the filer or SharePoint Web applications details
      page., click the **Select Action** drop-down corresponding to a path, and
      select **Scan Errors** to view the failed scans on that path.

2   On the Scan Errors page, review the time of the error, the error code, and the
    possible cause of the error.

Backing up and restoring data

This chapter includes the following topics:

- Selecting the backup and restore order
- Backing up and restoring the Data Insight Management Server
- Backing up and restoring the Indexer node

Selecting the backup and restore order

To maintain consistency in the configuration data during backup, the backup the Data Insight components in the following order:

- Nodes with Indexer role
- Management Server

Restore the Data Insight components in the following order:

- Management Server
- Nodes with Indexer role

Backing up and restoring the Data Insight Management Server

It is mandatory to backup the Management Server.

To backup Management Server configuration files

1   Log in Data Insight Management Server.

2   Backup the entire `$data` folder using backup tools such as Symantec NetBackup. The backup software should be capable of taking Volume Shadow Copy/Snapshot based backups. If your backup software does not have such a capability, you must stop all Data Insight services before backup, to avoid incomplete backup due to locked files.

To restore the Management Server files

1   Install the operating system. Use the same version, host name (recommended for ease of configuration), and architecture as was installed before the backup.

2   Install the same version of Data Insight that was installed before the backup.

3   Select the option to install the Management Server role while installing Data Insight .

4   Specify the original location of the `$data` directory as the previous install. By default, the `$data` directory is located at `C:\DataInsight\data`.

5   Complete the installation. Do not start the services at this time; clear the **Start services now** option when the installer prompts for it.

6   Delete the `$data` folder that is created as a part of the new installation and copy the backed up data to this location.

7   Start the Data Insight services, which include DataInsightComm, DataInsightWatchdog, DataInsightHttpd, DataInsightweb, and DataInsightConfig.

8   Check the status of the services and ensure that they come to running state.

Successful start of all services indicates that the Management Server is successfully restored.

To restore the Management Server with a different host name or IP address

1   Repeat steps 1 through 6 as described in the section, *Restoring the Management Server files*.

2   Edit `$data/conf/<nodename>.conf` and enter the new server name.

Open the file, `$data/conf/config.db.<N>` (N being the latest version of config.db) in an SQLITE editor.

Update the node_name and node_ip columns in node table with the host name and IP address of the new server.

3 Run the following SQL updates:

```
update node set node_name='<new node name>'
where node_name='<prev node name>';

update node set node_ip='<new node name>'
where node_name='<prev node name>';
```

4 Open a Windows command prompt and run the following command to increment the version of the config.db file that was changed in Step 3:

```
<INSTALL DIR>\DataInsight\bin\configdb −O -J dummy −j dummy
```

5 Start all Data Insight services.

6 On each worker node, except the Windows File Server agents, stop DataInsightComm and DataInsightConfig services.

7 Perform steps 2and 3 on the worker node's config.db.N

8 Start the DataInsightComm and DataInsightConfig services. Ensure that the worker nodes show online on the Data Insight Management Console.

Backing up and restoring the Indexer node

You must mandatorily backup the nodes that serve as the Indexer roles in a Data Insight deployment.

To back up the Data Insight server with Indexer role

1 Log in to the server with the Indexer role.

2 Backup the entire $data folder using backup tools such as Symantec NetBackup. The backup software should be capable of taking Volume Shadow Copy/Snapshot based backups. If your backup software does not have such a capability, you must stop all Data Insight services before backup, to avoid incomplete backup due to locked files.

To restore the Indexer node

1 Install the operating system. Use the same version, host name (recommended for ease of configuration), and architecture as was installed before the backup.

2 Install the same version of Data Insight that was installed before the backup.

3 Select the option to install the Indexer and Collector role while installing Data Insight. If installing Data Insight on a Linux server, select the option to install Indexer.

4   Specify the original location of the `$data` directory as the previous install. By default, the `$data` directory is located at `C:\DataInsight\data`.

5   Clear the **Launch worker node registration wizard after exit** checkbox. You do not need to register the worker node at this time as the registration information is already present in the data that you have backed up.

6   Complete the installation. Do not start the services at this time; clear the **Start services now** option when the installer prompts for it.

7   Delete the `$data` folder that is created as a part of the new installation and copy the backed up data to this location.

8   Start the Data Insight services, which include DataInsightComm, DataInsightWatchdog, and DataInsightConfig.

9   Check the status of the services and ensure that they come to running state.

     Successful start of all services indicates that the Indexer node is successfully restored.

To restore the Indexer node with a different host name or IP address

1   Repeat steps 1 through 6 as described in the section, *Restoring the Indexer node.*.

2   Edit `$data/conf/<nodename>.conf` and enter the new server name.

3   Open the file, `$data/conf/config.db.<N>` (N being the latest version of config.db) in an SQLITE editor.

     Update the node_name and node_ip columns in node table with the host name and IP address of the new server.

4   Run the following SQL updates:

```
update node set node_name='<new node name>'
 where node_name='<prev node name>';

update node set node_ip='<new node name>';
where node_name='<prev node name >';
```

5   Log in to the Management Server and stop the DataInsightComm, DataInsightWeb, and DataInsightConfig services.

6   Perform 3 on the Management Server.

7   Open a Windows command prompt and run the following command to increment the version of the config.db file that was changed in 2

```
<INSTALL DIR>\DataInsight\bin\configdb -O -J dummy -j dummy
```

8   Start all Data Insight services on the Management Server.

9   On each worker node, except the Windows File Server agents, stop
    DataInsightComm and DataInsightConfig services.

10 If this node is a Collector for one or more Windows File Server agents, log in
    to each Windows File Server, stop the DataInsightComm and DataInsightConfig
    services.

    Perform step 3 on the worker node's config.db.N

11 Start the DataInsightComm and DataInsightConfig services on the Indexer and
    all other worker nodes where configdb.N was changed. Ensure that the worker
    nodes show online on the Data Insight Management Console.

Troubleshooting

This appendix includes the following topics:

- About general troubleshooting procedures
- About the Health Audit report
- Location of Data Insight logs
- Downloading Data Insight logs
- Migrating the data directory to a new location
- Enterprise Vault exceptions and their meanings

About general troubleshooting procedures

This section provides an overview of the general troubleshooting procedures that you can use to help discover and troubleshoot common problems.

