

IBM StoredIQ Administrator

Administration Guide



Note

Before using this information and the product it supports, read the information in [Notices](#).

This edition applies to Version 7.6.0.22 of product number 5724M86 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2001, 2020.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

- Tables..... V**

- About this publication..... vii**
 - IBM StoredIQ product library.....vii
 - Contacting IBM StoredIQ customer support..... vii

- Overview of IBM StoredIQ Administrator..... 1**

- Logging in to IBM StoredIQ Administrator..... 3**

- Data Servers and Volumes..... 4**
 - DataServer - Distributed..... 6
 - Adding a primary volume..... 7
 - Adding a retention volume..... 32
 - Adding a discovery export volume..... 33
 - Editing a volume.....34
 - Viewing details of multiple volumes..... 35
 - Restarting services on a data server..... 35
 - Rebooting a data server.....35
 - Removing a volume.....36

- Harvests.....37**
 - Harvesting a volume..... 37
 - Searching for a specific harvest.....38
 - Stopping a harvest..... 38
 - Editing a harvest..... 38
 - Deleting a harvest..... 38
 - Reviewing harvest information..... 38

- System infosets..... 40**
 - Creating a system infoset..... 40
 - Limiting access to data..... 41

- User management..... 43**
 - Managing IBM StoredIQ user accounts..... 43
 - Managing user passwords..... 46

- Action management.....48**
 - Creating an action..... 49
 - Managing actions..... 52

- Target sets..... 53**
 - Managing target sets..... 53

- Reports..... 55**
 - Configuring report settings..... 55
 - Creating a report..... 56
 - Custom reports..... 57

Auto-classification models.....	61
Applying the auto-classification feature.....	62
Data server sizing-configuration guidelines.....	62
Downloading the classification export utility.....	63
Building an auto-classification model.....	63
Preparing an IBM Content Classification model for import.....	63
Importing an auto-classification model.....	64
Selecting a classification model.....	64
Searching for a specific classification model.....	65
Editing a model.....	65
Deleting a model.....	66
Retraining a classification model.....	66
Best practices for creating an Auto-classification model.....	67
Cartridges.....	70
Cartridge types.....	70
Deploying cartridges.....	70
Uploading cartridges.....	71
Updating cartridges.....	72
Deleting cartridges.....	73
IBM StoredIQ Connector operations.....	74
Auditing.....	88
Audit events.....	88
Configuring auditing on the AppStack.....	90
Viewing audit entries in the database.....	92
Administering IBM StoredIQ components.....	95
The system administrator account	95
Administration commands.....	95
Administration files.....	97
Notices.....	99
Trademarks.....	100
Terms and conditions for product documentation.....	101
IBM Online Privacy Statement.....	101
Index.....	103

Tables

- 1. Data Dashboard fields and descriptions..... 4
- 2. The Data Servers pane column names and descriptions..... 4
- 3. System Status and System Activity panes field names and descriptions..... 4
- 4. The Volumes pane column names and descriptions..... 5
- 5. Box primary volumes..... 7
- 6. CIFS/SMB, SMB2, or SMB3 primary volumes..... 8
- 7. Connections primary volumes..... 10
- 8. CMIS primary volumes..... 11
- 9. Documentum primary volumes..... 12
- 10. Exchange primary volumes..... 13
- 11. FileNet primary volumes..... 15
- 12. HDFS primary volumes..... 16
- 13. IBM Content Manager primary volumes..... 18
- 14. Livelink primary volumes..... 19
- 15. NFS primary volumes..... 20
- 16. NewsGator primary volumes..... 21
- 17. OneDrive primary volumes..... 22
- 18. SharePoint primary volumes..... 23
- 19. Mapping permission from CIFS to Box Collaboration Roles..... 28
- 20. Mapping permission from SharePoint to Box Collaboration Roles..... 29
- 21. Retention volume options..... 32
- 22. Discovery export volume options..... 33
- 23. List of Harvests fields and descriptions..... 37

24. Infoset management column names and descriptions.....	40
25. Actions column names and descriptions.....	49
26. Classification model column names and descriptions.....	64
27. Row names and descriptions of the <model name> Details panel.....	65
28. Cartridges.....	71
29. IBM StoredIQ Connector operations.....	74

About this publication

IBM® StoredIQ® Administrator Administration Guide describes how to manage the IBM StoredIQ Administrator application to configure shares, create system infosets, manage users, target sets, and actions, and create reports.

IBM StoredIQ product library

The following documents are available in the IBM StoredIQ product library.

- *IBM StoredIQ Overview Guide*
- *IBM StoredIQ Deployment and Configuration Guide*
- *IBM StoredIQ Data Server Administration Guide*
- *IBM StoredIQ Administrator Administration Guide*
- *IBM StoredIQ Data Workbench User Guide*
- *IBM StoredIQ Cognitive Data Assessment User Guide*
- *IBM StoredIQ Insights User Guide*
- *IBM StoredIQ Integration Guide*

The most current version of the product documentation can always be found online: https://www.ibm.com/support/knowledgecenter/en/SSSHEC_7.6.0/welcome/storediq.html

Contacting IBM StoredIQ customer support

For IBM StoredIQ technical support or to learn about available service options, contact IBM StoredIQ customer support at this phone number:

- 1-866-227-2068

Or, see the Contact IBM web site at <http://www.ibm.com/contact/us/>.

IBM Knowledge Center

The IBM StoredIQ documentation is available in [IBM Knowledge Center](#).

Contacting IBM

For general inquiries, call 800-IBM-4YOU (800-426-4968). To contact IBM customer service in the United States or Canada, call 1-800-IBM-SERV (1-800-426-7378).

For more information about how to contact IBM, including TTY service, see the Contact IBM website at <http://www.ibm.com/contact/us/>.

Overview of IBM StoredIQ Administrator

IBM StoredIQ Administrator provides at-a-glance understanding of the different issues that can crop up in the IBM StoredIQ environment. These views are unique to the IBM StoredIQ Administrator application, providing an overview of how the system is running. These views also allow access to various pieces of information that are shared across applications or allow for the management of resources in a centralized manner.

An administrator is the person responsible for managing the IBM StoredIQ installation at a customer site. This individual has strong understanding of data sources, harvests, indexes, data servers, jobs, infosets, reports, concepts, and actions. The following list provides an overview of the tasks that the administrator can accomplish with IBM StoredIQ Administrator:

Managing data servers

Identify what data servers are deployed, their location, what data is being managed, and the status of each data server in the system. Volume management is a central component of IBM StoredIQ. The administrator can see:

- What volumes are currently under management.
- Which data server is responsible for that volume.
- The state of the volume after index.
- The amount and size of information that is contained by each volume.
- Whether a volume is published to the governance catalog.

Administrators can also add volumes to data servers through this interface.

Administering harvests

Initiate the incremental or full harvest of a single volume through the IBM StoredIQ Administrator user interface without having to log in to IBM StoredIQ Data Server.

Creating system infosets

Create and manage system infosets that use only specific indexed volumes within IBM StoredIQ Administrator. Although infosets are a core component of IBM StoredIQ Data Workbench, system infosets are created as a shortcut for users in IBM StoredIQ Administrator.

Managing users

Create users and manage users' access to the various IBM StoredIQ applications:

- IBM StoredIQ Administrator
- IBM StoredIQ Data Workbench
- IBM StoredIQ Insights
- IBM StoredIQ Policy Manager

Configuring and managing actions

Create and manage actions which are processes upon the data that is represented by the indexes. Actions are run by data servers on indexed data objects. Any errors or warnings that are generated as a result of an action are recorded as exceptions in IBM StoredIQ Data Workbench.

Actions are created within IBM StoredIQ Administrator and then made available to other IBM StoredIQ applications such as IBM StoredIQ Data Workbench.

Managing target sets

Define and manage wanted targets for specific actions that require a destination volume for their actions.

Creating reports

Create reports and upload report packages, helping you produce reports about your environment.

Using auto-classification models

Deploy natural language processing-based document classification to help you identify relevant documents in cases when relevance is difficult to characterize with traditional methods, such as queries and set operations.

Managing cartridges

Upload cartridges for use in Step-up Analytics actions, update existing cartridges, or delete unused ones.

Logging in to IBM StoredIQ Administrator

The system comes with a default administrative account for use with the IBM StoredIQ Administrator user interface. Use this account for the initial setup of IBM StoredIQ.

The system administrator must ensure that these tasks are completed before you can use IBM StoredIQ Administrator:

- Install and configure IBM StoredIQ.
 - Ensure that the data servers are up and running.
1. Open the IBM StoredIQ user interface from a browser.
Ask your system administrator for the URL.
 2. In the login window, enter your user name or email address and password.
To log in for the first time after the deployment, use the default administrative account: enter `superadmin` in the email address field and `admin` in the password field. For security reasons, change the password for the `superadmin` account as soon as possible, preferably the first time you log in. Also, create extra administrators for routine administration so that their actions can be audited.
 3. Click **Log in** to open IBM StoredIQ Administrator.
The `superadmin` account has access to all IBM StoredIQ applications on the application stack. To switch applications, click the triangle arrow icon and select the application that you want to access from the list of the available applications.

Data Servers and Volumes

The **Data Servers and Volumes** page of IBM StoredIQ Administrator contains the Data Dashboard, where you see all data currently under management, **Data Servers** and **Volumes** sub-navigation. When a data server is selected, the System Status and System Activity panes provide detailed information about the selected data server.

Data Dashboard

The description of the Data Dashboard fields is listed in the following table.

<i>Table 1. Data Dashboard fields and descriptions</i>	
Data Dashboard	Description
Total Data Objects	This field lists the total number of data objects under management.
Total Data Size	This field lists the total size of data objects under management.
Number of Data Servers	This field lists the total number of data servers under management.
Number of Volumes	This field lists the total number of volumes under management.

Data Servers

The **Data Servers** pane lists detailed information about a data server.

<i>Table 2. The Data Servers pane column names and descriptions</i>	
All Data Servers Column Name	Description
Data server name	Lists the name of the data server.
Status	Lists the data server's status, which are Healthy, Vulnerable, Critical, Unavailable, Under maintenance, Rebooting, or Restarting.
IP address	Lists the IP address of the data server.
Data objects	Lists the number of data objects that are found on the data server.
Total data object Size	Lists the size of the data objects that are found on the data server.

The **System Status** and **System Activity** panes provide performance information on a selected data server. These panels appear only once a data server is selected. The ability to check the status of all deployed data servers is critical to understanding the system's performance metrics.

<i>Table 3. System Status and System Activity panes field names and descriptions</i>	
System Status Field Name	Description
Status	Lists the data server's status, namely Healthy, Vulnerable, Critical, Unavailable, Under maintenance, Rebooting, or Restarting.

Table 3. **System Status** and **System Activity** panes field names and descriptions (continued)

System Status Field Name	Description
Status message	Describes the data server's status. Status messages include these messages: <ul style="list-style-type: none"> • System and services running. • System is running but some services are still coming up. • Some query processes are not running. • There was a problem with the RAID controller. • Some services are in error. • Cannot connect to system and error.
IP address	Lists the IP address of the data server.
Software version	Lists the installed software version.
DataServer type	Lists the installed index type: Classic (DataServer - Classic) and Distributed (DataServer - Distributed).
DB version	Lists the installed database version.
System time	Lists the system's time, date, and utilized time zone.
System Activity Field Name	Description
Free RAM memory	Denotes the amount of available memory (RAM) for the selected data server.
Free swap memory	Denotes the amount of available swap memory for the selected data server.
Load average	Lists the average process load for the selected data server.
Available space	Lists the available space on the selected data server.
Active DB connections	Lists the number of database connections currently open for the selected data server.
System uptime	Lists the total duration of time for which the data server ran.

Volumes

The pane provides the detailed information about the volumes.

Table 4. The **Volumes** pane column names and descriptions

Volumes Column Name	Description
Volume name	Lists the name of the volume.
Data server	Lists the name of the parent data server.
Type	Lists the volume's type, namely Primary , Retention , or Discovery Export . Primary indicates a primary volume. Retention indicates a retention volume. Discovery Export indicates a discovery export volume.
Source type	Lists the source or connection type for the volume.
Server name	Lists the name of the server where the volume is found.
Data objects	Lists the number of data objects that are found on the volume.

<i>Table 4. The Volumes pane column names and descriptions (continued)</i>	
Volumes Column Name	Description
Total data object size	Lists the size of the data objects in the volume.
Last harvested	Denotes the last time that the volume was harvested.
Publish to catalog	Denotes whether the volume is published to the governance catalog: No , Yes , and Yes (Inferred) , which indicates that the volume was excluded from synchronization but is still synchronized because it contributes to at least one infoaset that is fully published to the catalog. Note: This field is available only if your system is enabled for synchronization with the governance catalog.

Select the volume and click **View Details**, a bottom window opens with the number of data objects, total data object size, harvest status, last harvest date and time, the harvest type, information whether the volume is published to the governance catalog, and, if so, a link to the respective information asset in the catalog.

DataServer - Distributed

Data servers of the type DataServer - Distributed do not support the exact same features as data servers of the type DataServer - Classic.

Data servers of the type DataServer - Distributed support the following features:

- Full and incremental harvest
- Creating, deleting infosets, creating infosets from a data map, and applying set, node, and duplicate operations to infosets
- Uploading and deleting cartridges and auto-classification models
- Running auto-classification enhancements
- Metadata and full text search
- Supporting these actions: copy, move, delete, export, Step-up Snippet, Step-up Full-Text, and Step-up Analytics
- Covering these data sources: Box, CIFS, CMIS, Connections, Desktop, Documentum, Exchange, FileNet, HDFS, IBM Content Manager, IBM Domino, NFS, OneDrive, Sharepoint
- Supporting retention volumes of the NFS and CIFS server types
- Exporting and importing volume data
- Generating all reports
- Using the same StoredIQ gateway to manage both DataServer - Classic and DataServer - Distributed

If an infoaset contains volumes from both DataServer - Classic and DataServer - Distributed, then the infoaset is a mixed infoaset. Operations on mixed infosets that are not supported on DataServer - Distributed run only on DataServer - Classic and vice versa.

To deploy the virtual appliance that contains preinstalled Elasticsearch and populate the gateway and data server, see the topic about deploying the virtual appliances in the IBM StoredIQ deployment documentation.

Adding a primary volume

A primary volume serves as a primary data source in IBM StoredIQ. You must have at least one primary volume within your configuration.

1. Click **Data Servers** > **a data server** > **Add Volume**.
2. In the Add Volume dialog box, complete all fields as required for the data source.
 - [Box primary volumes](#)
 - [CIFS/SMB, SMB2, or SMB3 primary volumes](#)
 - [Connections primary volumes](#)
 - [CMIS primary volumes](#)
 - [Documentum primary volumes](#)
 - [Exchange primary volumes](#)
 - [FileNet primary volumes](#)
 - [HDFS primary volumes](#)
 - [IBM Content Manager](#)
 - [Livelink primary volumes](#)
 - [NFS primary volumes](#)
 - [NewsGator primary volumes](#)
 - [OneDrive primary volumes](#)
 - [SharePoint primary volumes](#)

Except for Box and OneDrive volumes, volumes can also be added in IBM StoredIQ Data Server. However, the set of available configuration options slightly varies.

Chatter, Domino®, and Jive volumes can be added in IBM StoredIQ Data Server only.

Note: When you name a volume, avoid using commas (.). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.

<i>Table 5. Box primary volumes.</i>		
Prerequisites: For Box volume prerequisites and configuration information, see Box volume configuration notes .		
Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select Box .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	The server name <code>api.box.com</code> is automatically set and cannot be changed.	

Table 5. Box primary volumes.

Prerequisites: For Box volume prerequisites and configuration information, see [Box volume configuration notes](#).

(continued)

Field	Value	Notes
Authenticate with Box	Before a Box volume can be added, the user must be authenticated. Click the Authenticate with Box link, sign in to the Box account, and select Grant access .	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common, user-defined name of this volume.	
Include Users	Select this option to scope the volume.	Regular expressions are supported.
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 6. CIFS/SMB, SMB2, or SMB3 primary volumes

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select CIFS .	SMB, SMB2, and SMB3 are supported. Depending on the setup of your SMB server, some additional SMB configuration might be required on the IBM StoredIQ data server. For details, see the information about configuring SMB properties in the Data Server administration guide. If you want to preserve ownership of objects in Copy or Move actions between CIFS volumes, you can add an admin knob as described in the respective instructions in the Data Server administration guide.