You can use the **Events** page on the Data Insight Management Console to get a quick overview of the node on which the error has occurred.

To troubleshoot a problem, it helps to consider the following:

- Check for prior occurrence.
  Check existing troubleshooting information to see if the problem has occurred before and if there is a workaround available to troubleshoot the same. A good source for this type of information is the *Symantec Data Insight Release Notes*. The Release Notes contain a list of known issues for Data Insight and a list of possible workaround.

- Consider recent alterations.
  If a system is having problems immediately after some kind of maintenance, software upgrade, or other change, the problem might be linked to those changes.

Data Insight runs a daily health audit report that helps you identify potential problems in your environment.

About the Health Audit report

The Health Audit report gives you an overall status of the heath of Data Insight deployment. You can use the report to identify potential problems in the system and proactively take measures to fine tune the system settings. The report aids Symantec Support to troubleshoot issues on your setup.

The HealthAuditReportJob periodically creates the report at 5:00 A.M. daily and stores it on the `log/health_audit` folder.

Location of Data Insight logs

Symantec Data Insight log files are located in the Data Insight installation directory, `<INSTALLDIR>\log`. Typically the installation directory is located at `C:\Program Files\Symantec\Data Insight\log`. On Linux, the logs are located at `/INSTALL/DataInsight/log`

Table A-1 describes the logs that are relevant for troubleshooting.

Table A-1   Data Insight logs

`webserver0.0.log`	This file contains the log messages from the Web service process.
`commd0.0.log`	This file contains the log messages from the scheduler communication service.
`adcli.log`	This file contains the log messages from the Active Directory scanner process, `adcli.exe`.
`celerrad.log`	This file contains the log messages for DataInsightCelerra service.
`cli0.0.log`	This file contains the log messages for various command line utilities.
`collector.log.N`	This file contains the log messages for the audit pre-processor (collector.exe).
`dashboard.log`	This file contains the log messages for the Dashboard data generation report.
`dscli0.0.log`	This file contains the log messages for LDAP, NIS, NIS+ Directory scanner.
`scanner/extN_msuN.log`	This file contains the log messages for Full file system scans.

Table A-1   Data Insight logs *(continued)*

`scanner/extN_msuN.ilog`	This file contains the log messages for Incremental file system scans.
`fpolicyd.log`	This file contains the log messages for DataInsightFpolicy service.
`idxcheck.log`	This file contains log of index integrity check.
`indexer/index-N.log`	This file contains log messages for index updater process.
`localusers.log`	This file contains log messages for local users scanning process.
`mxpolicy.log`	This file contains log messages for policy evaluation process.
`queryd.log`	This file contains log messages for DataInsightConfig service.
`sharepoint_audit.log`	This file contains log messages for SharePoint audit fetching utility.
`upgradecli0.0.log`	This file contains log messages of the upgrade utility.
`upgrade_to_X.log`	This file contains log messages of the upgrade process.
`watchdog0.0.log`	This file contains log messages of DataInsightWatchdog service.
`winnasd.log.0`	This file contains log messages of DataInsightWinnas service.
`winnas_util.log`	This file contains log messages of windows share discovery utility.
`workflowd0.0.log`	This file contains log messages from the DataInsightWorkflow service.

Downloading Data Insight logs

To troubleshoot errors, you can download the Data Insight logs relevant to a file server, share, SharePoint Web application, or SharePoint site collection from the **Settings** tab of the Management Console.

To download Data Insight logs

1   On the relevant listing page, click the **Select Action** drop-down, and select **Download Logs** for the data repository you want to troubleshoot.

2   On the **Download Logs** pop-up, select the check box for the information that you want to include in the logs.

    You can select one or all of the following information:

- **Config database** - Select this option to include the configuration database in the download. Secret information, such as passwords are purged from the copied database.

- **Indexer database** - Select this option to include the index for the problematic shares or site collections in the download.

- **Error files** - Select this option to includes scan or audit files that have not been indexed in the download.

- **User database** - Select this option to include the cached Active Directory information in the download.

Note: Contact Symantec Support to help you determine which of these options you should select when troubleshooting an issue.

Migrating the data directory to a new location

The data directory is the location where a Data Insight server stores the product data. You specify the location of the data directory during the Data Insight installation. You can find out the current location of the data directory by reading the `datadir.conf` file that is located at `C:\Program Files\Symantec\DataInsight\`.

To move the data directory from its current location

1   Stop all the Data Insight services that are running on the server.

2   Navigate to `C:\Program Files\Symantec\DataInsight\` and open the file `datadir.conf` in a text editor. Note the current location of the data directory. For example, `matrix.datadir=C:/DataInsight/data`.

3   Edit the value for the parameter *matrix.datadir* to indicate the new location of the data directory. For example, `matrix.datadir=E:/DataInsight/data`.

4   Copy the folder,`$DATADIR/data`, from the old location to the new location. For example, copy the folder from the original location `C:/DataInsight` to the new location `E:/DataInsight`.

Note: If you choose to rename the data directory, do not use any space in the filename. Doing so will prevent the Data Insight services from starting.

5   Navigate to `C:\Program Files\Symantec\DataInsight\bin` using the command prompt. Execute the following command:

```
configdb -c
```

6   Verify that the command output points to the new data directory location.

7   Execute the command `configdb -p -T node`.

    Verify that the command output lists all the Data Insight servers that are in your deployment.

8   Start the Data Insight services on the server.

9   After all Data Insight services start successfully, delete the original data directory.

Enterprise Vault exceptions and their meanings

Sometimes, when you initiate an archive operation from the Data Insight Management Console, using the Enterprise Vault, you may encounter error messages. Here is a list of possible error messages and their description:

Table A-2   Symantec Enterprise Vault errors and descriptions

Exception	Description
Exception in method Archive:System. ServiceModel.CommunicationException: The maximum message size quota for incoming messages (65536) has been exceeded. To increase the quota, use the MaxReceivedMessageSize property on the appropriate binding element. ---> System.ServiceModel. QuotaExceededException: The maximum message size quota for incoming messages (65536) has been exceeded. To increase the quota, use the MaxReceivedMessageSize property on the appropriate binding element.	This error can arise when a large number of files are sent for archiving in a single batch. This causes an overflow of the product buffer.

Table A-2   Symantec Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`Exception in method GetFileServerShares: System.ServiceModel.Endpoint NotFoundException: There was no endpoint listening at http://10.209.108.73/EnterpriseVault/ FileSystemArchivingAPI that could accept the message. This is often caused by an incorrect address or SOAP action. See InnerException, if present, for more details. ---> System.Net.WebException: The remote server returned an error: (404) Not Found. 2013-05-07 03:37:51 INFO: #{27} [ArchiveStep.archiveFiles] [ARCHIVE OUT] --- End of inner exception stack trace --- 2013-05-07 03:37:51 INFO: #{27} [ArchiveStep.archiveFiles] [ARCHIVE OUT] at System.Service Model.Channels. HttpChannelUtilities.ProcessGet ResponseWebException(WebException webException, HttpWebRequest request, HttpAbortReason abortReason)`	This error occurs when the Enterprise Vault Admin service is down or the Enterprise Vault server is unreachable. For instance, when the disk on the Enterprise Vault server vault store is full, all the Enterprise Vault services are stopped. This could also occur if the name of Enterprise Vault server, protocol, or port are specified incorrectly. If the AdminService is not down, still this error can arise if the EVFileSvrArcMngr service is not running.

Table A-2 Symantec Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`[Exception in method Archive:System.Service Model.FaultException` `1[www.symantec.com.Enterprise Vault.API.FileSystemArchiving. Data.ExecutionFailedFault]: This request operation sent to net.tcp://evserver1.tulip.matrixad .local:5114/TaskService/ 10E7B53932AB5754980C95B4591BC 72171012f00evserver1 did not receive a reply within the configured timeout (00:30:00). The time allotted to this operation may have been a portion of a longer timeout. This may be because the service is still processing the operation or because the service was unable to send a reply message. Please consider increasing the operation timeout (by casting the channel/proxy to IContextChannel and setting the OperationTimeout property) and ensure that the service is able to connect to the client. (Fault Detail is equal to www.symantec.com.EnterpriseVault. API.FileSystem Archiving.Data.Execution FailedFault).]`	This error occurs when the internal services of the Enterprise Vault times out when you try to archive a batch of files. To resolve this problem, you can consider changing the size of batch to a smaller value. Also you can consider reducing the number of files that are sent in a batch. Alternatively, you can change the timeout values in the configuration for the Enterprise Vault. To change the timeout value: 1 Open the `EvFileSvrArcMngr.exe.config` file. 2 Set the appropriate value for the key "<add key="OperationTimeoutInMin" value = "60"/>

Table A-2   Symantec Enterprise Vault errors and descriptions *(continued)*

Exception	Description
`Executing Function Archive: Exception in method Archive:System.ServiceModel.Fault Exception`1[www.symantec.co m.EnterpriseVault.API.FileSystem Archiving.Data.TimeoutFault]: The File System Ar chiving task service failed to start. Check that the File System Archiving task service is enabled in the configuration file, <Enterprise_Vault_installation_fol der>\EvFSAArchiving Task.exe.config. (Fault Detail is equal to www.symantec.com.E nterpriseVault.API.FileSystem Archiving.Data.TimeoutFault).`	This error occurs at random. To resolve this problem, re-run the failed archive operation. If you are an Data Insight administrator user, you can re-run the archive operation from the **Settings** > **Action Status** page of the Management Console.See "Viewing and managing the status of an operation" on page 234.
`2013-03-21 11:17:27 WARNING: Archive: Got exception while archiving - System.ServiceModel.Fault Exception`1[www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ServerTemporary UnavailableFault]: Unable to contact the Enterprise Vault Task Controller service. Check that the service is running. (Fault Detail is equal to www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ServerTemporary UnavailableFault).`	This error occurs when the Enterprise Vault task controller service cannot be reached for some reason.

Table A-2 Symantec Enterprise Vault errors and descriptions *(continued)*

Exception	Description
```Exception in method Archive:System.ServiceModel.Fault Exception`1[www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.ExecutionFailedFault]: Internal error occurred. Internal Error : <Error checking if file: \\?\UNC\vnxemc.tulip.matrixad.local\ Share4EV\Public Share Copy\Ganesh\di automaation \logs3\DI_Automation\Library\Copy (3) of KeyAction\Action1\SnapShots\ ssf35f0.html.z is a placeholder file.> (Fault Detail is equal to www.symantec.com.Enterprise Vault.API.FileSystem Archiving.Data.Execution FailedFault)```	This error occurs when you attempt to re-archive an already archived file, but the placeholder service on the File System Archiving (FSA) agent is not running. To resolve this problem, run the placeholder services on the FSA agent and attempt to archive the file again.
```[Exception in method Archive: System.Web.Services.P rotocols.SoapException: The socket connection was aborted. This could be caused by an error processing your message or a receive timeout being exceeded by the remote host, or an underlying network resource issue. Local socket timeout was ''''00:29:59.9810000''''., at evClient.Program. Archive(FileSystemArchivingService channel)]```	This error occurs when you restart the Data Insight services, while Data Insight is still processing an archiving operation. To resolve this error, make sure that the File System Archiving (FSA) task and Enterprise Vault services are running.

Command File Reference

This appendix includes the following topics:

- fg.exe

- indexcli.exe

- reportcli.exe

- scancli.exe

- installcli.exe

fg.exe

`fg.exe` – A script that modifies the file group configuration for Data Insight.

SYNOPSIS

```
fg -C -N <name of file group>

fg -D -N <name of file group>

fg -L -d

fg -L -N <name of file group> -d

fg -R -N <name of file group> -t <name of extension>
```

Description

fg is a script used to modify the configuration for sorting files into file groups. By default, Data Insight sorts files into 18 file groups based on the file extensions.

Options

-i *<username>*

    (Required) The fully-qualified user name of the user running the command, for example, user@domian. This user should have Server Administrator privileges in Data Insight.

-A   Adds an extension to an existing file group.

-C   Creates a new file group.

-D   Deletes an existing file group.

-L   Lists existing file groups.

-R   Removes an extension from an existing file group.

-N   Name of the file group to be created or deleted.

-d   Shows file group details when listing existing file groups.

-t *<name of extension>*

    The file extension to add or delete from the file group (For example, doc).

-h   Prints the usage message.