Table 6. CIFS/SMB, SMB2, or SMB3 primary volumes (continued)		
Field	Value	Notes
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting. If you create a volume for use with Distributed File System (DFS) services, provide the following information: For a domain-based namespace, specify the fully qualified domain name (FQDN) of the server. For a standalone namespace, specify the hostname of the namespace server.	For using DFS services, the <code>jcifs.smb.client.dfs.disabled</code> SMB property in the <code>jcifs.properties</code> file must be set to false. For details, see the information about configuring this SMB property in the Data Server administration guide.
Username	Enter the user name that is used to connect to and mount the volume. If you create a volume for use with Distributed File System (DFS) services, enter the fully qualified domain name of the server and the user name for connecting and mounting the volume in the format <code>FQDN \user</code> .	The user must be in the backup operator group on the Windows Share server that shows the shares on IBM® StoredIQ® and also needs to have full control share-level permissions.
Password	Enter the password that is used to connect to and mount the volume.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Share	Enter the share name of this volume. If you create a volume for use with Distributed File System (DFS) services, enter the DFS namespace.	Data from file or directory symbolic links in a share cannot be harvested.

<i>Table 6. CIFS/SMB, SMB2, or SMB3 primary volumes (continued)</i>		
Field	Value	Notes
Initial Directory	<p>Optionally, enter the name of the initial directory from which the harvest must begin.</p> <p>If you create a volume for use with Distributed File System (DFS) services, enter the folder target.</p>	If you create a volume for use with Distributed File System (DFS) services and want to include all namespace folders, do not specify an initial directory.
Indexing Options	<p>Select the checkbox for the indexing options that you want to include:</p> <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

<i>Table 7. Connections primary volumes.</i>		
Prerequisites: For Connections volume prerequisites and configuration information, see the topic about the configuration of IBM Connections in the IBM StoredIQ Data Server administration information.		
Field	Value	Notes
Volume Type	Select Primary in the Volume type list.	
Source Type	Select Connections in the Source type list.	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	
User name	Enter the user name of the account that is set up with admin and search-admin privileges on the Connections server.	
Password	Enter the password of the account that is set up with admin and search-admin privileges on the Connections server.	

Table 7. Connections primary volumes.

Prerequisites: For Connections volume prerequisites and configuration information, see the topic about the configuration of IBM Connections in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter a name for the volume.	
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Class name	Enter <code>deepfile.fs.template.impl.ibmconnections.ibmconnectionsconn.IBMConnections</code>	Required
Repository	Enter <code>deepfile.fs.template.impl.ibmconnections.ibmconnectionsconn</code>	Required
Option string	Enter more option parameters.	
Indexing options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 8. CMIS primary volumes

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select CMIS .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	
Port	Enter the port number.	

<i>Table 8. CMIS primary volumes (continued)</i>		
Field	Value	Notes
Username	Enter the user name that is used to connect to and mount the volume.	
Password	Enter the password that is used to connect to and mount the volume.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Use SSL	Select the Use SSL checkbox.	
Service	Enter the service name.	
Repository	Enter the name of the repository.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

<i>Table 9. Documentum primary volumes.</i>		
<p>Prerequisites: Before you can add Documentum volumes, you must add the Documentum server. For more information, see the topic about adding a Documentum server as a data source in the IBM StoredIQ Data Server administration information.</p>		
Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select Documentum .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Doc base	Enter the name that was entered on the data server from the doc broker settings.	
Username	Enter the user name that is used to connect to and mount the volume.	

Table 9. Documentum primary volumes.

Prerequisites: Before you can add Documentum volumes, you must add the Documentum server. For more information, see the topic about adding a Documentum server as a data source in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Password	Enter the password that is used to connect to and mount the volume.	
Assign To Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Harvest all document versions	If you need to harvest all document versions, select the checkbox.	Important: If you do not select this option for the initial harvest, changing the setting later does not have an effect when the volume is reharvested. As a workaround, create a new volume and ensure that the Harvest all document versions option is set before you start harvesting.
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects. • Include content tagging and full-text index. 	These options are not selected by default.

Table 10. Exchange primary volumes.

Prerequisites: For Exchange volume prerequisites and configuration information in general, see the topic about the configuration of Exchange servers in the IBM StoredIQ Data Server administration information.

Exchange Online volumes require some additional prerequisite configuration. For more information, see the topic about registering IBM StoredIQ with Microsoft in the IBM StoredIQ Data Server administration information.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select Exchange .	

Table 10. Exchange primary volumes.

Prerequisites: For Exchange volume prerequisites and configuration information in general, see the topic about the configuration of Exchange servers in the IBM StoredIQ Data Server administration information.

Exchange Online volumes require some additional prerequisite configuration. For more information, see the topic about registering IBM StoredIQ with Microsoft in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	If Exchange Online is selected as the source type, the server name is automatically entered.
Username	Enter the user name that is used to connect to and mount the volume.	This field is not available if Exchange Online is selected as the source type.
Password	Enter the password that is used to connect to and mount the volume.	This field is not available if Exchange Online is selected as the source type.
Impersonation Account	Enter the user account to use for connecting to Exchange Online. This account must be authorized to impersonate the members of the specified impersonation scope.	This field is available only if Exchange Online is selected as the source type. For more information, see the topic about registering IBM StoredIQ with Microsoft in the IBM StoredIQ Data Server administration information.
Client ID	Enter the application (client) ID under which IBM StoredIQ is registered with Microsoft.	This field is available only if Exchange Online is selected as the source type. For more information, see the topic about registering IBM StoredIQ with Microsoft in the IBM StoredIQ Data Server administration information.
Client Secret	Enter the client secret that is associated with the client ID. The values make up the credentials for access to a Microsoft Exchange Online data source.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	

Table 10. Exchange primary volumes.

Prerequisites: For Exchange volume prerequisites and configuration information in general, see the topic about the configuration of Exchange servers in the IBM StoredIQ Data Server administration information.

Exchange Online volumes require some additional prerequisite configuration. For more information, see the topic about registering IBM StoredIQ with Microsoft in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Volume Name	Enter the common name of this volume.	
Server Version	Select the version of Microsoft Exchange, choosing from 2000/2003, 2007, 2010/2013/2016 , and Online .	
Mailbox Server	Enter the names of the mailbox servers, which are separated by commas.	If Exchange Online is selected as the Server Version , this option is not available.
Active Directory Server	Enter the name of the Active Directory server.	If Exchange Online is selected as the Server Version , this option is not available.
Use SSL	To use secure socket layer, select the Use SSL checkbox.	If Exchange Online is selected as the Server Version , this option is automatically selected and cannot be edited.
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Virtual Root	The name defaults to the correct endpoint for the selected Exchange version.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 11. FileNet® primary volumes.

Prerequisites: For FileNet volume prerequisites and configuration information, see the topic about configuring the FileNet domain in the IBM StoredIQ Data Server administration information.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select FileNet .	

Table 11. FileNet® primary volumes.

Prerequisites: For FileNet volume prerequisites and configuration information, see the topic about configuring the FileNet domain in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	
Port	Enter the port number.	
Username	Enter the user name that is used to connect to and mount the volume.	
Password	Enter the password that is used to connect to and mount the volume.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Object Store	Enter the object store.	
Connection Type	Select either HTTP or HTTPS .	
Path	Enter the appropriate directory path.	
Stanza	Enter the appropriate stanza.	
Scope	Optionally, enter the appropriate SQL where clause.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 12. HDFS primary volumes

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select HDFS .	

<i>Table 12. HDFS primary volumes (continued)</i>		
Field	Value	Notes
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified host name of the server or the IP address from which the volume is available for mounting.	Either NameNode service or Knox Gateway service is assumed to be running on this server.
Port	Enter the port number.	For NameNode service, use port 50070 and port 8443 for Knox Gateway service.
Username	Enter the user name that is used to connect to and mount the volume.	
Password	Enter the password that is used to connect to and mount the volume.	Authentication to HDFS is not supported for NameNode connectivity (port 50070). If your HDFS server requires a password, use Knox Gateway connectivity.
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Use SSL	To use SSL, select the checkbox.	See Option String for more certificate options.
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Repository	Enter the name of the repository.	
Option String	VerifiCertificate=True <ul style="list-style-type: none"> • This option is supported. • This option is optional. knox_prefix=/gateway/default <ul style="list-style-type: none"> • This option is supported. • This option must be used for Knox Gateway connectivity. 	This VerifiCertificate option is used to indicate that the validity of the HDFS server's SSL certificate is verified when SSL is used. Values are True, False, or default value. If no value is specified, value is False. To validate the certificate on the HDFS server, the user needs to specify this option and set the value to True.

Table 12. HDFS primary volumes (continued)

Field	Value	Notes
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 13. IBM Content Manager primary volumes.

Prerequisites: For IBM Content Manager configuration information, see the topic about IBM Content Manager attributes in the IBM StoredIQ Data Server administration information.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select IBM Content Manager .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified host name of the library server database.	
Port	Enter the port that is used to access the library server database.	
Username	Enter the user name that is used to connect to and mount the volume.	
Password	Enter the password that is used to connect to and mount the volume.	
Connection String	Optional: Enter connection-string parameters.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Repository	Enter the name of the library server database.	

Table 13. IBM Content Manager primary volumes.

Prerequisites: For IBM Content Manager configuration information, see the topic about IBM Content Manager attributes in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Server Type	Select the type of server that is associated with the volume. Options include DB2 and Oracle . By default, DB2 is selected.	
Schema	Enter the schema for this library server database.	
Remote Database	Enter the name of the remote database.	
Harvest Itemtype	Enter the names of the item types to be harvested, separated by commas.	Harvest type is required to harvest the CM8 volume.
Copy to Itemtype	Either enter <code>SiqDocument</code> as the item type or leave the field empty. If you do not specify the item type, the volume cannot be used for copy-to actions.	For more information, see the topic about IBM Content Manager attributes in the IBM StoredIQ Data Server administration information.
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 14. Livelink primary volumes.

Prerequisites: A copy of the `lapi.jar` file from a Livelink API installation must be available in the `/usr/local/IBM/ICI/vendor` directory on each data server in your deployment. Usually, you can find this file on the Livelink server in the `C:\OPENTEXT\application\WEB-INF\lib` directory. However, the path might be different in your Livelink installation.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select Livelink .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	

Table 14. Livelink primary volumes.

Prerequisites: A copy of the `lapi.jar` file from a Livelink API installation must be available in the `/usr/local/IBM/ICI/vendor` directory on each data server in your deployment. Usually, you can find this file on the Livelink server in the `C:\OPENTEXT\application\WEB-INF\lib` directory. However, the path might be different in your Livelink installation.

(continued)

Field	Value	Notes
Port	Enter the port number.	
Username	Enter the user name that is used to connect to and mount the volume.	
Password	Enter the password that is used to connect to and mount the volume.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Database	Enter the database name.	
Search Slice	Enter the search slice.	
Initial Directory	Optionally, enter the search slice and name of the initial directory from which the harvest must begin.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 15. NFS primary volumes.

Prerequisites: Root access must be enabled on the NFS server that is connected to IBM StoredIQ.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select NFS .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.

Table 15. NFS primary volumes.

Prerequisites: Root access must be enabled on the NFS server that is connected to IBM StoredIQ.
(continued)

Field	Value	Notes
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Export	Enter the export name for this volume.	
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 16. NewsGator primary volumes.

Prerequisites: For NewsGator volume prerequisites and configuration information, see the topic about configuring NewsGator in the IBM StoredIQ Data Server administration information.

Field	Value	Notes
Volume Type	Select Primary .	
Source Type	Select NewsGator .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.	
Username	Enter the user name that is used to connect to and mount the volume.	

Table 16. NewsGator primary volumes.

Prerequisites: For NewsGator volume prerequisites and configuration information, see the topic about configuring NewsGator in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Password	Enter the password that is used to connect to and mount the volume.	
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	
Volume Name	Enter the common name of this volume.	
Initial Directory	Optionally, enter the name of the initial directory from which the harvest must begin.	
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.

Table 17. OneDrive primary volumes.

Prerequisites: For OneDrive volume prerequisites and configuration information, see [OneDrive volumes configuration notes](#).

Field	Value	Notes
Volume Type	Select Primary	
Source Type	Select OneDrive .	
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	For OneDrive volumes, enter the server name.	
Authenticate with OneDrive	Before a OneDrive volume can be added, the user must be authenticated. Click the Authenticate with OneDrive link, sign in with your Global Administrator account.	Required. You will also need to approve the requested permissions. Select the Consent on behalf of your organization checkbox and click Accept .

Table 17. OneDrive primary volumes.

Prerequisites: For OneDrive volume prerequisites and configuration information, see [OneDrive volumes configuration notes](#).

(continued)

Field	Value	Notes
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	Required.
Volume Name	Enter the common, user-defined name of this volume.	
Initial Directory	<p>Select one of these options:</p> <ul style="list-style-type: none"> To harvest the data of all sites including subsites and all private files on these sites, leave the field empty. <p>This is the default for new volumes.</p> <ul style="list-style-type: none"> To harvest the entire data of a specific site including its subsites, specify the site name. To harvest the entire data of a specific subsite, specify the name of the subsite in the format <i>site/subsite</i>. To harvest all private files of a specific user, specify this user's email address. <p>To include private folders in harvests of existing volumes, update the volume. Then, reharvest the volume to have the folders indexed.</p>	
Indexing Options	<p>Select the checkbox for the indexing options that you want to include:</p> <ul style="list-style-type: none"> Include metadata for contained objects Include content tagging and full-text index 	These options are not selected by default. Facets, personal drives, and notifications are not harvested.

Table 18. SharePoint primary volumes.

Prerequisites: For SharePoint volume prerequisites and configuration information, see the topic about the configuration of SharePoint in the IBM StoredIQ Data Server administration information.

Field	Value	Notes
Volume Type	Select Primary .	Required.

Table 18. SharePoint primary volumes.

Prerequisites: For SharePoint volume prerequisites and configuration information, see the topic about the configuration of SharePoint in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Source Type	Select SharePoint .	Required.
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.
Server	Enter the fully qualified domain name of the SharePoint server.	Required.
Username	Enter the name of a user with the required permissions for that site collections. Use the following syntax: <ul style="list-style-type: none"> SharePoint Online: <pre>userid@Microsoft_cloudname.com</pre> Other SharePoint versions: <pre>Active Directory Domain Name\username</pre> 	Required. Use a site collection administrator account. No volume can be added if the validation of the credentials fails, which can happen for the following reasons: <ul style="list-style-type: none"> The user does not exist or does not have the required permissions. The password is not correct. The HTTP status code is usually 401 Unauthorized. However, for SharePoint Online, the HTTP status code 400 Bad Request is returned for insufficient permissions.
Password	Enter the password for the user specified as Username .	Required.
Assign to Data Server	Select the IBM StoredIQ data server that you want to obtain and index the data.	Required.
Volume Name	Enter a meaningful name for this volume.	Required.
Server Version	Select the applicable SharePoint version: 2003, 2007, 2010, 2013, 2016, or Online .	Required.
Site URL	Enter the URL of the SharePoint site collection, for example: /portal/site	Required. Do not include the SharePoint server name in the URL, otherwise the URL cannot be located on the server and thus no volume is created.
Recurse into subsites	To check all sites and subsites of the site collection for data objects, select this option.	Optional.

Table 18. SharePoint primary volumes.

Prerequisites: For SharePoint volume prerequisites and configuration information, see the topic about the configuration of SharePoint in the IBM StoredIQ Data Server administration information.

(continued)

Field	Value	Notes
Use SSL	Select Use SSL only if SSL is enabled for this SharePoint server. For SharePoint Online, this option is automatically selected and cannot be edited.	Optional. If SSL is enabled on the SharePoint server and you do not select this option, no volume is created and the HTTP status code 301 Moved Permanently is returned. To fix the issue, select the option. If SSL is not enabled on the SharePoint server and you select this option, no volume is created and the socket error [Errno 111] Connection Refused is returned. To fix this issue, clear the Use SSL checkbox.
Include all versions	To harvest all versions of a document, select this option.	Optional.
Initial Directory	Enter the name of the subsite from which you want the harvest to start.	Optional.
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	Optional. Tip: Select Include metadata for contained objects to have metadata for objects in containers added to the metadata index. To avoid creating a full-text index for the entire volume, make sure the Include content tagging and full-text index checkbox is not selected. Create a full-text index for a subset of data later by running a Step-up Analytics action. For SharePoint Online, full-text indexing of OneNote notebook objects, that is, Notes®, is not supported currently. FSMD-based searches for these files are supported.

3. Click **Save** to save your configurations and add the volume.
4. Click **View Volumes**. Notice that the added volume appears therein, listed as a primary volume. To harvest this newly added volume, select that volume and then click **Harvest**.

Box volume configuration notes

Learn about the prerequisites and special considerations for Box volumes. Box volumes can be added only from IBM StoredIQ Administrator Administrator.

To complete the prerequisite configuration steps on the IBM StoredIQ application stack, you must be logged in as `siqadmin` user.

The prerequisites are as follows:

- Every application stack requires a dedicated Client ID. Each Client ID has its own Redirect URL that points to the IBM StoredIQ application stack. To create the application, log in to `developers.box.com`. Click **My Apps > Create Application** to get `client_id`, `client_secret`, and `redirect_uri`
- In the Box application, set `redirect_uri` to `https://x.x.x.x/proxy/enamel/1.0/oauthtokengenerator/box`, where `x.x.x.x` is the hostname/IP address of the application stack with which the user logs in.
- In the Box application, select the following **Scopes** options:
 - **Read and write all files and folders**
 - **Manage an enterprise**
 - **Manage an enterprise's managed users**
 - **Manage an enterprise's groups**
 - **Manage an enterprise's properties**
 - **Manage an enterprise's retention policies**
- Before Box volumes are added, the Client ID, Client Secret, and Redirect URI must be stored in an application stack `oauth.conf` file. You can edit this file by using the **`edit_etc_siq_oauth_conf`** command. After the `oauth.config` file is modified, restart the `uwsgi` service by running this command: **`systemctl restart appstack-uwsgi.service`**
- Each Client ID must have the **As-User** capability that is enabled to allow harvests and data discovery.
Note: Customers must contact their Box support representative to enable this permission. They also need their Client ID.
- The application stack must be configured for SSL access, and it must have access to the Box authentication endpoint on `api.box.com`. If the application stack was deployed without SSL enabled, run **`certcfg`** and follow the instructions on the Certificate Configuration screen.
- A data server that manages a Box volume must have access to `api.box.com`.
- The **Application Settings** for the Box Enterprise Account must not disable unpublished applications from accessing the account. Clear the **Unpublished Applications** checkbox when you configure **Application Settings**.

When you copy to Box, by default the owner of a source content is mapped to the Box user only if the corresponding user is administered by the Box administration account. This way the contents can be copied to the mapped user account in Box. Additionally, the folder hierarchy of the source volume can be reproduced, but it is rooted in the home folder of the user.

When the user configures a copy action with Box as the target set from IBM StoredIQ Administrator, the user can optionally enter a destination directory. The user also can harvest the copied files by selecting or deselecting **Do not auto-harvest destination volume after copy**.

Two other options are available for the copy action, which has Box as target data source.

Mapping

Map permissions from source to destination maps access permissions from the source to Box folder collaborators. Mapping preserves owners from source to destination volume.

Preserve version chains on destination

Creates version chains in Box during the copy action.

Both options are not enabled by default and they can be modified based on your needs. If you clear **Preserve version chains on destination**, version chains are not created in Box, but all versions are copied as separate files.

If an infoset that is copied to Box contains an archive file along with members, then IBM StoredIQ copies only the archive file. The members of the archive file are skipped and audited in Policy Audits.

See the following retained Box metadata attributes to copy to Box volumes.

File name

Same as source.

Description

Same as source.

Owner

Same as source based on mapping heuristics that apply only to CIFS and SharePoint. Box administrator's user name in all other cases.

For CIFS, if single sign-on (SSO) that is based on Active Directory is configured and the IBM StoredIQ data server can locate the Active Directory Domain Controller by using DNS, IBM StoredIQ uses the E-mail attribute from the Active Directory user profile of the owner of a source data object and tries to match it with the email address of a user in Box. If the owner of a data object does not have an email address in Active Directory Server, or if owners of data objects that are not defined in Active Directory Server cannot be mapped, owner mapping fails.

If the IBM StoredIQ data server cannot locate the Active Directory Domain Controller, IBM StoredIQ uses the owner name itself and tries to map it to a Box User ID. Owner names in CIFS are typically in the form of `domain\username`, whereas email addresses in Box are in the form of `localpart@domainname`. In this particular case, IBM StoredIQ maps only `username` to `localpart` of an email address in Box in a case insensitive manner.

For SharePoint, the mail address property of the source content owner must be matched with the Box User ID, where Box User ID must be managed by the Box administrator account and Box User ID is the mail address of the user. If the mail address is not available, the display name of the source content owner must be matched exactly with the display name of the Box user. If IBM StoredIQ cannot find a display name for the SharePoint user, then the `username` part of a SharePoint login name is mapped to the `localpart` of an email address in Box. For example, `username` in `domain\username` is mapped to `localpart` in `localpart@domainname`.

No user mapping takes place if the conditions that are described are not met. Then, all content is copied to the administrator account of Box.

Size

Same as source.

Created

Same as source.

Modified

Same as source.

Copy From Box is supported. IBM StoredIQ supports copying files from Box to CIFS, NFS, FileNet, and Box as target sets.

- Box Notes and Bookmarks are skipped during a Copy From Box. Skipped information is recorded in Policy Audits.
- If multiple versions exist, all versions are copied to the target. If the target data source does not support versioning, then the copied files have version numbers that are appended to the file names.
- Box metadata, except for **Created At** and **Modified At**, is not copied when files are copied from Box.
- Source file owner name is not mapped in the destination volume. Instead, all of the copied files have an `owner_name`: user name of the target volume.

IBM StoredIQ supports copying to Box from CIFS, NFS, and SharePoint source volumes. However, only copying documents and files to Box is supported. Copying social or collaboration content types such as Wiki pages, blog posts are not supported. IBM StoredIQ also supports Discovery Exports from Box. Box Notes and Bookmarks are exported as MHTML files. An export for a Box Note carries the textual content of the Box Note. An export of the Box Bookmark has no textual content but only metadata.

Each StoredIQ data server contains a CSV mapping file, which includes a set of CIFS and SharePoint standard properties that are mapped to the custom Box properties. The mapping file is on the data server at `/deepfs/data/mapping_file/box_mapping.csv`. The mapping file can be edited as required but for the changes to take effect, services need to be restarted on the StoredIQ data server.

StoredIQ can index a list of collaborators who work with a document in Box. If an ancestor folder of a document is shared for collaboration, then each collaborator is indexed by StoredIQ with information about the User ID, email address, name, and role of the user. Incremental harvests in StoredIQ currently do not pick up changes to collaborators for a document. A full harvest might be required to get the index up-to-date concerning collaborators for documents.

Box notes cannot be viewed through the Data Object Viewer in IBM StoredIQ Data Workbench.

When an external user who is not native to the organization collaborates a folder with users who are managed by the administrator, the folder is not indexed if the Box volume was added with the credentials of the administrator. This kind of collaborated folder can be harvested only if the Box volume is added with the credentials of the managed users who were invited to collaborate on the folder by the external user.

When files or folders are deleted in Box, they are moved to the Trash folder in Box. Currently, contents in Trash are not indexed by StoredIQ.

Mapping of access permissions to Box folder collaborators

When you copy data into Box, besides the ownership of individual files, access permissions might also need to be retained based on the access control settings on the source volumes. The **Mapping: Map permissions from source to destination** option in the **Copy Action** is designed to support this use-case. IBM StoredIQ applies its proprietary mapping heuristic for mapping access control from CIFS and SharePoint volumes to Box folder collaborators.

Box supports only collaboration at the folder-level. Collaborators cannot be defined for individual files. The IBM StoredIQ proprietary mapping heuristic for mapping access control from CIFS and SharePoint volumes to Box folder collaborators works as follows.

Note: Currently, Box ACL is supported on the SharePoint 2013 server.

Individual file permissions within a source folder are mapped to Box collaborators that can be applied to the corresponding target folder. The folder collaborators on the target folder are an aggregate of the permissions for all files within the source folder. The permissions are aggregated in such a way that no user accidentally has permissions to a file on Box that they did not have access to on the source. However, this aggregation might, in certain cases, force some users to lose access to documents on Box that they were able to access on the source. For example, if User A had access to one file in the source folder but not to another file within the same folder, the aggregation forces User A to lose access to both files in Box. It happens because Box collaborators can be applied only at the folder-level.

Permissions are mapped from CIFS and SharePoint to Box Collaboration Roles by using the following mapping tables:

CIFS	BOX
Full Control	Co-owner, Owner
Modify	Editor
Read & Execute	Viewer
List Folder Contents	Viewer

<i>Table 19. Mapping permission from CIFS to Box Collaboration Roles (continued)</i>	
CIFS	BOX
Read	Viewer
Write	Uploader

<i>Table 20. Mapping permission from SharePoint to Box Collaboration Roles</i>	
SharePoint	BOX
Full Control, Design	Co-owner, Owner
Edit, Contribute	Editor
Read	Viewer
View Only	Pre-Viewer

The **Mapping: Map permissions from source to destination** option works in tandem with the **Mapping: Preserve owners from source to destination** option:

- If only **Map permissions from source to destination** is selected, all files and folders are copied to the Box Administrator account and access control is mapped from the source volume to Box folder collaborators.
- If both **Map permissions from source to destination** and **Preserve owners from source to destination** are selected, then each file or folder is copied to Box user accounts that map to the source file owners and access control is mapped from the source volume to Box folder collaborators.
- If only **Preserve owners from source to destination** is selected, then each file or folder is copied to Box user accounts that map to the source file owners, but no folder collaborators are added on Box.
- If neither of these options are selected, the files and folders are all copied to the Box Administrator account and no folder collaborators are added on Box.

As an example, if a Public File Share is being copied to Box, **Map permissions from source to destination** can be selected while **Preserve owners from source to destination** can be deselected so that all the contents land in the Box Administrator account, while individual users still continue to have collaboration roles for this content. Similarly, if a particular User Share is being copied to Box, **Map permissions from source to destination** can be deselected while **Preserve owners from source to destination** can be selected so that all the contents land in the mapped Box user's account.

Before you run the actual copy action, it is a good practice to run a simulate action to preview the mappings and determine whether any access control changes are necessary before you run the copy action.

Note: Any files directly within a source file share (not contained within other folders), do not have collaboration enabled in Box since they do not have a parent folder to collaborate on. Setting a **Destination Directory** for the Copy Action ensures that these files have a parent folder and collaboration can be enabled in Box. Any files directly within a source file share (not contained within other folders), do not have collaboration enabled in Box since they do not have a parent folder to collaborate on. Setting a **Destination Directory** for the Copy Action ensures that these files have a parent folder and collaboration can be enabled in Box.

Registering IBM StoredIQ as a Microsoft service application for access to Exchange Online

For harvesting Exchange Online mailboxes, IBM StoredIQ must be registered as a service application with Microsoft.

This procedure requires administrator access to the Microsoft Azure portal.

Register IBM StoredIQ to integrate it with the Microsoft identity platform for secure login and authorization when accessing Microsoft Exchange Online volumes.

To set up OAuth authentication and obtain the authentication information for use in IBM StoredIQ, complete these steps:

1. Log in to <https://portal.azure.com>.

Make note of the login ID. This user ID is set up as the mailbox impersonator. You'll need this information for the Exchange Online volume configuration in IBM StoredIQ.

2. Select **Azure Active Directory** service.
3. Select **App registrations > New registration**.
4. On the **Register an application** page, enter IBM StoredIQ as the display name for the application.
5. Under **Supported account types**, select **Accounts in any organizational directory (Any Azure directory - Multitenant)**.
6. Do not specify a redirect URI.
7. Click **Register**.

This generates a unique application (client) ID that you can see on the **Overview** page. Make a note of this ID, you'll need it for the volume configuration in IBM StoredIQ.

8. Configure permissions for the application.
 - a) On the **Overview** page, select **API permissions** from the menu on the left and then click **Add a permission**.
 - b) On the **Request API permissions** page, look for **Exchange** in the **Supported legacy APIs** section and select it.
 - c) On the **Request API permissions** page for Exchange, select **Application permissions**.
 - d) Select the **full_access_as_app** checkbox and click **Add permissions**.
You return to the **API permissions** page, where you can now see the granted permissions listed.
 - e) Click **Grant admin consent for *tenant_name***, where *tenant_name* is your company's tenant name.
9. Add credentials to the application.
 - a) On the **Overview** page, select **Certificates & secrets** from the menu on the left.
 - b) Click **New client secret**.
 - c) Optionally, add a description for your client secret.
 - d) Set the client secret to never expire.
 - e) Click **Add**.

The table now shows the description, the expiration date, and the newly created client secret value.

Important: Be sure to copy this value for use in IBM StoredIQ. As soon as you leave this page, the value is obscured and no longer accessible. You'll need the client secret value for configuring Exchange Online volumes in IBM StoredIQ.

10. Save all changes.

IBM StoredIQ is now properly registered with Microsoft.

OneDrive volumes configuration notes

OneDrive for Business volumes can be added only from IBM StoredIQ Administrator and have several configuration prerequisites that must be met.

To complete the configuration steps on the IBM StoredIQ application stack, you must be logged in as siqadmin user.

Register IBM StoredIQ with Microsoft

To set up OAuth authentication and obtain the authentication information for use in IBM StoredIQ, complete these steps:

1. Log in to <https://portal.azure.com>.
2. Select **Azure Active Directory** service.
3. Select **App registrations > New registration**.
4. Enter IBM StoredIQ as the display name for the application.
5. Under **Redirect URIs**, select **Web** and set **Redirect URI** to `https://x.x.x.x/proxy/enamel/1.0/oauthtokengenerator/onedrive`, where `x.x.x.x` is the IP address of the IBM StoredIQ application stack.
Make a note of the redirect URI, you'll need it for the configuration in IBM StoredIQ.
6. Click **Register**. This generates an application ID (or client ID) that you can see on the **Overview** page. Make a note of this ID, you'll need it for the configuration in IBM StoredIQ.
7. On the **Overview** page, select **Authentication** from the menu on the left. Then, set the account type to **Accounts in any organizational directory (Any Azure directory - Multitenant)** and save your change.
8. Configure permissions for the application. On the **Overview** page, select **API permissions** from the menu on the left and then click **Grant admin consent for *tenant_name***, where *tenant_name* is your company's tenant name.
9. Add credentials to the application. On the **Overview** page, select **Certificates & secrets** from the menu on the left and add a client secret:
 - a. Click **New client secret**.
 - b. Optionally, add a description for your client secret.
 - c. Set the secret to never expire and click **Add**.
 - d. Make a note of the password (client secret), you'll need it for the configuration in IBM StoredIQ.
10. Save the changes.

Create the OneDrive configuration in IBM StoredIQ

Before OneDrive volumes are added, the client ID, client secret, and redirect URI must be stored in an application stack `oauth.conf` file. You can edit this file by using the **edit_etc_siq_oauth_conf** command. The OneDrive configuration must be a separate configuration under the Box configuration with the section header `[onedrive]`.

- Include the key for the OneDrive configuration:
`OAUTHLIB_RELAX_TOKEN_SCOPE=1`
- Set the properties `client_id`, `client_secret`, and `redirect_uri` to the values obtained from registering IBM StoredIQ with Microsoft.
- Set the scope:
`scope=offline_access Files.ReadWrite.All Group.ReadWrite.All Notes.ReadWrite.All`
- Set the authentication and token URLs:
`auth_url=https://login.microsoftonline.com/common/oauth2/v2.0/authorize`
`token_url=https://login.microsoftonline.com/common/oauth2/v2.0/token`

A sample `oauth.conf` file is as follows:

```
[onedrive]
client_id=d1106b42-ec8e-40b9-9832-d5c970aaa7bc
client_secret=bmSuRLeFzaDKemMrDNmdFS5
redirect_uri=https://198.51.100.0/proxy/enamel/1.0/oauthtokengenerator/onedrive
auth_url=https://login.microsoftonline.com/common/oauth2/v2.0/authorize
token_url=https://login.microsoftonline.com/common/oauth2/v2.0/token
scope=offline_access Files.ReadWrite.All Group.ReadWrite.All Notes.ReadWrite.All Sites.ReadWrite.All User.ReadWrite.All
OAUTHLIB_RELAX_TOKEN_SCOPE=1
```

After the `oauth.config` file is modified, restart all appstack services. Using an SSH tool, log in to the application stack and run **systemctl stop appstack.target** and then **systemctl start appstack.target**.

Further considerations

- The application stack must be configured for the SSL access; it must have access to the OneDrive authentication endpoint on `graph.microsoft.com`. If the application stack was deployed without SSL enabled, run **certcfg** and follow the instructions on the Certificate Configuration screen.
- A data server that manages a OneDrive volume must have access to `graph.microsoft.com`.
- IBM StoredIQ supports index, copying from, and exporting from OneDrive for Business to CIFS and NFS shares. However, OneNote is not supported in the OneDrive harvest.
- When you add the OneDrive volume in the IBM StoredIQ Administrator UI, the **Authentication with OneDrive** link is enabled only if the `oauth.conf` file contains a `[onedrive]` section with the proper settings. To complete the setup of a OneDrive volume, you must click the link and provide administrator credentials. Then, approve the requested permissions by selecting the **Consent on behalf of your organization** checkbox and clicking **Accept**.
- If you want to harvest users' private files, you must permit access to their OneDrive files. Complete the following steps:
 1. Log in to the Microsoft 365 admin center at <https://admin.microsoft.com/AdminPortal>.
 2. Go to **Users > Active Users**.
 3. Search for the user whose OneDrive files are to be harvested and double-click the user ID.
 4. Go to the **OneDrive** tab in the details pane.
 5. In the **Get access to files** section, click **Create link to files**.
 6. Click the **X** to close the details pane.