EXAMPLES

EXAMPLE 1: The following command creates a new file group.

```
fg -i <username> -C -N <name of file group>
```

EXAMPLE 2: The following example adds a new extension to an existing file group.

```
fg -i <username> -A -N <name of file group>  -t <name of extension>
```

EXAMPLE 3: The following example deletes an extension from an existing file group.

```
fg -i <username> -R -N <name of file group>  -t <name of extension>
```

EXAMPLE 4: The following command deletes a file group.

```
fg -i <username> -D -N <name of file group>
```

EXAMPLE 5: The following command displays a detailed listing of all configured file groups.

```
fg -i <username> -L -d
```

EXAMPLE 6: The following command displays a detailed listing of a particular file group.

```
fg -i <username> -L -N <name of file group> -d
```

—

indexcli.exe

`indexcli.exe` – a utility that manages the index segments available on an Indexer
worker node.

SYNOPSIS

```
indexcli.exe
            --display|--archive|--purge|--restore|--rearchive|--list-jobs
            |--stop-jobs [OPTIONS]
indexcli.exe -A <name of the index segments to be archived>
indexcli.exe -c
indexcli.exe -D <name of the index segments to be purged>
indexcli.exe -d
indexcli.exe -h
indexcli.exe -j
indexcli.exe -r
indexcli.exe -t
indexcli.exe -u
```

Archive options

```
indexcli.exe -A -a ¦ -f <FILERS> ¦ -m
<SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦ -I
 <MONTHS>
```

`-a` Archives all index segments older than the specified interval.

`-f <name of filer(s)>`
  Archives all index segments for the specified list of filers.

`-I <interval in months>`
  Archives segments older than the specified interval. The segments which have
  been restored earlier are not archived.

`-m <name of share(s)>`
  Archives all index segments for the specified list of shares.

`-S,--sitecoll <`*`SITECOLLS`*`>`

Archives segments for specified list of Microsoft SharePoint site collections.

`-w, --webapp <`*`WEBAPPS`*`><`

Archives segments for specified list of Microsoft SharePoint Web applications.

Purge options

```
indexcli.exe -D -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦
-I <MONTHS>
```

`-a`   Purges all index segments older than the specified interval.

`-f <`*`name of filer(s).`*

Purges all index segments for the specified list of filers.

`-I ,`*`interval in months.`*

Purges segments older than the specified interval. The segments which have been explicitly restored earlier, for which the lease is still valid, are not purged.

`-m <`*`name of share(s)>`*

Purges all index segments for the specified list of shares.

`-S,--sitecoll <`*`SITECOLLS`*`>`

Purges segments for specified list of Microsoft SharePoint site collections.

`-w, --webapp <`*`WEBAPPS`*`><`

Purges segments for specified list of Microsoft SharePoint Web applications.

Display options

```
indexcli.exe -d -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> ¦
-s <STATES>
```

`-a`   Displays information for all shares.

`-f <`*`name of filer(s)>`*

Displays information for the specified list of filers.

`-m <`*`name of share(s)>`*

Displays information for the specified list of shares.

-s *<name of state>*

Displays index segments for the given state only. Multiple stars can be separated by comma. Possible states are, ARCHIVING, RE-ARCHIVING, ARCHIVED, RESTORING, RESTORED, RESTORE, FAILED, or DELETED.

-S,--sitecoll *<SITECOLLS>*

Displays information for a specified list of Microsoft SharePoint site collections.

-w, --webapp *<WEBAPPS><*

Displays information for a specified list of Microsoft SharePoint Web applications.

Restore options

```
indexcli.exe -r -a ¦ -f <FILERS> ¦
-m <SHARES> ¦ -S <SITECOLLS> ¦ -w <WEBAPPS> -C -F <FROM> ¦ >
¦ -R <RANGE>
[-L <MONTHS> ¦ -l <MONTHS> ¦ -y]
```

-a   Restores the index segments for all shares.

-C   If the continue-on-error option is not specified, the restore command fails if the segment files required to restore data for the specified parameters are not available.

-f *<name of filer(s)>*

Restores all index segments for the specified list of filers.

-F *<month from which the segments need to be restored>*

Specify the month in the format, YYYY/MM. For example, indexcli.exe -r -F 2010/01 restores segments from January 2010 till date.

-L *<interval in number of months>*

Resets lease on segments that are restored earlier using -l option. Specify the new lease interval in months. This command replaces the previous lease interval. Setting the value to 0 will make the lease permanent.

-l *<interval in number of months>*

Restores segments for a temporary lease in months. After the lease expires, restored segments are automatically re-archived. If this option is not specified, segments remain restored till you re-archive them with the -u, --rearchive, option.

-m *<name of share(s)>*

Restores all index segments for the specified list of shares.

-S,--sitecoll *<SITECOLLS>*

    Restores segments for a a specified list of Microsoft SharePoint site collections.

-w, --webapp *<WEBAPPS>*<

    Restores segments for a specified list of Microsoft SharePoint Web applications.

-R *<range in months>*

    Restore all index segments for the specified month range. Specify the month in the format, YYYY/MM-YYYY/MM. For example, `indexcli.exe -r -R 2010/01-2010-03` restores segments from January 2010 to March 2010.

-y   Instead of restoring segments, this option displays the list of files that must be available before restoring the specified segments.

Re-archive options

```
indexcli.exe -u -a ¦ -f <FILERS> ¦
-m <SHARES>  ¦ -S <SITECOLLS> ¦ -w <WEBAPPS>
-F <FROM> ¦ -R <RANGE>
```

-a   Re-archives all previously restored index segments.

-f *<name of filer(s)>*

    Re-archives previously restored index segments for the specified list of filers.

-F *<month FROM which the segments need to be restored>*

    Specify the month in the format, YYYY/MM. For example, `indexcli.exe -u -F 2010/01` restores segments from January 2010 till date.

-m *<name of share(s)>*

    Re-archives previously restored index segments for specified list of shares.

-S,--sitecoll *<SITECOLLS>*

    Re-archives previously restored segments for a specified list of Microsoft SharePoint site collections.

-w, --webapp *<WEBAPPS>*<

    Re-archives previously restored segments for a specified list of Microsoft SharePoint Web applications.

-R *<range in months>*

    Restore all index segments for the specified month range. Specify the month in the format, YYYY/MM-YYYY/MM. For example, `indexcli.exe -u -R 2010/01-2010-03` restores segments from January 2010 to March 2010.