Complete steps “3” on page 32 through “6” on page 32 for each user whose files are to be included in the harvest.

Adding a retention volume

A retention volume stores data objects that are placed under retention, which means that the object is retained for a specified period.

1. Click **Data Servers** and select a data server, then click **Add Volume**.
2. In the Add Volume dialog box, complete these fields:

<i>Table 21. Retention volume options</i>			
Field	Value	Notes	Source type
Volume Type	Select Retention .		<ul style="list-style-type: none"> • CIFS • NFS
Source Type	Select the source or connection type.		<ul style="list-style-type: none"> • CIFS • NFS
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.	<ul style="list-style-type: none"> • CIFS • NFS

<i>Table 21. Retention volume options (continued)</i>			
Field	Value	Notes	Source type
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.		<ul style="list-style-type: none"> • CIFS • NFS
Username	Enter the user name that is used to connect to and mount the volume.		<ul style="list-style-type: none"> • CIFS
Password	Enter the password that is used to connect to and mount the volume.		<ul style="list-style-type: none"> • CIFS
Assign to Data Server	Select the data server to which you want to assign the volume.		<ul style="list-style-type: none"> • CIFS • NFS
Volume Name	Enter the common name of this volume. Avoid using commas (,) in the name. Although a comma is a valid character, it causes queries to fail when specified as part of the search text.		<ul style="list-style-type: none"> • CIFS • NFS
Share	Enter the share name for this volume.		<ul style="list-style-type: none"> • CIFS
Export	In the Export text box, enter the export name for this volume.		<ul style="list-style-type: none"> • NFS
Indexing Options	Select the checkbox for the indexing options that you want to include: <ul style="list-style-type: none"> • Include metadata for contained objects • Include content tagging and full-text index 	These options are not selected by default.	<ul style="list-style-type: none"> • CIFS • NFS

3. Click **Save** to save your configurations and add the volume. Retention volumes do not have any data to harvest until a copy to retention runs.
4. Click **View Volumes**. The added volume appears, listed as a retention volume.

Adding a discovery export volume

A discovery export volume contains data that is produced from a discovery export action. That data is kept so that it can be exported as a load file and uploaded into a legal review tool.

1. Click **Data Servers & Volumes > All Data Servers**, and then click **Add Volume**.
2. In the Add Volume dialog box, complete these fields:

<i>Table 22. Discovery export volume options</i>			
Field	Value	Notes	Source type
Volume Type	Select Discovery Export .		<ul style="list-style-type: none"> • CIFS • NFS
Source Type	Select the source or connection type.		<ul style="list-style-type: none"> • CIFS • NFS

Table 22. Discovery export volume options (continued)

Field	Value	Notes	Source type
Unified Governance	If you want to exclude the volume from synchronization with the governance catalog, clear the Publish to catalog checkbox.	This option is available only if synchronization with the governance catalog is enabled.	<ul style="list-style-type: none"> • CIFS • NFS
Server	Enter the fully qualified domain name of the server from which the volume is available for mounting.		<ul style="list-style-type: none"> • CIFS • NFS
Username	Enter the user name that is used to connect to and mount the volume.		<ul style="list-style-type: none"> • CIFS
Password	Enter the password that is used to connect to and mount the volume.		<ul style="list-style-type: none"> • CIFS
Assign to Data Server	Select the data server to which you want to assign the volume.		<ul style="list-style-type: none"> • CIFS • NFS
Volume Name	Enter the name of this volume. Avoid using commas (,) in the name. Although a comma is a valid character, it causes queries to fail when specified as part of the search text.		<ul style="list-style-type: none"> • CIFS • NFS
Share	Enter the share name for this volume.		<ul style="list-style-type: none"> • CIFS
Export	Enter the export name for this volume.		<ul style="list-style-type: none"> • NFS

3. Click **Save** to save your configurations and add the volume.
4. Click the **All Volumes** tab. The added volume appears, listed as a discovery export volume.

Editing a volume

You can edit the existing volumes from the **Volumes** pane.

To edit a volume, follow these steps.

1. Select a volume from the existing volume list and click **Edit Volume**.
The **Edit Volume** window appears.
2. Make changes as needed.

For a volume that is synchronized because data objects from this data source are part of a published infoset (marked **Yes (Inferred)** in the list of volumes), the **Publish to catalog** checkbox is not selected. If you decide to have the volume fully synchronized to the governance catalog, select the option.

Restriction: For some source types, editing from IBM StoredIQ Administrator is limited or not possible at all:

- For volumes of source type Domino, you can change only the setting of the **Unified Governance** option. All other settings must be modified from IBM StoredIQ Data Server.
- Chatter and Jive volumes cannot be edited from IBM StoredIQ Administrator at all. They can be edited only from IBM StoredIQ Data Server. If you want to exclude such a volume from

synchronization to the governance catalog, contact your IBM Services representative or IBM Support.

3. Click **Save**.

Tip: In some cases, you might need to specify a password to be able to save your changes.

Depending on the changes you made, you might need to reharvest the volume to see the changes reflected.

Viewing details of multiple volumes

Detailed, combined information for multiple volumes can be viewed.

1. Click **View Volumes**, and then press and hold the **CTRL** key or **SHIFT** key to select the volumes for which you want to view combined detailed information.

The Combined Volume Details page appears, showing the combined details of the selected volumes, including the number of data objects, the size of the total number of data objects, the number of data servers, and the number of volumes.

2. Add or remove volumes from the Combined Volume Details view using the **CTRL** key.

The volume is added or removed, and the Combined Volume Details view is updated automatically.

Restarting services on a data server

You can restart services on a data server from the data server's details view in IBM StoredIQ Administrator.

1. Click **Data > Data Servers**, then select a data server.

The **Details: data server name** page appears and the **Restart Services** and **Reboot Data Server** buttons become active.

2. Click **Restart Services**.

The **Confirm Restart Services** window appears, where you must confirm that you want to restart services on the selected data server.

3. Click **Restart Services**, and services restart.

When the services are restarting, the data server's status is changed to **Restarting**, and the status message is one of the following:

- System and services running.
- System is running but some services are still coming up.
- Some query processes are not running.
- There was a problem with the RAID controller.
- Some services are in error.
- Cannot connect to system and error.

All other details appear as **N/A** until the services restart, at which point the data server's details are once again depicted.

Rebooting a data server

You can reboot a data server from the data server's details view in IBM StoredIQ Administrator.

1. In IBM StoredIQ Administrator, click **Data > Data Servers** and select a data server.

The **Restart Services** and **Reboot Data Server** buttons become active.

2. Click **Reboot Data Server**.

The **Confirm Data Server Reboot** window appears, where you must confirm that you want to reboot the selected data server.

3. Click **Reboot Data Server**, and the data server reboots.

When the data server is rebooting, its status is changed to **Maintenance Mode**, and the status message is **Rebooting-system**. All other details appear as **N/A** until the reboot completes, at which point the data server's details are once again depicted.

Removing a volume

Volumes can be deleted from the list of available data sources. However, volume deletion can cause data loss in existing infosets or indices that reference the deleted volumes.

1. Select **Volumes** on the side navigation bar or **View Volumes** on the Data Servers pane, and then select the volume you want to delete.
2. Click **Remove Volume**.

The volume is deleted, removing it from the list of available volumes.

If the volume was published to the governance catalog, the respective asset will be deleted from the catalog on the next synchronization run.

Harvests

The **Harvests** page of IBM StoredIQ Administrator displays the complete list of harvests, where you see all harvests currently under management.

When you click **Harvests** from the side navigation bar, the menu bar expands to show the **Scheduled**, **Current**, and **Completed** options. Select either **Scheduled**, **Current**, or **Completed**. Depending on the option that is selected, only those harvests in that category are shown Scheduled in the List of Harvests table.

The description of the different harvest fields is listed in the following table.

List of Harvests	Description
Name	Lists the name of the harvest instance.
Type	Lists the type of harvest, either Full or Incremental .
Start time	Lists the date and time at which the harvest was initiated.
End time	Lists the date and time at which the harvest was ended.
Total time	Lists the time that the harvest takes.
Owner	Lists the owner of the harvest.

Harvesting a volume

By clicking **Harvest**, an administrator can initiate a full or incremental harvest of a primary volume without having to utilize IBM StoredIQ Data Server.

1. Click **Data** on the side navigation bar.
The menu expands.
2. Click **Volumes**. Or, from the Data Servers pane, click **View Volumes**.
3. Select a volume and click **Harvest**.
The **Harvest Volume** dialog box appears.
4. In the **Harvest Name** text box, enter a unique name for this harvest.
When you name the harvest, avoid using commas (.). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.
5. In the **Schedule Harvest** area, select either of the following options.
 - Select **Immediate** to harvest the volume immediately.
 - Click **Schedule** to schedule a date and time for the harvest to occur. Enter the date (YYYY-MM-DD) and time in the text boxes.
6. In the area of **Harvest Options**, select **Incremental** or **Full**.
 - An **Incremental** harvest only harvests any changes on the selected volume since its last harvest.
 - A **Full** harvest performs a full harvest of the selected volume, regardless of when it was last harvested.
7. Click **Save**.
The volume is harvested.

Searching for a specific harvest

When you have multiple harvests available to choose, you can search for a specific harvest instead of using the slider to move through available options.

1. Click **Harvests** on the side navigation bar, select either **Scheduled**, **Current**, or **Completed**.
Depending on the option that is selected, only harvests in the selected category are shown in the table.
2. In the **Enter key term(s)** text box, enter the name of the harvest for which you want to search, and then click **Search**. You can enter either the full harvest name or a portion of it.
To remove the search term, click the **X** to the right of the **Enter key term(s)....** text box.
Harvests that match the entered search term are returned.
3. Select a harvest from the returned list of harvests.

Stopping a harvest

Harvests that are not yet completed can be stopped.

1. In the **Scheduled Harvests** list, select **Current Harvests**. The list of harvests that are currently running opens.
2. Select the harvest that you want to stop and click **Stop Harvest**.
The **Confirm Stopping Harvest** window appears.
3. Click **Stop Harvest** to confirm that the harvest needs to be stopped and not completed.

Editing a harvest

Existing harvests can be modified as needed.

1. In the **Scheduled harvests** pane, select the harvest that you want to edit, and then click **Edit Harvest**.
The Edit Harvest dialog box appears.
Note: Only scheduled harvests can be edited.
2. In the Edit Harvest dialog box, edit the harvest as needed.
3. Click **Save** to save your changes.

Deleting a harvest

Scheduled harvests can be deleted from the **List of Harvests** page.

1. In the **Scheduled harvests** pane, select the scheduled harvest that you want to delete, and then click **Delete Harvest**.
The **Delete Harvest** window appears.
2. Click **Delete**.
The harvest is deleted from the list of scheduled harvests.

Reviewing harvest information

Harvest information is available in the **Completed** harvests page. For each harvest that has exceptions, you can drill down to examine the errors.

Take the following information into account when you review harvest details:

- Only harvests that are initiated on the IBM StoredIQ Administrator server are reflected. Harvests that are requested from the data server do not show up in this list.

- Step-up Full-Text and Step-up Snippet action do not contribute to the harvest details shown here.
- Contrary to the data server, skipped directories contribute to the processing exceptions that account for the value in the **Data Objects Not Harvested** field in the harvest details.
- For harvests that were completed before IBM StoredIQ release 7.6.0.8, the **Skipped directories** and **Skipped - user configuration** categories display zero.

The **Harvest Details** view presents summary information about a harvest. The value for **Data Objects Not Harvested** indicates the number of exceptions that the selected harvest has. Any number greater than zero is a link that you can click it to drill down to the exceptions.

1. Select **Harvest > Completed** on the side navigation bar.

To monitor running harvests, use the Harvest Tracker tool. For more information, see the topic about monitoring harvests in the data server administration documentation.

2. From the list of completed harvests, select a harvest that you want to review.

3. In the harvest details, check the number next to **Data Objects Not Harvested**.

Zero indicates that this harvest has no exceptions. For any number greater than zero, click the number to proceed to the harvest exception page.

Note: If a harvest from the link expired or a volume was deleted, the exception page shows the categories and their respective counts, but no exceptions are listed. An appropriate message also appears.

4. Select from these exception categories to filter the list for review.

- **Binary text extracted, full processing complete**
- **Binary text extracted, partial processing complete**
- **Content Skipped - user configuration**
- **Content type known, but cannot extract content**
- **Content type known, but error processing content**
- **Content type known, partial processing complete**
- **Content type unknown, not processed**
- **Error gathering ACLs**
- **Error processing binary content**
- **Skipped - cannot access data object**
- **Skipped directories**
- **Skipped - user configuration**

5. To return to the **Completed** harvest page. either click **Harvest Management** or **Harvests > Completed**.

System infosets

System infosets can be created in IBM StoredIQ Administrator to allow users to have a different starting point than the All Data Objects infoset.

System infosets are the basis for user-created infosets. System infosets allow administrators to select which indexed volumes the infoset can draw upon, creating a shortcut for IBM StoredIQ Data Workbench users. Actions cannot be run on system infosets, but only on user-created infosets. Therefore, you must first create the system infoset here in IBM StoredIQ Administrator and then create a user infoset that can be acted upon in IBM StoredIQ Data Workbench.

Infosets can be public or private. The administrator determines users' access to system infosets at the time of creation. The All Data Objects and All System-Level Objects infosets, which are generated by the application, are set to admin by default and can be viewed only by an administrator. Also, the All Data Objects and All System-Level Objects system infosets cannot be deleted. Other system infoset can be deleted if it is not in use by any other infosets.

You can search the list of system infosets for a specific infoset by entering either the complete user name in the search field or by entering a partial string in this field to filter the list.

Note: System infosets must not be used to manage volumes or indexes. Additionally, system infosets must be created judiciously. If several system infosets through which a user must cull exists, the shortcut aspect is removed.

Infoset Management Column Name	Description
Name	Lists the name of the system infoset. This name must be unique against all infosets.
Total objects	Lists the total number of data objects that are contained within the system infoset.
Infoset size	Lists size of the system infoset.
In use by	Indicates the infosets in which this system infoset is used. If the system infoset is in use, a link is shown, indicating the number of infosets by which the system infoset is being used. Click the link to see which infosets use this system infoset.
Access	Indicates whether the system infoset is public or private. If the system infoset is public, it can be viewed by any user. If the system infoset is private, it can be viewed by only those users granted permission by the administrator. If users are granted permission, the number of users is listed with a hyperlink to the Access Members modal, which lists the name and role of users who have access.
Description	This column lists the optional description of the system infoset.

Creating a system infoset

System infosets can be created with selected volumes, creating a shortcut for IBM StoredIQ Data Workbench users.

1. Click **System Infosets**, and then click **Create Infoset**.
2. In the **Add System Infoset** window, complete these fields.
 - a) In the **Infoset Name** text box, enter the name of this system infoset.

When you name the infoset, avoid using commas (.). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.

- b) In the **Description of Infoset** text box, enter a brief description of this system infoset. This infoset might need to be used by several people, so an accurate description can be helpful to others.
- c) In the Access area, click either **Public** or **Private**.

The All Data Objects and All System-Level Objects infosets, which are generated by the application, can be viewed only by an administrator. System infosets that are created by an administrator default to public. Public system infosets are visible and available to all users, including users that are added at a future time. For public system infosets, the user list is disabled.

To limit access to the system infoset and all of the user infosets created from it, create a private system infoset. From the list of available users, select the ones that you want to have access to the private system infoset. If you do not select any users, only administrators have access to that system infoset because they have administrative rights to every infoset.

Important: Access cannot be edited. It must be determined at the time of infoset creation.

- d) Add volumes to this infoset. In the **Available Volumes** area, select the volumes that you want to add to this system infoset, and then click **Add**.

To remove volumes from the infoset, select the volume in **Selected Volumes** and then click **Remove**.

- 3. Click **Save**.

The infoset is now listed in the System Infosets pane.

Limiting access to data

You might want to restrict the number of users who have access to the data on a specific volume, for example, for volumes that contain sensitive data.

Access to the data on a volume is determined by the setup of the system infoset that you create and by the roles that those users have who have access to the system infoset.

To limit access to the data on a specific volume:

- 1. Determine which users you want to have access to the data.

User roles determine whether a user can view the data or just the list of data objects in an infoset:

User role	Access rights
Admin	This role grants access to all data on a volume. Users with this role automatically have access to every infoset.
Data User Insights User	This role grants access to all public system infosets, specifically assigned private system infosets, and the user infosets created from them, and grants permission to preview data objects.
Data User (No Preview)	This role grants access to all public system infosets, specifically assigned private system infosets, and the user infosets created from them, but does not grant permission to preview data objects.

- 2. Create a private system infoset for the volume.

From the list of available users, select the ones that you want to have access to the infoset. Only these users can see the infoset and create user infosets from it besides any administrators. Remember that you cannot edit the access after the infoset is created.

The access rights for any new user info set based on this system info set are inherited from the system info set.

Important: The right to view the content of data objects is tied to the user role. If the role includes the viewing permission, users with this role can view the data objects of all info sets to which they have access. You cannot limit the viewing permission at the info set level.

User management

A user account is required for logging in to and working with IBM StoredIQ.

The **Users** pane lists all user accounts that exist on the application stack and provides the following information:

- The user name.
- The user's given name and surname.
- The user's email address. System notifications are sent to this email address.
- The user's roles. A role is a collection of permissions and is used to grant a user access to data and specific IBM StoredIQ applications.
- The user's status, which can be **Active** or **Inactive**.
- The date and time of the most recent user login. If the user is not logged in to the system, **N/A** is shown.

Select a user to view more details about the account, or to edit or delete it.

The **View User Details** pane summarizes this information for a selected user. In addition, the pane provides a link to the user's access log. The log contains audit information for system activities, such as the IP address, the user name, the type of activity, and the date and time when the activity occurred.

To search the list of users for a specific user, either enter the complete user name in the search field or enter a partial string in this field to filter the list of users.

Tip: At any time, users can check their account settings by selecting **View Profile** from the drop-down menu next to the user name in the top navigation bar. For local user accounts, the profile also has a link for changing the password.

IBM StoredIQ Data Server users are managed separately from the IBM StoredIQ Administrator, IBM StoredIQ Data Workbench, IBM StoredIQ Insights, and IBM StoredIQ Cognitive Data Assessment user accounts. IBM StoredIQ Data Server administrative users are managed in the IBM StoredIQ Data Server user interface.

Managing IBM StoredIQ user accounts

Define user accounts and provide users with login information so that they can use the IBM StoredIQ applications IBM StoredIQ Administrator, IBM StoredIQ Data Workbench, and IBM StoredIQ Insights. You can create, import, edit, or delete such user accounts from the **Users** pane in IBM StoredIQ Administrator.

You can create user accounts manually or by importing users from an LDAP server or from LDAP-supported authentication systems, enabling enterprise security policy enforcement. Before you can do so, you must configure a connection to the LDAP server. For more information, see [“Connecting to the LDAP server”](#) on page 45.

To manage an account:

- Create an account either manually or by importing users.
 - a) When you create the account manually, you must provide the following information:
 - A short, unique name as the user name. This must not be the person's full name.
 - The user's given name.
 - The user's surname.
 - The user's email address.

When you import a user, these fields are filled when you select a user from the **LDAP Users** list and click **Save**.

For a manually created account (referred to as *local account*), set a password. Passwords must be at least eight characters in length and contain at least one of each of these characters: an uppercase letter, a lowercase letter, a number, and a special character. For an imported user, you cannot set a password. The user authenticates with the password that is stored on the LDAP server.

A user can log in to IBM StoredIQ with the user name or the email address.

b) Assign one or more of these roles:

Admin

This role grants access to all of the IBM StoredIQ components. A user with this role can perform all tasks within IBM StoredIQ but primarily this role grants the right to perform administration tasks.

For IBM StoredIQ Insights administration, this role must be combined with the **Insights User** role.

Data User

This role grants access to IBM StoredIQ Data Workbench. A user with this role can assess and manage the company's data. The rights granted with this role include the permission to preview data objects. Therefore, assign this role on a need-to-know basis.

Data User (No Preview)

This role also grants access to IBM StoredIQ Data Workbench. A user with this role can assess and manage the company's data. The rights granted with this role exclude the permission to preview data objects.

Insights User

This role grants access to IBM StoredIQ Insights. A user with this role can perform dynamic queries on the company's data for interactive data assessment. The rights granted with this role include the permission to preview the content and metadata of data objects.

To be able to save IBM StoredIQ Insights search results as infosets, the user must also have one of the **Data User** roles or the **Admin** role.

To grant IBM StoredIQ Insights administration rights, this role must be combined with the **Admin** role.

Insights Reviewer

This role grants access to IBM StoredIQ Insights but access is restricted to a subset of documents. The set of accessible documents (content and meta data) is determined by a reviewer tag assigned to documents. This reviewer tag is created in the **Reviewer** tag group when the role is assigned and corresponds to the user's first, last, and user name. Users with this role, the **Insights User** role, or administrators can assign reviewers to documents. If this role is removed from a user account, this user can no longer be assigned as a reviewer.

To be able to save IBM StoredIQ Insights search results as infosets, the user must also have one of the **Data User** roles or the **Admin** role.

CDA Project Contributor

This role grants access to the IBM StoredIQ Cognitive Data Assessment dashboard if the application is enabled. A user with this role can participate in the training and validation of CDA classification models.

CDA Project Owner

This role grants access to IBM StoredIQ Cognitive Data Assessment if the application is enabled. A user with this role can set up and manage training projects for CDA classification models. Users with this role can also make approved models available for use in other applications.

SDK User

This role grants access to the IBM StoredIQ SDK. A user with this role also must be assigned the **Admin** role to be able to make full use of the SDK. A user with just the **SDK User** role cannot log in to any of the IBM StoredIQ web user interfaces.

Important: When you assign more than one role to a user, the permissions are cumulative; the user has all of the permissions that are granted by the assigned roles.

c) Set the user's status.

The default status is **active**, which means that the user can log in and use the product. Users with the status **inactive** cannot log in to IBM StoredIQ. When inactive users try to log in, they are notified that their login is invalid and that they must contact an administrator.

- Edit an account to modify the settings, to activate a disabled account, or to reset the password.

For an imported user, you can change only the role and the status.

Note: User accounts become disabled after three failed log-in attempts. A disabled account is activated by resetting the password. However, as an administrator you can also change the account status from active to inactive and vice versa. For example, you might make an account inactive because the user is on a leave of absence.

- Delete an account.

Any log files that are associated with this user are also deleted. To avoid that, consider changing the user's status to **inactive** instead.

Connecting to the LDAP server

Before users can be imported and authenticated from an LDAP or LDAP-supported authentication system, you must connect to the LDAP server.

Your IBM StoredIQ application stack must be at version 7.6.0.6 or later.

1. Using an SSH tool, log in to the application stack as the `siqadmin` user.
2. Enter the following command: `ldapcfg`.
3. Enter these configuration details.

a) To enable LDAP, select **Allow External LDAP User**.

b) Provide the following information:

Parameter	Value
LDAP URL	<p>The host name or IP address of your directory server:</p> <ul style="list-style-type: none"> • For non-SSL connections: <pre>ldap://ldap-server-hostname ldap://ip-address</pre> • For SSL connections: <pre>ldaps://ldap-server-hostname ldaps://ip-address</pre>
LDAP User	The user account for accessing the directory server, for example, <code>cn=user,dc=example,dc=com</code>
LDAP Password	The LDAP user's password
Base DN	The base domain name, for example, <code>dc=siqdomain,dc=com</code>

The **LDAP Configuration** window also contains the following attribute-mapping details.

Attribute	Predefined mapping
First Name	<code>givenName</code>

Attribute	Predefined mapping
Last Name	sn
Email	mail
Username	cn This field cannot contain spaces or special characters. Additionally, it must be a part of your DN for users to import successfully.

Important: Do not modify these predefined attribute mappings unless your schema is different. If you have questions about the schema or these changes, contact your company's LDAP administrator.

c) Select **Test connection** to test whether you are connected to the LDAP server.

The **LDAP Status** information shows whether the connection test passed or failed.

4. Save and exit the configuration.

5. Restart the application stack by using this command: **systemctl restart appstack-uwsgi.service**

6. Check the status by using this command: **systemctl list-dependencies appstack.target**

The list of LDAP users is now available for import.

To avoid that the information in the application stack database becomes stale, the user details are synchronized with the directory server on a daily basis. A recurring background service avoids overloading a synchronization request on the LDAP server. However, you can trigger an immediate synchronization by running the `sync_ldap` script as `siqadmin` user.

Note: During synchronization, only active LDAP users' details are updated. If LDAP users are deleted from the LDAP server, those users' details are marked as inactive in the application stack database.

Managing user passwords

As an administrator, you set the initial passwords for local users, and you can later change or reset the passwords.

You set the initial password when you create a local user account. Valid passwords are at least 8 characters in length and contain at least one of each of these characters: an uppercase letter, a lowercase letter, a number, and a special character. At any time, you can change a local user's password, which might become necessary for one of the following reasons:

- An active user account became disabled after three failed login attempts and must be re-enabled.
- A user requested a new password either via the **Forgot your password?** link or by contacting you as directed in the **Did not receive password details** or **Changed your email address** information in the login window.

Local users can change the password by selecting **View Profile** from the drop-down menu next to the user name and clicking the provided link.

Passwords of LDAP users must be changed by using LDAP administration tools.

To change a user password:

1. From the **Users** pane, select the user for whom you want to change the password.
2. Edit the account and specify a new password.

The new password must be 8 to 64 characters long and must contain characters from at least three of these categories:

- Uppercase letters: A - Z
- Lowercase letters: a - z

- Digits: 0 - 9
 - Punctuation marks or special characters: . : = * ^ / \$ # _ + @ & % -
3. To notify the user of the change, click **Send reset password email**.
This email requests the user to contact the system administrator.
 4. Save the settings.

Action management

Actions represent executable processes that act upon indexed data within data servers. Use the **Action Management** page to create and modify actions that are used in IBM StoredIQ Data Workbench.

With IBM StoredIQ Administrator, you can create these types of actions:

- **Copy.** The Copy action copies infosets to a target set.
 - You can specify the directory structure for copied data objects.
 - If a data object exists with the same name, then the subsequent data objects that are encountered are renamed.

Tip: When you copy data between CIFS volumes, the owner of the objects on the target server usually is the user whose credentials are used for connecting to the target server. To preserve the ownership of the objects, enable that capability as described in the respective instructions in the Data Server administration guide.
- **Copy to Retention.** The Copy to Retention action copies data objects from source volumes to a target set for a specified period.
- **Delete.** The Delete action removes data objects from the source volume.
 - Only data objects, not directories, are deleted.
 - Use caution when you are deleting data objects.
 - Delete actions are limited to system level objects. If a responsive object targeted for deletion is within a container object (such as a .zip or .tar file), the container level object is the object that is deleted. Thus, any other objects within the container file are also deleted. The respective logging information reflects the container level object, indicating success or failure of the action.
 - A Delete action on a responsive object attached to a CIFS based .msg, .eml, or .pst file also results in the container object being deleted including any other objects attached to or contained within the container object. The respective logging information reflects the container level object, indicating success or failure of the action.
 - Delete actions are supported only for .
- **Discovery Export EDRM XML.** This Discovery Export action copies data objects and generates an EDRM XML file for loading into third-party legal review tools.
- **Discovery Export DAT.** This Discovery Export action copies data objects and generates a Concordance DAT file for loading into third-party legal review tools.
- **Discovery Export DAT Light.** This Discovery Export action copies data objects and generates a Concordance DAT (Light) file for loading into third-party legal review tools.
- **Modify Attribute.** The Modify Attribute action simulates setting retention by manipulating specific attributes on objects.
- **Move.** The Move action moves data objects in an infoset to another volume.
 - You can specify the directory structure for moved data objects.
 - If a data object exists with the same name, then the data object is renamed.
 - When the source equals the destination, the move is not allowed.

Tip: When you move data between CIFS volumes, the owner of the objects on the target server usually is the user whose credentials are used for connecting to the target server. To preserve the ownership of the objects, enable that capability as described in the respective instructions in the Data Server administration guide.
- **Step-up Analytics.** The Step-up Analytics action applies the analysis logic of a cartridge to a given infoset.

In addition, the following built-in actions are available:

- **Step-up Full-Text.** The Step-up Full-Text action can be run on an infoset, providing content to the full-text index for objects that are infoset members.
- **Step-up Snippet.** The Step-up Snippet action can be run on an infoset, extracting and storing text for any data object to which it is applied. These data snippets are then used by auto-classification. Infosets objects without snippets are not classified for auto-classification.

Note: When scheduled harvests or actions are run against a data source, snippets are not updated or generated if they are missing.

Table 25. Actions column names and descriptions

Actions Column Name	Description
Action Name	Lists the name of the action.
Type	Lists the action type. They are Copy, Copy to Retention, Delete, Modify Attribute, Move, Discovery Export XML, Discovery Export DAT, Discovery Export DAT Light, Step-up Analytics, Step-up Full-Text, and Step-up Snippet.
Last Modified	Lists the date and time at which the action was last modified.
Target Set Name	Lists the destined target set of the action.
Description	Lists the action description.

To search for a specific action, either enter the complete action name in the search field or enter a partial string in this field to filter the list of actions.

Creating an action

Actions are processes that are enacted on data objects, which are defined by infosets. Actions must be defined here in IBM StoredIQ Administrator.

The Step-up Snippet and Step-up Full-Text actions are built-in actions, which means that no options are provided within the user interface. You cannot create or edit the Step-up Snippet or Step-up Full-Text action as they are provided as prepopulated options on the **Actions** tab.

1. Click **Actions > Create Action**.

2. In the **Add Action** dialog box, complete these fields.

a) Enter the name of this action and select the action type.

You can create a Step-up Analytics action only if at least one cartridge was uploaded.

When you name the action, avoid using commas (.). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.

b) Enter a brief description of this action.

Make it as accurate as possible, so that data experts can easily discern the purpose of the action.

c) If you are creating a copy, copy to retention, discovery export, modify attribute, or move action, select a target set or click **Create a new target set**.

If the action is copy to Box, you must select or create a target set with **Box** as its source type.

d) Click **Next**.

Depending on the selected action type, different additional options are available.

Option	Description	Action types
Attribute Type	Select the access mode to set for files in an infoset on CIFS and NFS data sources: to read-only or read/write.	Modify Attribute

Option	Description	Action types
Cartridges	Select one or more cartridges to apply with the action.	Step-up Analytics
Copy Modified Files	Select this option if you want to copy only those files that were modified after the last harvest.	Copy Copy to Retention Move
Destination Directory	Specify the file path to the target directory.	Copy Discovery Export EDRMXML Discovery Export DAT Discovery Export DAT-Light Move
Directory Structure > Recreate directory structure on destination	Select this option to replicate the source directory structure on the target. This option is enabled by default when a Copy to Box target set action is configured.	Copy Move
Do not delete data objects that have been accessed since last harvest.	Select this option if you do not want any data objects that were accessed since the last harvest to be deleted.	Delete
Email Message Format	Select the format of the email message.	Discovery Export EDRMXML Discovery Export DAT Discovery Export DAT-Light
Export data objects modified since last harvested	Select this option to export only those data objects that were modified since last harvest.	Discovery Export DAT
Export Modified Files	Select this option to export only those data objects that were modified since last harvest.	Discovery Export EDRMXML Discovery Export DAT-Light
Harvesting > Do not auto-harvest destination volume after copy	Select this option to suppress auto-harvesting of target volume after the copy is complete. For Copy to Retention actions, you can later harvest the retention volume with the discover retention volumes job on the data server.	Copy Copy to Retention
Hashes > Compute a hash value for each data object	Select this option to have a hash value generated for each copied data object.	Copy Copy to Retention

Option	Description	Action types
Limit Export Process to	Select this option and specify the maximum number of data objects to limit the number of exported files.	Discovery Export EDRMXML Discovery Export DAT Discovery Export DAT-Light
Mapping <ul style="list-style-type: none"> • Map permissions from source to destination • Preserve owners from source to destination 	The mapping options in general are available for Box target sets only. Permission mapping is supported only for CIFS source volumes. Preserving owners is supported only for SharePoint and CIFS source volumes.	Copy
Others > Preserve version chains on destination	This option is available for Box target sets only.	Copy
Retention Period	Define the time period for retention.	Copy to Retention
Retention Tag	Specify a tag that is attached to and retained in the copied and harvested data object.	Copy to Retention
Save a text copy	Select this option to save a plain text version of the data objects.	Discovery Export EDRMXML
SharePoint User Profile Items <ul style="list-style-type: none"> • Notes • Libraries • Blog Posts • Wikis • Misc 	Select the SharePoint object types to which the action should be applied.	Copy Copy to Retention Discovery Export EDRMXML Discovery Export DAT Discovery Export DAT-Light
Unique ID Prefix and Number of Digits Pad Zeros	Enter a unique prefix and the number of digits to follow that prefix to generate unique document IDs. You can also select to pad the document ID with zeros.	Discovery Export EDRMXML Discovery Export DAT Discovery Export DAT-Light

3. Select **Save**, **Back**, or **Cancel** as needed.

Creating a Step-up Analytics action

A Step-up Analytics action is required to apply the logic of a cartridge.