EXAMPLES

EXAMPLE 1: The following command archives index segments for specified list of filers.

```
indexcli.exe -A -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 2: The following command archives index segments for specified list of shares.

```
indexcli.exe -A -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 3: The following command purges index segments for specified list of filers.

```
indexcli.exe -D -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 4: The following command purges segments for specified list of shares.

```
indexcli.exe -D -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 5: The following command restores index segments for specified list of filers.

```
indexcli.exe -r -f <\\filer1,\\filer2,ID1,ID2>
```

EXAMPLE 6: The following command restores index segments for specified list of shares.

```
indexcli.exe -r -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 7: The following command re-archives previously restored index segments for specified list of filers.

```
indexcli.exe -u -f \\filer1,\\filer2,ID1,ID2
```

EXAMPLE 8: The following command re-archives previously restored index segments for specified list of shares.

```
indexcli.exe -u -m \\filer1\share1,\\filer2\shares2,ID3,ID4
```

EXAMPLE 9: The following command archives segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 10: The following command archives segments for specified list of Microsoft SharePoint Web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 11: The following command purges segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 12: The following command purges segments for specified list of Microsoft SharePoint Web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 13: The following command displays information for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 14: The following command displays information for specified list of Microsoft SharePoint Web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 15: The following command restores segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 16: The following command restores segments for specified list of Microsoft SharePoint Web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

EXAMPLE 17: The following command re-archives previously RESTORED segments for specified list of Microsoft SharePoint site collections.

```
indexcli.exe -S,--sitecoll<http://sp_webapp:8000/sc1,ID2,ID3...>
```

EXAMPLE 18: The following command re-archives previously RESTORED segments for specified list of Microsoft SharePoint Web applications.

```
indexcli.exe - w,--webapp<http://sp_webapp:8000,ID2,ID3,...>
```

reportcli.exe

`reportcli.exe` – a utility to create reports using a properties file that contains the input parameters, execute and list configured reports, check the status of the reports, and cancel report runs.

SYNOPSIS

```
reportcli.exe --list-jobs¦──list-reports¦--list-outputs¦--create
--execute¦--cancel¦--help [OPTIONS]

reportcli.exe -c

reportcli.exe -e

reportcli.exe -h

reportcli.exe -j

reportcli.exe -l

reportcli.exe -o
```

Options

```
reportcli.exe -n -r <name of report> -p <property file path> -u <user
name of creator> [-rt <report type>] [--users <path of users' .csv
file>] [-t <path of .csv file of paths>] [--custodian <path of
custodian' .csv file>]
```

Creates a report using the properties file in which the input parameters are specified. The following attributes apply:

`-r --report <name of report>`	Creates a report with the specified name.
`-p --properties <property file path>`	Property file containing the input parameters for the report. **Note:** By default, a sample properties file is installed in the INSTALL_DIR/DataInsight/reports folder.
`-u --creator <user name of creator>`	Creator of the report.
`-rt --type <report_type>`	Creates the specified report type. For example, Access Details report.

`--users <`*`path of users' .csv`* *`file>`*	Path of the .csv file containing the names of users in the user@domain,<user group> format.
	Specify the path to input user information for the report.
`--paths <`*`path of .csv file of`* *`paths>`*	Path of the .csv file containing the fully qualified paths of the data for which you want to create the report.
`--custodian <`*`path of custodian'`* *`.csv file>`*	Path of the .csv file that contains information about custodians on configured paths. The names of the custodians are specified in the .csv file in the format user@domain.
	Specify the path to the custodian.csv file to include custodian information in the report.

`-j`   Lists the report jobs that are currently running.

`-l`   Lists all configured reports.

`-o -m <`*`TOP_N`*`> -r <`*`Report Name`*`>`

  Lists all report outputs. The following attributes apply:

`-m --max <`*`TOP_N`*`>`	Limits output to specified number of records, and lists the latest output first. If the number of records is not specified, prints status for the last run.
`-r - -report <`*`Report Name`*`>`	Prints the status of jobs for the specified report. You can either specify the report ID or the report name.
`-rt - -type<`*`Report Type`*`>`	Prints the status of jobs for the specified report type.

`report.exe -e [-d <`*`Output_Dir`*`> -r <`*`Report Name`*`> -w <`*`Max_Wait`*`>`

  Executes report. The following attributes apply:

`-d --output <`*`Output_Dir`*`>`	The generated report output, including the SQLite database is copied to the specified directory. If you specify this option, you do not have to pass the –-w option.

`-r --report<Report Name>`	Executes the specified report. You can either specify the report ID or the report name.
`--w --wait <Max_Wait>`	Returns the report output only after the report execution is complete or the specified wait time in minutes is exceeded. Specify –1 to wait forever.

```
reportcli.exe -e -r <name of report> [-rt <report type>]--p <property
file path> --creator <user name of creator> [--output <Output_Dir>]
[--users <path of users' .csv file>] [--paths <path of .csv file >]
[--wait <MAX_WAIT>] [--custodian <path of custodian' .csv file>]
```

Executes a report using the properties file in which the input parameters are specified. The following attributes apply:

`-r --report <name of report>`	Creates a report with the specified name.
`-p -properties <property file path>`	Property file containing the input parameters for the report.
`-u --creator <user name of creator>`	Creator of the report.
`--custodian <path of custodian.csv>`	Path of the .csv file containing the names of the custodians
`-d --output <OUTPUT_DIR>`	Copies the report database and output to specified location.
`-rt --type <report_type>`	Creates the specified report type. For example, Access Details report.
`--users <path of users' .csv file>`	Path of the .csv file containing the names of users in the user@domain,<user¦group> format.
`--paths <path of .csv file >`	Path of the .csv file containing the fully qualified paths of the data for which you want to create the report.
`-w --wait<MAX_WAIT>`	If the wait time is specified, the command returns only after the report is executed OR the specified wait time in minutes is exceeded. Specify -1 to wait forever. Data Insight cancels the report execution if the wait time is exceeded.

`--custodian <`*`path of custodian'`* *`.csv file>`*	Path of the .csv file that contains information about custodians on configured paths. The names of the custodians are specified in the .csv file in the format user@domain.
	Specify the path to the custodian.csv file to include custodian information in the report.