Follow these steps to create a Step-up Analytics action.

Remember: You don't have to create an action to apply the logic contained in the governance cartridge that is created when data is synchronized with the governance catalog. The required action named **Governance Analytics** is automatically created on the first synchronization run and updated as required during subsequent runs.

1. Click **Actions**, and then click **Create Action**.

The **Add Action** dialog box appears.

2. In the Add Action dialog box, complete these fields.

- a) Enter a name for the Step-up Analytics in the **Action Name** field.
- b) Select Step-up Analytics from the **Action Type** list.
- c) Type a description for the action in the **Description of Action** field.
- d) Click **Next**.

A list of available cartridges displays for you to choose. One or more cartridges must be selected. Otherwise, you cannot save the action.

- e) Click **Save**.

Note: You cannot change the cartridges that are selected for a Step-up Analytics action after you create the Step-up Analytics action. If you do, an error message displays and the action cannot be saved.

Managing actions

You can edit, clone, or delete actions from the **Action Management** pane.

To manage actions, navigate to the list of available actions by clicking **Actions**.

You can select either of the following options:

- Clone an action to duplicate actions.
 - a) Select the action that you want to clone and click **Clone**.
 - b) Enter the name of the new action and update the description as required.

The name of the cloned action must be different from the name of the original action. You cannot change the action type. When you update the description, make it as accurate as possible, so that data experts can easily discern the purpose of the action.
 - c) The action type determines what other options are available. Update these options as applicable.
 - d) Save your settings.
- Edit an existing action.
 - a) Select the action that you want to edit and click **Edit**.
 - b) Edit the action as needed.
 - c) Save your settings.

- Delete an action.

The action is removed from the action list and it is no longer available for use.

Target sets

Some actions require a destination to complete, and a target set represents a mapping for each data server of the wanted target volume. To copy files, the data server must know where to copy the file to. Each data server must be configured with a list of target volumes, such as retention and discovery export volumes. Target sets are managed in IBM StoredIQ Administrator.

Target Set Example

Data Server 1 has three retention volumes: Volume A, Volume B, and Volume C.

Data Server 2 has two retention volumes: Volume B, Volume X, and Volume Z.

Data Server 3 has two retention volumes: Volume B and Volume Z.

Resulting Copy Action: You would like to create a target set for a copy action

Data Server 1 copies to Volume A.

Data Server 2 copies to Volume X.

Data Server 3 copies to Volume B.

Another user might want to create a target set for a copy action where Data Server 1, Data Server 2, and Data Server 3 copies to Volume B.

The **Target Sets** pane lists all available target sets and provides the following information:

- The name of the target set.
- The type of the target set, namely **Primary**, **Retention**, or **Export**.
- The target set's source type: **Box**, **CIFS**, **CMIS**, **Documentum**, **FileNet**, **HDFS**, **IBM Content Manager**, **NFS**, **OneDrive**, or **SharePoint**.
- The number of volumes in that target set.
- The date and time that the target set was last modified.
- The description of the target set.

Select a target set to edit, clone, or delete it.

To search the list of target sets for a specific set, either enter the complete target set name in the search field or enter a partial string in this field to filter the list of target sets.

Managing target sets

Retention, copy, move, and discovery export actions require a named destination to which data objects can be copied, moved, or exported. Target sets associate data servers and appropriate volumes, providing your actions with a destination.

Manage target sets in IBM StoredIQ Administrator: create new target sets or edit, clone, or delete existing ones.

In IBM StoredIQ Administrator, click **Target Sets** from the side navigation bar:

- To create a target set, click **Create Set**.

In the **Add Target Set** dialog box, provide the following information:

- a) Enter a name and a description for the target set.

When you name the target set, avoid using commas (,). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.

- b) Select the type of this target set from **Primary**, **Retention**, or **Export**.

- c) Select the source type of this target set. The list to choose from varies depending on the selected type.

- d) In the **Target Set List** area, click **Add Data Servers**.

Select a data server and then click **Add** to move it to the **Target Set List** box. To delete a data server from the list, select the entry and click **Remove**. Use the **Ctrl** key to add or remove several selections simultaneously.

Click **Done** when you're finished adding data servers for the target set. At any time, you can modify the list by clicking **Add/Remove Data Server(s)** in the **Create Target Set** dialog box.

- e) Back in the **Create Target Set** dialog box, select a data server, and then select a volume on that data server. Repeat this step for each data server.
 - f) Save the definition when you're done.
- To edit a target set, select the target set and click **Edit Set**.
 - To duplicate a target set, select the target set and click **Clone Set**.
The cloned target set's name must be unique.
 - To delete a target set, select the target set and click **Delete Set**.

Reports

IBM StoredIQ contains a number of default reports and the ability to upload a report package.

IBM StoredIQ provides features to inventory unstructured data. It can then collect and evaluate unstructured data, compiling reports and jobs to act on the data according to an organization's needs. Reports can then be printed and disseminated, allowing teams to analyze data in an ad hoc manner. New reports can be added by uploading a report package.

IBM StoredIQ Administrator provides a number of built-in reports, such as summaries of data objects in the system, storage use, and the number of identical documents in the system. You can create custom reports, including Query Analysis Reports for e-discovery purposes, and automatically send email report notifications to administrators and other interested parties. By default, these types of reports are available.

- Compliance Report
- Content Collector Manifest File CSV Export
- CSV All Audited Object Exports
- CSV Attribute Summary
- CSV Exception List Export
- CSV Infoset Data Object Export
- CSV Term Hit Details Export
- Data Assessment Report
- Data Privacy Report Card
- Data Topology Report
- Duplication Summary Report
- Overlay Hit Report
- Term Hit Report

For some report types, you can configure specific report settings. For details, see [“Configuring report settings”](#) on page 55.

Configuring report settings

For some reports, you can modify the configuration to determine the scope information that is included in the report.

1. Using an SSH tool, log in to the AppStack as `siqadmin` user.
2. Run the **`edit_etc_siq_report_conf`** command to edit the `report.conf` file and update the settings as required.

- To include the IBM StoredIQ unique node ID in the Custom CSV Export, CSV Infoset Data Object Export, and CSV All Audited Objects Export reports, set `use_node_id` to `true`.

The CSV Term Hit Details Export and CSV Term Hit Count reports include the IBM StoredIQ node ID column by default.

- To change the maximum number of rows that are exported to a CSV report, edit the `csv_row_max_limit` setting. The default value is 64000.
- To change the interval at which the application stack checks the status of a CSV Term Hit Details Export or CSV Term Hit Count report on the data server, edit the `as_report_sync_interval` setting. Specify the value in seconds. The default value is 60.

- To define the minimum disk space that must be available in the /tmp directory on the application stack when CSV Term Hit Details Export or CSV Term Hit Count report files are transferred from the data server, edit the `min_diskspace_threshold` setting. The default value is 2147483648 (2 GB).
- To define the maximum number of terms that are to be retrieved from each document when running the CSV Term Hit Details Export report, edit the `term_hit_limit` setting. The default value is -1, which means that no term limit is enforced.
- To change the search criteria for the Data Privacy Report Card, edit the entries for the categories `high_sensitive`, `medium_sensitive`, and `low_sensitive`. If you remove a category or leave it empty, the category will still show up in the report with a value of zero. Data that does not match any of the categories will be listed as not sensitive.

These types of search terms are allowed:

- Custom terms and filters following the full-text syntax for filters. For details about the handling of special characters or punctuation in search terms, see the topic about extended ASCII characters.
- Full-text macros in the format `{macro}`
- Indexed annotations in the format `ia:filter_term`

Search terms within a category are ORed whether they are separated by commas or by spaces. For example, for the search criteria `high_sensitive={SSN}, John Doe, ia:BankAccountNumber`, the report returns results for `{SSN} OR John OR Doe OR ia:BankAccountNumber`.

The asterisk (*) and the question mark (?) are treated as wildcard characters.

Usage of single or double quotation marks to search for exact matches is not supported.

Make sure to inform your data workbench users about the configured search criteria, so that they can run the appropriate Step-up actions before running the Data Privacy Report Card. Otherwise, the report will not return meaningful results.

- To include a document digest column in CSV All Audited Objects Export reports, set `use_digest_id` to true.

3. Save your changes.

4. To put the changes into effect, restart the AppStack by running this command:

```
systemctl restart appstack.target
```

Creating a report

Provide custom reports in addition to the built-in reports.

You can create and upload custom report definitions so that IBM StoredIQ Data Workbench users can generate additional reports. For more information about developing such reports, see [“Custom reports” on page 57](#).

To upload a custom report definition:

1. Click **Reports**, and then click **Create**.

2. In the **Create Report** window:

a) Enter a name for this report.

When you name the report, avoid using commas (,). Although a comma is a valid character, it causes queries to fail when specified as part of the search text.

b) Enter a description of this report so that users can easily discern its purpose.

c) Click **Choose File**, and then find and select the appropriate report package.

Report packages are compressed files that contain at least a report definition in the format of a `.rptdesign` file and, possibly, any stylesheets and images that are required for proper rendering of the report.

3. Click **Create**.

The newly created report is added to the report list and is available in IBM StoredIQ Data Workbench.

Custom reports

Reporting is a key step within the data-management process as it validates that processes were completed correctly within IBM StoredIQ.

You can customize reports in any of these scenarios:

- Design reports to carry your organization's custom styles, logos, aligning them with other organization-based artifacts and documentation.
- Design reports to contain information that is not found in other, existing reports.

Using ODBC or any client that supports remote access to PostgreSQL, you can connect to the "enamel" database IBM StoredIQ AppStack, locating a schema that is named `report_schema`. This schema can be accessed remotely by logging in as a `reportuser`, which is an account that is created for the purposes of reporting on IBM StoredIQ data.

The default password for the `reportuser` account is `St0red1q`. For security purposes, you should change the password. You can do so at any time by running the `change_reportuser_password` script as `siqadmin` user on the AppStack. You are asked to enter and confirm the new password. Your input is not visible on the screen.

To allow for access to the schema, port 5432 must be opened. Port 5432 is the port for communication with the PostgreSQL instance. Thus, opening this port allows for communication with the entire instance in general, which means with all database tables. Therefore, you should use caution and make sure to close the port again after completing the report design. Running a report does not require this port to be open.

To open or close the port, run the `postgres_port_open` or `postgres_port_close` script as `siqadmin` user.

Note: For further help in customizing IBM StoredIQ reports, contact your technical service representative.

The following report views are available for reporting and primarily carry information about infosets and filters that are created in IBM StoredIQ Data Workbench. They can be accessed by means of SQL SELECT queries.

infoset

This view contains the available infosets with their metadata.

infoset_overlay

This view contains the infoset overlays that are defined for an infoset.

named_filter

This view contains the saved filters that are available in IBM StoredIQ with their metadata and their internal IDs. With these IDs, you can look up the filters referenced in the "infoset_filter" view.

query_filter

This view contains the `ia`: query filters that are available in IBM StoredIQ.

Working with the BIRT Report Designer

The Eclipse-based BIRT Report Designer provides core reporting features such as report layout, data source configuration, and scripting features that render content and layout dynamically during report generation. You must set up the BIRT Report Designer before you can develop custom reports.

The BIRT Report Designer is a stand-alone Eclipse project bundle that can be downloaded and installed on any machine. Details about the BIRT Designer are available here:

<http://www.eclipse.org/birt/about/designer.php>

Download and install the Report Designer Full Eclipse Install package. The latest versions of the BIRT packages are available from this download page:

<http://download.eclipse.org/birt/downloads/>

However, the instructions apply to BIRT version 4.6. To download this specific version, navigate to the BIRT 4.6.0 download site by selecting **Recent Builds > 4.6.0**.

You can also extend an existing Eclipse installation by adding the Report Development feature. Within the Eclipse application, go to **Help > Install New Software** and click **Add** to add a site for BIRT:

<http://download.eclipse.org/birt/update-site/4.6>.

This URL is for BIRT version 4.6.

A JDBC driver is required to access the "enamel" database from BIRT. Download the latest PostgreSQL JDBC driver and save it in a location that is accessible by BIRT:

<http://jdbc.postgresql.org/download.html>

To create a report definition:

1. Create a report.

After you set up your development environment, you must create a new report in the BIRT Report Designer tool before you can develop your custom report.

- a) Open BIRT Report Designer and start a new project by selecting **File > New > Project**.
- b) Expand **Business Intelligence and Reporting Tools**, select **Report Project** and click **Next**.
- c) Select a suitable project name and location, and click **Finish**.
- d) Now, start a new report by selecting **File > New > Report**.
- e) Select the project that you created previously.
Enter a suitable name for the new report.
- f) Select a template to start with, for example, the **Blank Report** template.
- g) Click **Finish**.

2. Connect to the reporting database:

- a) Create a data source by selecting **Data > New Data Source**.
- b) Select **JDBC Data Source**, specify a suitable name and click **Next**.
- c) Click **Manage Drivers**, then click **Add** and select the JDBC driver that you downloaded as a prerequisite. Click **OK**.
- d) In the **Create a new data source** window, select **org.postgresql.Driver** in the **Driver Class** list and provide the rest of the information as described:
 - **Database URL:** `jdbc:postgresql://appstack-host:5432/enamel`
 - **User name:** `reportuser`
 - **Password:** Contact your administrator for the `reportuser` password to specify here.
 - **JNDI URL:** You can leave this field empty.
- e) To confirm that the connection settings are working, click **Test Connection**.
- f) Click **Finish**.

Alternatively, you can create a connection profile for repeated use:

- a) From the main menu, select **File > New > Other**.
- b) From **Connection Profiles**, select **Connection Profile** and click **Next**.
- c) Select **BIRT JDBC Data Source** as connection profile type.
- d) Enter a unique name for the connection profile and optionally a description. Then, click **Next**.
- e) Provide the required connection information as described before.

- f) Test the connection to verify that the connection profile is working.
 - g) Click **Finish**.
3. Define required report parameters.
- Such report parameters must be provided when a user renders the report. Currently, IBM StoredIQ supports only one report parameter, which must map to the infoSet ID.
- a) From the main menu, select **Data > New Report Parameter > New Parameter**.
 - b) Enter a descriptive name such as `infoSet_id`.
 - c) In the **Prompt text** field, enter the message that must be displayed to the user who renders the report.
 - d) Set other properties as follows.
 - For this parameter, the data type must be set to **String**, the display type must be **Text Box**, and the **Is Required** option must be selected.
 - e) Click **OK**.
4. Create data sets.

A **Data Set** is subset of the data available within a data source. It is typically defined by providing a SQL query including sort and filter criteria. This defines the data on which the report is based.

After you set up a data source and created report parameters, a new data set can be created in the following way:

- a) In the Data Explorer view, right-click **Data Set** and select **New Data Set**.
- b) In the **New Data Set** window, select the data source that you created previously.
- c) Select **SQL Select Query** from the **Data Set Type** list and enter a suitable name in the **Data Set Name** field. Then, click **Next**.
- d) Under **Query Text**, enter the SQL query that is used to fetch the data specific to the report.
 - If you want to target the data of a specific infoSet, include a where clause:

```
select * from report_schema.infoSet where infoSet_id = ?;
```

The ? denotes that a report parameter must be used here. In this case, you must add a parameter to the data set.

- e) Optional: To add a report parameter, select **Parameters** and click **New**.
 - Then, change the **Linked to Report Parameter** value from **None** to the name of your report parameter that is created during data set creation and click **OK**.
 - You can refine the data set at any time by editing it.
 - f) Click **Preview Results** to view the data fetched from the query.
 - If everything looks correct, click **OK** to close the window.
5. Design the report.
- BIRT Designer provides a variety of controls and display elements that can be used to design engaging reports. These controls are available in the **Palette** window and can be added to the report by dragging them there.
6. Test your design.
- Before you upload the customized report, preview the report as both HTML and PDF output to ensure that the report renders properly. From the main menu, select **Run > View Report** and then select the appropriate rendering: **as HTML** or **as PDF**.
7. If everything looks fine, you can upload the report package in IBM StoredIQ Administrator.
- IBM StoredIQ supports uploading a report design in a compressed file. Report designs are expected to be in the `rptdesign` format within the compressed file. Currently, only one `rptdesign` is expected to be present within a compressed file. The design can use one or more `.css` files and one or more image files, which can all be packaged within the compressed file.

To better manage the artifacts such as the images and stylesheets for a single report design package, you must do one of the following tasks:

- Create a project for each report design. Each project contains only one `rptdesign` file.
- Create a folder for each report design within the project. Each folder contains only one `rptdesign` file.

By selecting one of the report-design management options, you can manage all images and stylesheets. They are used by the design in a single location, making the creation of the compressed file easier as you compress a folder. Additionally, the uploaded compressed file must meet the following requirements:

- It contains only one `rptdesign` file and uses the file extension `.rptdesign`.
- All CSS and images that are referenced from the `rptdesign` file are also present within the compressed file.
- Only one report parameter is defined in the `rptdesign` file.
- Images use these extensions: `.bmp`, `.jpg`, `.jpeg`, `.jpe`, `.jfif`, `.gif`, `.png`, `.tif`, `.tiff`, `.ico`, or `.svg`.
- Stylesheets must use the `.css` file extension.

Deploying additional fonts

You can deploy additional fonts for report rendering.

You must be logged in to the IBM StoredIQ AppStack as `root` user.

By default, the IBM StoredIQ reporting engine supports the following fonts that are provided with the product:

- `icon-webfont.ttf`
- `ubuntu-bold-italic.ttf`
- `ubuntu-bold.ttf`
- `ubuntu-medium-italic.ttf`
- `ubuntu-medium.ttf`
- `ubuntu-regular-italic.ttf`
- `ubuntu-regular.ttf`

Restriction: Only True Type Fonts are supported. If a font is used in the report design and is not deployed correctly on the application stack, the default Times New Roman font is used when you render the report.

1. Deploy the fonts to the AppStack.

Copy the fonts to this folder on the application stack: `/usr/share/fonts/default/TrueType`

2. Restart IBM StoredIQ services:

- a) Issue the command **`bootstrap server.mode idle`** and wait until **`bootstrap server.mode`** shows idle.
- b) Run this command: **`bootstrap server.mode production`**
- c) To check the status, run the command **`bootstrap server.mode`** again.

Auto-classification models

Auto-classification models, also called automated document categorization, integrates the IBM® Content Classification's classification model into the IBM StoredIQ info-set-generation process.

IBM Content Classification

IBM Content Classification helps organize unstructured content by analyzing the full text of documents and emails and applying rules that automate classification decisions. By managing documents and email, you organize and act on content every day. IBM Content Classification reduces the burden of manual decision making that is done by employees by accurately and automatically organizing information. Embedded with natural language processing and semantic analysis capabilities, IBM Content Classification determines the true intent of words and then uses that knowledge to automate decision making. Unlike other classification systems that are based on rules only, IBM Content Classification combines rules and contextual analysis to incorporate synchronous learning that adapts to changing business needs. As a result, classification becomes accurate over time.

IBM Content Classification can organize information by policies or keywords. The classification process not only searches for a single word or phrase, but also analyzes the entire document. It then distills the main point of the text and assigns the text to a category. When it analyzes content, IBM Content Classification can recognize misspellings, abbreviations, jargon, and technical terms.

Accuracy improves over time because the system adapts to the unique nature of your business by identifying different categories from examples that you provide. You provide feedback, and a deferred retraining operation adjusts the model and implements corrections based on that feedback. The accuracy of the classification results keeps pace with changes in your business.

IBM Content Classification combines this context-based approach with a rule-based, decision-making approach. The system can identify keywords, patterns, such as account numbers and phone numbers of case identifiers, and words within certain proximity of each other. For example, the system can identify occurrences of the phrase "Attorney General" in the same sentence as the word "California". When content that matches a condition in a rule is detected, the action that is defined for the rule is applied, and the document or email is classified.

Examples of classification applications

You can use IBM Content Classification to achieve various business goals.

- Enterprise content standardization. To support document classification and taxonomy automation within your content management system, document properties or metadata can be automatically assigned when the content is classified. Documents can be automatically moved to the correct enterprise repository.
- Compliance and records management. Documents and email can be declared as records when they are classified and placed under the control of record retention policies and compliance standards.
- E-discovery readiness. Documents and email can be filtered to ensure that only items with business value are classified and archived. You can quickly and cost effectively prepare content for potential legal notices.
- Business process optimization. Automated decision making ensures more consistent outcomes and reduced costs. With content-based analysis, you accomplish these actions.
 - Insert documents into the workflow of a business plan.
 - Reroute email.
 - Suggest and apply agent responses within a customer-support center.

Using auto-classification, you can create and update classification models to help map data within info-sets. The Step-up Snippet action extracts and stores text on data objects within an info-set. For auto-

classification to work, you must use the Step-up Snippet action since infoSet objects without extracted snippets do not have auto-classification scores.

The Data Expert can use IBM Content Classification to train a classification model. The classification model is then registered with IBM StoredIQ Administrator. The registered classification model can be applied to an existing infoSet in IBM StoredIQ Data Workbench to generate new metadata for the objects in the infoSet. Finally, this metadata can be used in rule-based filters to create new infoSets.

Note: To use the IBM StoredIQ Administrator auto-classification function, users must have a licensed installation of the IBM Content Classification on their desktop.

Applying the auto-classification feature

To use the auto-classification feature successfully, you must work in both IBM StoredIQ Administrator and IBM StoredIQ Data Workbench.

Prior to preparing an IBM Content Classification model for import, you must finish downloading the classification export utility, which must be done only one time. The export utility is used for each IBM Content Classification model that you want to import.

IBM Content Classification works by taking a sample set of documents that you adjust, constructing a classification model, and then classifying a new document. The IBM StoredIQ defines the classification model as a decision plan with at least one knowledge base that is referenced by the decision plan.

1. In IBM StoredIQ Administrator, in the **Action** pane, the Step-up Snippet action is created automatically and appears in IBM StoredIQ Data Workbench.
2. In IBM StoredIQ Administrator, create an auto-classification model by importing a compressed file from IBM Content Classification.
3. In IBM StoredIQ Data Workbench, start the Step-up Snippet action, which is available as **Action** on the side navigation bar.

Note: The infoSet must be a user infoSet, not a system infoSet.

4. While still in IBM StoredIQ Data Workbench, apply the wanted auto-classification enhancement, which is available on the **Enhance** tab, to the same infoSet.
5. After the enhancement completes in IBM StoredIQ Data Workbench, go to **Create > Build** and apply the auto-class filter to create a new infoSet.
6. Review the results of the newly created infoSet.

Data server sizing-configuration guidelines

The inclusion of the IBM Content Classification server on the data server changed the required memory and processors. Data server sizing-configuration guidelines are provided here.

To support the auto-classification function, you must increase the virtual machine settings for processors and memory on data servers to these new minimums:

- Memory: 16 GB
- vCPUs: 4

Required storage on the data server for the extracted text (snippet cache) is roughly 13% of the size of the uncompressed content, or 130 GB per 1 TB. The size can vary depending on the object type. The maximum amount of extracted text per item is configured to not exceed 2 MB by default.

A data server with 2 TB of configured storage can support approximately 10 TB of managed storage, assuming 50 million data objects with an average size of 200 KB. Approximately 300 GB is metadata storage with the remaining 1,700 GB supporting a combination of full text and extracted text.

Downloading the classification export utility

The classification export utility is a one-time-only download that must be installed before you import a model.

1. In IBM StoredIQ Administrator on the **Auto-Classification** tab, click **Download Classification Export Utility**.

The `siq-classification-export.zip` file is downloaded.

2. Extract the export utility to the `bin` directory of your IBM Content Classification product installation. For Windows, it is typically in the `C:\IBM\ContentClassification\Bin` directory.

Building an auto-classification model

Creating rules to find documents that fit differing categories is time-consuming and requires constant, meticulous adjustments. However, importing a classification model with sets of training documents helps find other, similar documents.

Using previously harvested data, you can create an auto-classification model.

1. Determine the categories into which you want the auto-classification model to classify documents.
2. Using IBM StoredIQ Data Workbench, create a filter for each category to capture documents that are representatives of the category.
3. For each filter, create an info set. The members of the resulting info set become the "training corpus" for the category.
4. For each info set, run a copy action with IBM StoredIQ Data Workbench onto a folder that is accessible by the IBM Content Classification application.
5. Use the IBM Content Classification application to create a decision plan and knowledge base by importing the training corpus that you created.

Note: A classification model consists of one decision plan and at least one knowledge base, which is a requirement of the IBM StoredIQ auto-classification feature.

For best practices to create an Auto-classification model, see [“Best practices for creating an Auto-classification model”](#) on page 67.

Preparing an IBM Content Classification model for import

You must prepare the IBM Content Classification model for importing into IBM StoredIQ with the IBM Content Classification tool `BundleDPKB.exe`.

IBM StoredIQ defines the IBM Content Classification model as a set of a single decision plan and one or more knowledge bases that are referenced by the decision plan. You must prepare each classification model for import by bundling the relevant decision plan and knowledge base with the IBM Content Classification export utility `BundleDPKB.exe`.

Note: Before you prepare an IBM Content Classification model for import, you must finish downloading and installing the classification export utility.

1. In a command interface, go to the `BundleDPKB` installation directory and then enter `cscript.exe ClassificationModelExporter.vbs <location of decision plan projects and knowledge bases> <decision plan directory> <optional empty directory> <optional zip file name>`.

If you did not specify a directory to which the compressed file must be exported, an export directory is created in the `<location of decision plan projects and knowledge bases>` location.

2. Note the location and name of the generated compressed file.

Importing an auto-classification model

You must complete these tasks before you can import an auto-classification model.

- Download the classification export utility. The procedure must be done only one time.
- Prepare the IBM Content Classification model for import. This procedure must be completed for every model you want to import.

To import a classification model, select **Auto Classification** from the side navigation bar. The right pane displays the imported models as described in this table.

Classification model column name	Description
Classification model name	Lists the name of the classification model.
Attribute name	Lists the name of the attribute that is used by the classification model.
Status	Lists the classification model's status, specifically Available , or Pending .
Creation date	Lists the date and time at which the classification model was created.
Description	Lists the description of the classification model.

1. In the Auto Classification pane, click **Create Model**.

The Create Model dialog box appears.

2. In the **Classification Model Name** text box, enter a unique name for this auto-classification model.
3. In the **Attribute Name** text box, enter a unique name for this attribute.

The **Attribute Name** must be unique within auto-classifications as it cannot be changed after the classification model is imported.

4. In the **Description** text box, enter a description of this model.
5. In the **Auto-Classification Model File** text box, click **Browse** and select the classification model file that you want to use when you are creating the model. Generally, this file is a compressed file that is created in the classification export utility. It consists of elements that are created in IBM Content Classification by the Data Expert.
6. Click **Save**, and the newly created model appears in the list of classification models.

Selecting a classification model

By selecting a classification model, you can view the details of that model, including the status of the last retrain, how much feedback it has received, and how many infosets have been enhanced using this model.

1. Within IBM StoredIQ Administrator, click **Auto-Classification**.

The List of Classification Models page appears.

2. Select the model for which you want to view details. The **<model name> Details** panel appears.

<i>Table 27. Row names and descriptions of the <model name> Details panel</i>	
Row name	Description
Status	Lists the classification model's status, specifically Available , Defunct , Deleting , Failed , Invalid , Pending , Processing , or Retraining . Note: In the Processing or Retraining states, models cannot be edited.
Last retrained	Lists the date and time this classification model was last retrained.
Last retrain status	Lists the classification model's retrain status. Status messages are Success , Failed , and N/A .
Total feedbacks	Lists the number of times the model has received feedback.
Total feedbacks after last retrain	Lists the number of times this model received feedback after its last retrain date.
Total infosets created using this model	Lists the number of infosets that use this classification model.
Total infosets enhanced by this model	Lists the number of infosets that have been enhanced by this classification model.

Searching for a specific classification model

When you have multiple classification models to choose, you can search for a specific classification model instead of using the slider to move through available options.

1. In the **Enter key terms(s)** text box, enter the name of the classification model for which you want to search, and then press **Search**. You can enter either the full classification model name or a portion of it.

To remove the search term, click the **X** to the right of the text box.

Classification models that match the entered search term are returned.

2. Edit, delete, or retrain the classification model as needed.

Editing a model

The terms and formats that are associated with a particular classification model change over time, which means that the classification model must also change. Classification models can be edited or updated, modifying the model file that is associated with that particular classification model.

1. In the Auto Classification pane, click the name of the classification model that you want to edit, and then click **Edit**.

The **Edit Model** window appears.

2. Edit the classification model fields as needed.

Note: Attribute names cannot be changed, but you can edit the name of the model without specifying a new model (.ZIP) file. If you attempt to modify a classification model that is being used, a warning appears. Click **OK** to return to the **Edit Model** window.

3. Click **Save**.

Deleting a model

Existing classification models can be deleted from the list of available models, if it was not used to enhance an info set and that info set still exists.

1. In the Auto Classification pane, select the classification model that you want to delete, and then click **Delete Model**.

The **Delete Model** window appears.

2. Click **Delete**.

The classification model is removed from the **Auto Classification** pane. It is no longer available for use.

Retraining a classification model

Retraining a classification model takes two personas to complete. The administrator must create a classification model for the data expert to prepare feedback. The administrator starts the retraining process and the data expert runs the enhancement to improve its accuracy.

With the assistance of IBM software support, modify the data-learning settings.

1. Click **Project > Project Options**.
2. In the **Project Options** window, click **Advanced**.
3. Select **Maintain learning data set with knowledge base**. Accept all other default values.
4. Click **OK**.

The data-learning settings can also be modified in the **Create, Analyze, and Learn** dialog box by clicking **Settings**. For more information, see http://www-01.ibm.com/support/knowledgecenter/SSBRAM_8.8.0/com.ibm.classify.workbench.doc/c_WBG_Saving_Learning_Data.htm.

1. Click **Auto Classification > Create Model**.

You must set an attribute when you create a classification model.

2. Enter the information to complete the fields and save it.
3. At this point, the data expert needs to complete the following actions from IBM StoredIQ Data Workbench.
 - Create a user info set.
 - Run the Step-up Snippet action.
 - Apply an enhancement.
 - Apply a filter.
 - View the objects in the Data Object Viewer.
 - Provide feedback for this data object.

For detailed descriptions of these actions, see the topic about retraining a classification model in the IBM StoredIQ Data Workbench documentation.

4. In IBM StoredIQ Administrator, click **Auto-Classification** and select the classification model. The **<model name> Details** window shows that it received feedback. Click **Retrain** to start the retrain process for this model.

If the model was uploaded without a learning archive (SARC file) or if it received no feedback, then the **Retrain** button is disabled.

Next, the data expert runs the enhancement against the user info set again from IBM StoredIQ Data Workbench to see the new scores within the **Data Object Viewer**. The improved scores indicate greater validity and accuracy.

Best practices for creating an Auto-classification model

The following topics address some best practices specific for creating an Auto-classification model.

IBM Content Classifier Workbench

IBM Content Classifier Workbench is an application of the IBM Content Classification product. When you create an Auto-classification model with this application, you must ensure that you follow these practices.

Preparing the sample content sets for training

- The more number of samples per category in the content set you provide, the better it is for training. However, the documents in the content set for a single category in the sample must not be self-contradictory as it can confuse the classifier.
- The best number is 40 - 50 items per category.
- The sample content set must be prepared or reviewed by some subject matter experts.

Creating a Knowledge Base

You must specify at least one field from the content set to be analyzed by the Classification Module. This field must contain meaningful text.

To tell the Classification Module that this field is to be analyzed, you assign a content type value to the field. Do not assign a content type value for the fields that contain non-textual values, such as account numbers or telephone numbers, or non-meaningful text, such as a content field that contains arbitrary administrative comments about each content item.

The number of documents per category

The number of documents per category in the Knowledge Base must be inversely proportional to the number of categories. The fewer categories that you have, the more documents you want to maintain within the Knowledge Base. The default value is 80. To change this value, follow these steps:

- Open an existing Knowledge Base or create a new one. Click **Project Options**.
- Click **Advanced** and edit **Optimally maintain X documents per category**.
- Increase the number as necessary. If you have five categories, for example, this number can be set to a range of 200 to 300. If you have many categories, it is best to keep this number around 100 rather than 200 or 300.

Training your Knowledge Base

Follow these steps to train your Knowledge Base:

- Click **Create, analyze and learn**.
- Use any other option besides **Create using all, analyze using all**.
- Use **Create using all, analyze using all** to assess only the data consistency of the document set contents and to check whether the information in the content can statistically predict the categories. A good document set can yield in this test more than 95% for the top ranking category.
- Use one of the other methods on the list, such as **Create using even, analyze using odd** to have a more realistic prediction estimation after the assessment is done.
- The right option depends on how the content set was prepared and shared by the Subject Matter Experts. Some degree of randomness must be introduced to show how the content set is partitioned for training and analysis. In general, the best option is **Create Using Even, Analyze using Odd**.
- In the next page, enable **Save learning data (SARC File) with Knowledge base** for the model to accept feedback.