```
Reportcli.exe –sa op_excl_path [-ap <FILE_PATH>] –at <ATTRIBUTE NAME>
-av <ATTRIBUTE VALUE>] -ty <ATTRIBUTE TYPE> --path –r <name of report>
-rt <report type>
```

Creates a report by excluding the paths specified in the `.csv` file for computing open shares data on the Data Insight dashboard. The following attributes apply:

`– ap <`*`file path>`*	Path of the .csv file containing the fully qualified paths that you want to exclude from open share computation for the dashboard.
`-rt --type <`*`report_type>`*	Creates the specified report type.
`-r --report <`*`name of report>`*	Creates a report with the specified name.

You can exclude any path on a file server or a SharePoint server from the dashboard data computation. However, Data Insight does not support the exclusion of DFS paths using this method.

```
report.exe –c –i <JOB_ID>
```

Cancels execution of the specified report job.

scancli.exe

scancli.exe – scancli.exe - a utility that scans shares and site collections.

SYNOPSIS

scancli.exe --start¦ --stop¦ --list-jobs¦ --help [OPTIONS]

-s --start

Scans the specified shares or site collections.

-c --stop

Cancels the scans for specified shares or site collections.

-l --list-jobs

Lists currently running jobs.

-d --display

Displays the scan status for specified shares or site collections. To view real time scan queue information, use the -l --list-jobs option.

-h --help

Displays help.

Scan options

scancli.exe -s -a ¦ -f <FILERS> ¦ -m <SHARES> ¦ -S <SITECOLLS> ¦w <WEBAPPS>
[-D] [-e <EXCLUDE] [-F ¦ -N ¦ -p] [-I <INCLUDE>] [-i <DAYS>] [-t]

-a - - all

Scans all shares and site collections.

-D - -disabled

By default, disabled devices or those for which scanning has been disabled are not included in the scan. Specify this option to include shares or site collections of disabled devices.

-e - -exclude <EXCLUDE>

Exclude shares or site collections matching specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern, for example, vol*,*$.

`-f - -filer <FILERS>`

Scans shares of the specified filers. For example, `-f - -filer >\\filer1, filer2, ID1,..>`.

`-F - -failed`

Select the shares or site collections whose last scan failed completely. This does not include those shares or site collections that have never been scanned before or those which succeeded partially (*).

`-I - -Include <INCLUDE>`

Include the shares or site collections matching the specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern. For example, `-I - -Include >vol*,*$ >`

`-i - -interval <DAYS>`

Select the shares or site collections that have not been scanned for specified number of days. This includes shares or site collections which have never been scanned before (*).

`-m - -share <SHARES>`

Scans specified list of shares. For example, `-m - -share >\\filer1\share1, share2, ID3...>`.

`-n - -never`

Select the shares or site collections that have never been scanned before (*).

`-p - -partial`

Select the shares or site collections whose last scan succeeded partially, that is, those shares or site collections for which the scan is complete but with failure to fetch information for some paths (*).

`-S - -sitecoll <SITECOLLS>`

Scans the specified list of Microsoft SharePoint site collections.

`- t - -top`

Adds shares or site collections to top of the scan queue.

`-w - -webapp <WEBAPPS>`

Scans site collections for specified list of Microsoft SharePoint Web applications.

Note: (*) indicates that the option can only be used on the Management Server.

Stop scan options

```
scancli.exe -l -a ¦ -f <FILERS> ¦ -m <SHARES> ¦ -S <SITECOLLS> ¦w <WEBAPPS>
[-D] [-e <EXCLUDE] [-I <INCLUDE>]
```

`-a – – all`

>   Stops scans for all shares and site collections.

`-e – –exclude <EXCLUDE>`

>   Exclude shares or site collections matching specified patterns. Separate multiple
>   patterns with a comma. You can specify one or more wildcards in the pattern,
>   for example, vol*,*$.

`-f – –filer <FILERS>`

>   Stops scans for shares of the specified filers.

`-I – –Include <INCLUDE>`

>   Include shares or site collections matching the specified patterns. Separate
>   multiple patterns with a comma. You can specify one or more wildcards in the
>   pattern.

`-m – –share <SHARES>`

>   Stops scans for the specified list of shares.

`-S – –sitecoll <SITECOLLS>`

>   Stops scans for the specified list of Microsoft SharePoint site collections.

`-w – –webapp <WEBAPPS>`

>   Stops scans for site collections for specified list of Microsoft SharePoint Web
>   applications.

List job options

```
scancli.exe –l [–n --node <NODE>]
```

`-n --node <Node ID or Node name>`

>   Lists scan jobs on the specified node. Specify either node ID or node name. If
>   not specified, localnode is assumed.

Display options

```
scancli.exe –d –a ¦ –f <FILERS> ¦ –m <SHARES> ¦ –S <SITECOLLS> ¦w <WEBAPPS>
[-D] [–e <EXCLUDE] [-F ¦ –N ¦ -p] [-I <INCLUDE>] [-i <DAYS>]
```

`-a – – all`

>   Displays scan status for all shares and site collections.

`-e – –exclude <EXCLUDE>`

>   Exclude shares or site collections matching specified patterns. Separate multiple
>   patterns with a comma. You can specify one or more wildcards in the pattern,
>   for example, vol*,*$.

`-f - -filer <FILERS>`

> Displays scan status for the shares of the specified filers.

`-F - -failed`

> Displays scan status for the shares or site collections whose last scan failed completely. The scan status does not include those that have never been scanned before or those which succeeded partially (*).

`-I - -Include <INCLUDE>`

> Include shares or site collections matching the specified patterns. Separate multiple patterns with a comma. You can specify one or more wildcards in the pattern.

`-i - -interval <DAYS>`

> Displays scan status for the shares or site collections that have not been scanned for specified number of days. The scan status includes the shares which have never been scanned before (*).

`-m - -share <SHARES>`

> Displays scan status for specified list of shares.

`-n - -never`

> Displays scan status for the shares or site collections that have never been scanned before (*).

`-p - -partial`

> Displays scan status for the shares or site collections whose last scan succeeded partially, that is, those shares or site collections for which the scan is complete but with failure to fetch information for some paths (*).

`-S - -sitecoll <SITECOLLS>`

> Displays scan status for the specified list of Microsoft SharePoint site collections.

`-w - -webapp <WEBAPPS>`

> Displays scan status for the site collections for specified list of Microsoft SharePoint Web applications.

Note: `-w - -webapp <WEBAPPS>` option can only be used on the Management Server.

Examples

EXAMPLE 1: The following command scans all shares of a filer, netapp1.

```
scancli - -start - - filer <netapp1>
```

EXAMPLE 2: The following command scans all shares and site collections for which a full scan failed 3 or more days ago.

```
scancli – –start – –all – –failed – –ineterval <3>
```

The following command scans all site collections of a Web application that have not been scanned for the past 30 days or have never been scanned.

```
scancli – –start – –webapp https://sitecoll:8080 – –interval 30
```

installcli.exe

`installcli.exe` – A utility that is used to configure multiple Windows File Servers and Data Insight worker nodes simultaneously.

SYNOPSIS

```
installcli [-w winnas_csv [-q]] [-n node_csv [-q]] [-p operation_token]
[-l] [-h]
```

Options

`-w --winnas winnas_csv`

Installs Data Insight Windows File Server agents and configures the corresponding filer.

-w option uses a `.csv` file with the following details as input:

- The host name or IP address of the Windows File Server that you want Data Insight to monitor.

- The host name, IP address, or ID of the Collector node that is configured to scan the filer.

- The host name, IP address, or ID of the Indexer node that is configured for the filer.

- The credentials that Data Insight should use to install the agent on the Windows File Server. The credential should be in the format user@domain. installcli.exe also accepts Local System credentials as value _LOCAL_. The same credentials must be added to Data Insight as a saved credential previously.

- True or false value indicating if the filer is clustered.

- The IP addresses of the agents. Separate multiple IP addresses with a semi-colon. If you do not want to use an agent to monitor the filer, indicate this with a hyphen (-).

- The credentials required to scan the filer. The credential should be in the format user@domain. The same credentials should be added to Data Insight as a saved credential previously.

- True or false value indicating whether the scan should be enabled according to the specified schedule.

- True or false value indicating whether event monitoring should be enabled. For example,

winnas.company.com,collector.company.com,indexer.company.com,
Administrator@DOMAIN,FALSE,winnas.company.com,
Administrator@DOMAIN,TRUE,TRUE,RP,
Symantec_DataInsight_windows_winnas_4_0_0_3002_x64.exe.

- In case of a Windows File Server agent upgrade, RP or Full value indicating the type of upgrade you want to perform. This parameter is optional.

- Optionally, the name of the installer. If not specified, an appropriate one will be picked up from installers folder on the collector.

`-n --node node_csv`

Installs the Data Insight Collector and Indexer nodes. The `node_csv` file must be in the following format:

- The host name or IP address of the worker node that you want to install or upgrade.

- True or false value indicating whether the node is a Collector. Since Data Insight does not currently support push-install on Linux nodes, you must specify *true* as that value for this column. I

- True or false value indicating whether the node is a Indexer.

- The credentials that Data Insight should use to install the package on the worker node. The credential should be in the format user@domain. The same credentials must be added to Data Insight as a saved credential previously.

- The port used by the DataInsightComm service.

- The port used by DataInsightConfig service.

- The destination directory where you want Data Insight to be installed.

- The location where you want to store the product data.

The values for the Communication service port, query service port, the installation path, and the data directory are optional. You can enter ? to use default values.

`-p --poll operation_token`

Starts polling for the latest status of an operation.

`-l --list`

Lists status and progress information of all currently running and historic operations.

`-q --nowait`

Forks off an operation and does not wait for it to complete.

`-h --help`

Displays help.

Data Insight jobs

This appendix includes the following topics:

■ Scheduled Data Insight jobs

Scheduled Data Insight jobs

Each Data Insight service performs several actions on a scheduled basis. These are called jobs. The section explains the function of the important jobs that run in various services. The schedule for few jobs can be changed from the **Advanced Settings** tab of the Server details page.

Table C-1   Communication service jobs

Job	Description
ADScanJob	Initiates the `adcli` process on the Management Server to scan the directory servers. Ensure the following: ■ The directory servers are added to the Data Insight configuration. ■ The credentials specified when adding the directory server have permissions to scan the directory server.
CollectorJob	Initiates the collector process to pre-process raw audit events received from storage devices. The job applies exclude rules and heuristics to generate audit files to be sent to the Indexers. It also generates change-logs that are used for incremental scanning.
ChangeLogJob	The CollectorJob generates `changelog` files containing list of changed paths, one per device, in the `changelog` folder. There cab be multiple files with different timestamps for each device. The ChangeLogJob merges all `changelog` files for a device.

Table C-1   Communication service jobs *(continued)*

Job	Description
ScannerJob	Initiates the scanner process to scan the shares and site collections added to Data Insight.
	Creates the scan database for each share that it scanned in the `data\outbox` folder.
IScannerJob	Intiates the incremental scan process for shares or site-collections for paths that have changed on those devices since the last scan.
CreateWorkflowDBJob	Runs only on the Management Server. It creates the database containing the data for DLP Incident Management, Entitlement Review, and Ownership Confirmation workflows based on the input provided by users.
DlpSensitiveFilesJob	Retrieves policies and sensitive file information from Data Loss Prevention (DLP).
FileTransferJob	Transfers the files from the `data\outbox` folder from a node to the `inbox` folder of the appropriate node.
FileTransferJob_Evt	Sends Data Insight events database from the worker node to the Management Server.
FileTransferJob_WF	Transfers workflow files from Management Server to the Portal service.
IndexWriterJob	Runs on the Indexer node; initiates the idxwriter process to update the Indexer database with scan (incremental and full), tags, and audit data.
	After this process runs, you can view newly added or deleted folders and recent access events on shares on the Management Console.
ActivityIndexJob	Runs on the Indexer node; It updates the activity index every time the index for a share or site collection is updated.
	The Activity index is used to speed up the computation of ownership of data.
IndexCheckJob	Verifies the integrity of the index databases on an Indexer node.
PingHeartBeatJob	Sends the heartbeat every minute from the worker node to the Data Insight Management Server.

Table C-1 Communication service jobs *(continued)*

Job	Description
PingMonitorJob	Runs on the Management Server. It monitors the heartbeat from the worker nodes; sends notifications in case it does not get a heartbeat from the worker node.
SystemMonitorJob	Runs on the worker nodes and on the Management Server. Monitors the CPU, memory, and disk space utilization at a scheduled interval. The process sends notifications to the user when the utilization exceeds a certain threshold value.
DiscoverSharesJob	Discovers shares or site collections on the devices for which you have selected the **Automatically discover and monitor shares on this filer** check box when configuring the device in Data Insight
ScanPauseResumeJob	Checks the changes to the pause and resume settings on the Data Insight servers, and accordingly pauses or resumes scans.
DataRetentionJob	Enforces the data retention policies, which include archiving old index segments and deleting old segments, indexes for deleted objects, old system events, and old alerts.
IndexVoldbJob	Runs on the Management Server and executes the command voldb.exe --index which consumes the device volume utilization information it receives from various Collector nodes.
SendNodeInfoJob	Sends the node information, such as the operating system, and the Data Insight version running on the node to the Management Server. You can view this information on the Data Insight Server > Overview page of the Management Console.
EmailAlertsJob	Runs on the Management Server and sends email notifications as configured in Data Insight.The email notifications pertain to events happening in the product, for example, a directory scan failure. You can view them on the **Settings** > **System Overview** page of the Management Console.
LocalUsersScanJob	Runs on the Collector node that monitors configured file servers and SharePoint servers. In case of a Windows File Server that uses agent to monitor access events, it runs on the node on which the agent is installed. It scans the local users and groups on the storage devices.

Table C-1 Communication service jobs *(continued)*

Job	Description
UpdateCustodiansJob	Runs on the Indexer node and updates the custodian information in the Data Insight configuration.
CompactJob	Compresses the `attic` folder and `err` folders in `<datadir>\collector`, `<datadir>\scanner`, and `<datadir>\indexer` folders. The process uses the Windows compression feature to set the "compression" attribute for the folders.
	The job also deletes stale data that's no longer being used.
Compact_Job_Report	Compresses the folders that store report output.
StatsJob	On the Indexer node, it records index size statistics to `lstats.db`. The information is used to display the filer statistics on the Data Insight Management Console.
MergeStatsJob	Rolls up (into hourly, daily and weekly periods) the published statistics. On the Collector nodes for Windows Filer Server, the job consolidates statistics from the filer nodes.
StatsJob_Index_Size	Publishes statistics related to the size of the index.
StatsJob_Latency	On the Collector node, it records the filer latency statistics for NetApp filers.
SyncScansJob	Gets current scan status from all Collector nodes. The scan status is displayed on the **Settings** >**Scanning Dashboard** > **In-progress Scans** tab of the Management Console.
SPEnableAuditJob	Enables auditing for site collections (within the web application), which have been added to Data Insight for monitoring.
	By default, the job runs every 10 minutes.
SPAuditJob	Collects the audit logs from the SQL Server database for a SharePoint web application and generates SharePoint audit databases in Data Insight.
SPScannerJob	Scans the site collections at the scheduled time and fetch data about the document and picture libraries within a site collection and within the sites in the site collection.

Table C-1   Communication service jobs *(continued)*

Job	Description
NFSUserMappingJob	Maps every UID in raw audit files for NFS and VxFS with an ID generated for use in Data Insight. Or generates an ID corresponding to each User and Group ID in raw audit files received from NFS/VxFS.
MsuAuditJob	Collects statistics information for all indexers on the Indexer.
MsuMigrationJob	Checks whether a filer migration is in process and carries it out.
ProcessEventsJob	Processes all the Data Insight events received from worker nodes and adds them to the yyyy-mm-dd_events.db file on the Management Server.
ProcessEventsJob_SE	Processes scan error files.
SpoolEventsJob	Spools events on worker nodes to be sent to Management Server.
WFStatusMergeJob	Merges the workflow and action status updates for remediation workflows (DLP Incident Remediation, Entitlement Reviews, Ownership Confirmation), Enterprise Vault archiving, and custom actions and update the master workflow database with the details so that users can monitor the progress of workflows and actions from the Management Console.
UpdateConfigJob	Reconfigures jobs based on the configuration changes made on the Management Server.

The following processes run in the Data Insight WatchDog service

Table C-2   WatchDog service processes

Job	Description
SyncPerformanceStatsJob -	Runs only on the Management server. Fetches performance related statistics from all other servers.
SystemMonitorJob	Gathers statistics like disk usage, CPU, memory usage.
SystemMonitorJob_backlog	Gathers statistics for unprocessed backlog files.

Table C-2   WatchDog service processes *(continued)*

Job	Description
UpdateConfigJob	Reconfigures its own jobs based on configuration updates from the Management Server.

The following processes run in the Data Insight Workflow service

Table C-3   Workflow service processes

Job	Description
WFStepExecutorJob	Processes actions for Enterprise Vault archiving, requests for permission remediation, and custom actions configured in Data Insight.
WFStepExecutorJob_im	Processes workflows of type Entitlement Reviews, DLP Incident Remediation, and Ownership confirmation. It also sends email reminders containing links to the remediation portal to the custodians at a specified interval.
UpdateConfigJob	Updates its schedules based on the configuration changes made on the Management Server.
WFSpoolStatusJob	Reads the workflow data every minute, and if there are any new updates in last minute, it creates a status database with the new updates.
FileTransferJob_WF	Transfer workflow status databases from the Self-Service portal nodes to the Management Server.

The following processes run in the Data Insight Webserver service.

Table C-4 Webserver service processes

Job	Description
CustodianSummaryReportJob	Periodically runs the custodian summary report, which is used to determine the custodians assigned in Data Insight for various resources. The output produced by this report is used in DLP Incident Remediation, Entitlement Review, and Ownership Confirmation workflows.
HealthAuditReportJob	Periodically creates a report summarizing health of the entire deployment, and stores it to `log/health_audit` folder on the Management Server. The report aids Symantec Support to troubleshoot issues on your setup.
PolicyJob	Evaluates configured policies in the system and raises alerts.
PurgeReportsJob	Deletes older report outputs.
UpdateConfigJob	Updates configuration database on the worker nodes based on the configuration changes made on the Management Server.
UserIndexJob_merge	Consolidates user activity and permission map from all indexers.
UserIndexJob_split	Requests each Indexer for user activity and permission map.

See "Monitoring Data Insight jobs" on page 214.

Index