Analyzing and reviewing your Knowledge Base by using the built-in reports

- Click **View Reports** on the toolbar to open the **View Reports** window.
- To understand the overall accuracy of your Knowledge Base, view **Knowledge Base Data Sheet**, **Cumulative Success** summary reports, and **Total Precision vs. Recall graph**.

- In **Knowledge Base Data Sheet**, set the first column to report a number greater than 95. It indicates that performance of the Knowledge Base is good during analysis. The less this number is than 95, the more self-contradictory the Knowledge Base becomes. This expectation is related to the data consistency verification when you use **Create using all, analysis using all**, while for the regular creating and analysis case, the number can be lower.
- In the **Total Precision vs. Recall graph**, the curve must be in the upper-right portion. It indicates that performance of the Knowledge Base is good.
- If the curve is on the lower-left portion, it means that performance of the Knowledge Base is poor.

Choosing to reserve items from the training set in the Knowledge Base

- Open an existing Knowledge Base or create a new one.
- Click **Knowledge Base**.
- Right click **Reserve Items**. It ensures that no matter how much feedback you give to this Knowledge Base in the future, the reliable sample data set that is used for training is always retained within the Knowledge Base.

Another way to reserve items is to use the **Freeze** option when you use the BundleDPKB utility to prepare a model and upload it into StoredIQ. An example of adding the **Freeze** option to the BundleDPKB command is: `BundleDPKB "C:\IBM\ContentClassification\Classification Workbench\Projects_Unicode" project_DP Output freeze`

Working with Learning Data (SARC)

In the Knowledge Base Workbench projects, you can set the Knowledge Base to work with the Learning Data method. During the training process, a file that retains the important training content information (SARC) is generated along with the Knowledge Base.

Assessing the data quality in IBM Content Classifier Workbench

Run **Create using all, analysis using all** on the data to verify that it is statistically consistent. Results are expected to be unrealistically high, which is about 95% correct for the top category. If you get less than that it means the data sets of the categories are contradictory.

When you pass that initial test, you must run a real test by training a part of the data and matching the other part. It gives you an idea of the results that you get in deployed systems. You must use either even or odd sets or random cuts.

For more information about how to work with IBM Content Classification, see this Redbook at <http://www.redbooks.ibm.com/redbooks/pdfs/sg247707.pdf>

IBM StoredIQ

When you create an Auto-classification model with IBM StoredIQ, you must ensure that you follow these practices.

- Run a Step-up Snippet enhancement on the infoset before you run an auto-classification enhancement with your model.
- Every time a model is retrained, another auto-classification enhancement must be run on the infoset to compute the classification results of the newly trained model.
- Step-up Snippet needs to be run only once for an infoset.
- After you run a Step-up Snippet, exceptions can be generated from the data servers that participate in the Step-up Snippet operation. Exceptions affect not only the classification, but also any feedback on the documents.
- It is assumed that auto-classification training and feedback work without human error. You must not provide feedback with any doubt about which category the document belongs to.
- Single feedback does not have a significant impact on a model. For example, giving feedback for the `foo.doc` file as Category A does not classify the `foo.doc` file as A after a retrain and subsequent enhancement.

- Repetitive feedback for the same document does not increase impact on the Knowledge Base. If you give feedback for the `foo.doc` file as Category A, after every retrain you keep giving the same feedback, the subsequent feedback has no additional effect on the Knowledge Base.

Cartridges

Cartridges are units of deployment and management for analytic plugin extensions in the IBM StoredIQ application stack. Cartridges enable IBM StoredIQ to detect additional information in documents.

Cartridges are basically compressed files that contain logic like regular expressions, rules for natural language processing, or other code for analyzing document content. When you add a cartridge to IBM StoredIQ application stack, it can detect the new data in documents during indexing and make these new insights searchable. For example, a sensitive pattern cartridge can enable IBM StoredIQ to detect passport numbers, phone numbers, and other IDs.

To apply the analysis logic contained in the cartridge, you must run a Step-up Analytics action that uses the cartridge on an infoset. IBM StoredIQ examines all documents in the infoset, applies the analytics, and then stores the analysis results in the IBM StoredIQ index. For details, see the topics about creating a Step-up Analytics action and about scheduling and running actions in the IBM StoredIQ Data Workbench documentation.

If a cartridge contains language specific rules, make sure the respective languages are enabled. For details, see the topic about the configuration of multi-language settings in the IBM StoredIQ administration information.

Cartridge types

Cartridges (also referred to as *analysis plugins*) can contain analysis logic based on different technologies that range from simple regular expressions to full blown cognitive approaches like natural language processing (NLP).

However, regex-based cartridges should no longer be used starting with IBM StoredIQ 7.6.0.22. Detection logic in your cartridges should now be based on AQL technology and use the Watson SystemT technology like later versions of the sample cartridges provided with IBM StoredIQ.

The cartridges that are provided by IBM for use with IBM StoredIQ can be downloaded from IBM Fix Central as interim fix packs. Check the [list of available fix packs](#) for the latest version of the cartridges. IBM Business Partners and IBM Services can also create cartridges for specific purposes and provide them through their channels.

IBM StoredIQ Cognitive Data Assessment models for document classification can also be made available as cartridges for upload to IBM StoredIQ.

In addition to these cartridges that you must upload to IBM StoredIQ and tie to a Step-up Analytics action manually, one specific governance cartridge and a corresponding Step-up Analytics action might be automatically created in IBM StoredIQ. For details, see the topic about the `GovernanceDataClasses` cartridge in the IBM StoredIQ integration documentation.

Deploying cartridges

Before you can use a cartridge to detect new data in documents, you must upload the cartridge into the IBM StoredIQ application stack and make it available in a Step-up Analytics action.

The first two steps of this procedure do not apply to the system-generated **GovernanceDataClasses** cartridge.

1. Upload or update a cartridge in the Cartridges pane.

See “[Uploading cartridges](#)” on page 71 for instructions to upload a cartridge. After the cartridge is uploaded, the cartridge is validated and information that is pulled from the cartridge is shown in the Details section.

See “Updating cartridges” on page 72 for instructions to update a cartridge. When you update a cartridge, you can skip step 2.

2. Tie a cartridge to a Step-up Analytics action.

Create a new Step-up Analytics action and select the appropriate cartridge. One Step-up Analytics action can use one or more cartridges as well as a cartridge can be included in more than one Step-up Analytics action.

For the **GovernanceDataClasses** cartridge, which is available only if IBM StoredIQ is configured for synchronization with a governance catalog, the **Governance Analytics** action is automatically created. For details, see the topic about the Governance Analytics action in the IBM StoredIQ integration documentation.

After the Step-up Analytics action is created, you can neither add cartridges to nor remove cartridges from the action.

3. From IBM StoredIQ Data Workbench, run the Step-up Analytics action on an infoset.

During this process, the analysis logic that is encapsulated in the cartridge is applied to the documents in the infoset that is selected for the Step-up Analytics. For example, the regex in the cartridge is run against the document content. The results of the analysis are then indexed to be available for quick searches through filters.

Every cartridge for which you run a Step-up Analytics action adds to the full-text index. After a volume is reharvested or the Step-up Full-text action is run, these cartridges are automatically reapplied to ensure that the results of any Step-up Analytics action are still available in the index.

When cartridges are reapplied to an infoset, documents that haven't been modified since the last run are processed only if a cartridge was updated. Otherwise, unmodified documents are skipped during processing.

To search for analysis results from a cartridge, the data expert can create a filter by using the cartridge syntax. Filters for analysis results from the **GovernanceDataClasses** cartridge are automatically created.

Uploading cartridges

Upload a cartridge for use in Step-up Analytics actions.

Download IBM StoredIQ cartridges from IBM Fix Central to your local drive. Cartridges can also be acquired from IBM Services or IBM Business Partners.

To upload a cartridge:

1. Select **Cartridges** from the side navigation bar in the IBM StoredIQ Administrator user interface.

For each cartridge in IBM StoredIQ, the following information is displayed:

Cartridge column name	Description
Name	The display name of the cartridge. You enter that name when you upload a cartridge.
Cartridge name	The cartridge name, which must be unique within the IBM StoredIQ instance. The name is retrieved from the <code>cartridge.properties</code> file in the cartridge.
Supported results	The search filter terms that the cartridge logic makes available as indexed annotations <code>ia:filterterm</code> . This information is retrieved from the cartridge.

<i>Table 28. Cartridges (continued)</i>	
Cartridge column name	Description
Status	The cartridge status, such as Pending, Processing, Failed, and Available.
Creation date	The date that a cartridge is uploaded.

2. Click **Upload Cartridge** in the **Cartridge** pane.
The **Upload Cartridge** window appears.
3. Enter a display name for the cartridge in the **Name** field.
You can change the display name when you update the cartridge.
4. Accept the terms of the license that is contained in the cartridge.
The **Choose File** button is not active until the license is accepted.
5. Click **Save**.
The cartridge validates automatically.
6. Select a cartridge after validation is complete.

The **Details** section is populated with this information:

Details field	Description
Status	The cartridge status, such as Pending, Processing, Failed, and Available
Last Run	The date that the cartridge analytics were last applied by running a Step-up Analytics action
Last Update	The date when the cartridge was last updated
Last error	The error message for a failed upload or update
Version	The version of the cartridge
Description	The cartridge description which is retrieved from the cartridge
Language supported	The document languages that the cartridge can annotate
Total infosets using cartridge	The total number of infosets to which the cartridge logic was applied
Total active step-up analytics executions using cartridge	The number of active Step-ups Analytics executions that use this cartridge
Total step-up analytics actions using cartridge	The number of Step-up Analytics actions that use this cartridge

Updating cartridges

Follow these steps to update a cartridge from an older version.

This procedure applies to all cartridges but the **GovernanceDataClasses** cartridge, which is created only if IBM StoredIQ is configured for integration with a governance catalog. This cartridge is automatically updated whenever relevant information in the governance catalog changes. For details, see the topic about the **GovernanceDataClasses** cartridge in the IBM StoredIQ integration documentation.

1. Select a cartridge from the list.
2. Click **Update** from the top of the window.

The **Update Cartridge** window opens.

3. Optionally, you can change the display name of the cartridge.
4. Make sure the license agreement is accepted.
5. Click **Choose File** to select a new version of the previously uploaded cartridge from the directory of your desktop.

The new cartridge must have the same cartridge name as the cartridge that it is being updated. The cartridge name is given in the `cartridge.properties` file that is contained in the cartridge. However, previously processed documents are reindexed during a rerun of the respective actions only if the cartridge's version number changed. Otherwise, any new cartridge logic is applied only to newly processed documents, which can result in inconsistencies. The cartridge version is also set in the `cartridge.properties` file. For more information about the properties file, see the developerWorks recipe [Developing Regular Expression Cartridges for StoredIQ](#).

6. Click **Save**.

If the cartridge that is provided in the update is invalid or if an unexpected error occurs during the update process, an error message is returned and displayed in the **Last Error** field of the **Details** pane. If an error occurs, the previously uploaded version of the cartridge remains available.

If Step-up Analytics executions are active, the **Update** button is not active. The number of active Step-up Analytics executions is given in the **Total active Step-up Analytics executions using this cartridge** field in the **Details** section.

When cartridges are updated, use your best judgment to decide whether you must rerun the respective Step-up Analytics actions. While you should rerun them when the updates provide new supported results, you might want to hold off on doing so for performance or detection quality improvements.

When you rerun the respective actions, all documents are processed whether or not they were modified since the last run.

Do not forget to create new infosets after you apply the new logic.

Deleting cartridges

Follow these steps to delete a cartridge.

1. Select a cartridge from the list.
2. Click **Delete**.

Note: The delete function is active only for failed cartridges or cartridges that are not in use by a Step-up Analytics action and infoset.

3. Click **Delete** to confirm your action on the Delete confirmation dialog.

IBM StoredIQ Connector operations

The following table provides information about IBM StoredIQ Connector operations, their parameters, and descriptions.

Connector operation	Parameter	Description
cancelExecution	id	ID of execution
createActionCopyPlain	action_name	Name of the action
	description	Description of the action
	targetset_id	Target set ID
	destination_path	Destination path
	copy_sharepoint_userprofile_notes	Boolean value to copy SharePoint user profile notes
	copy_sharepoint_userprofile_libraries	Boolean value to copy SharePoint user profile libraries
	copy_sharepoint_userprofile_posts	Boolean value to copy SharePoint user profile posts
	copy_sharepoint_userprofile_wikis	Boolean value to copy SharePoint user profile wikis
	copy_sharepoint_userprofile_misc	Boolean value to copy SharePoint user profile misc
	compute_hash_value	Boolean value to compute the hash value
	recreate_directory_structure	Boolean value to re-create directory structure
no_auto_harvest	Boolean value to indicate that there is no auto harvest	
createActionCopyRetention	action_name	Name of the action
	description	Description of the action
	targetset_id	Target set ID
	retention_type	Type of retention
	retention_value	Retention value
	retention_tag	Tag of retention
	compute_hash_value	Boolean value to compute hash value
	no_auto_harvest	Boolean value to indicate that there is no auto harvest

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createActionDeleteNodes	action_name	Name of the action
	description	Description of the action
	ignore_accessed_objects	Boolean to ignore accessed objects
createActionDiscoveryExportDat	action_name	Name of the action
	description	Description of the action
	targetset_id	Target set ID
	destination_path	Destination path
	data_object_limit	Data object limit
	copy_sharepoint_userprofile_notes	Boolean value to copy SharePoint user profile notes
	copy_sharepoint_userprofile_libraries	Boolean value to copy SharePoint user profile libraries
	copy_sharepoint_userprofile_posts	Boolean value to copy SharePoint user profile posts
	copy_sharepoint_userprofile_wikis	Boolean value to copy SharePoint user profile wikis
	copy_sharepoint_userprofile_misc	Boolean value to copy SharePoint user profile misc
	unique_document_id_prefix	Unique document ID prefix
	unique_document_id_digits	Unique document ID digits
	unique_document_id_pad_digits	Boolean value to unique document ID pad digits
email_item_disposition	Email item disposition	

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createActionDiscoveryExportDatLight	action_name	Name of the action
	description	Description of the action
	targetset_id	Target set ID
	destination_path	Destination path
	data_object_limit	Data object limit
	copy_sharepoint_userprofile_notes	Boolean value to copy SharePoint user profile notes
	copy_sharepoint_userprofile_libraries	Boolean value to copy SharePoint user profile libraries
	copy_sharepoint_userprofile_posts	Boolean value to copy SharePoint user profile posts
	copy_sharepoint_userprofile_wikis	Boolean value to copy SharePoint user profile wikis
	copy_sharepoint_userprofile_misc	Boolean value to copy SharePoint user profile misc
	unique_document_id_prefix	Unique document ID prefix
	unique_document_id_digits	Unique document ID digits
	email_item_disposition	Email item disposition

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createActionDiscoveryExportEdrxml	action_name	Name of the action
	description	Description of the action
	targetset_id	Target set ID
	destination_path	Destination path
	data_object_limit	Data object limit
	copy_sharepoint_userprofile_notes	Boolean value to copy SharePoint user profile notes
	copy_sharepoint_userprofile_libraries	Boolean value to copy SharePoint user profile libraries
	copy_sharepoint_userprofile_posts	Boolean value to copy SharePoint user profile posts
	copy_sharepoint_userprofile_wikis	Boolean value to copy SharePoint user profile wikis
	copy_sharepoint_userprofile_misc	Boolean value to copy SharePoint user profile misc
	unique_document_id_prefix	Unique document ID prefix
	unique_document_id_digits	Unique document ID digits
	unique_document_id_pad_digits	Boolean value to unique document ID pad digits
	save_text_copy	Boolean value to save the text copy
email_item_disposition	Email item disposition	
createActionExportExceptions	action_name	Name of the action
	description	Description of the action
	export_format	Export format
createActionModifyAttribute	action_name	Name of the action
	description	Description of the action
	attribute	Modification type
createActionMovePlain	action_name	Name of the action
	description	Description of the action
	targetset_id	ID of the target set
	destination_path	Destination path
	recreate_directory_structure	Boolean value to re-create a directory structure

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createActionMoveRetention	action_name	Name of the action
	description	Description of the action
	targetset_id	ID of the target set
	retention_type	Type of the retention
	retention_value	Retention value
createDatasourceBox Before you add a Box data source, see “Box volume configuration notes” on page 26.	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	box_server	Box server IP/hostname (default=api.box.com)
	box_include_users	Box server option to include users in a regular expression
	box_username	Box OAuth user name
	box_password	Box OAuth password
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index
createDatasourceCifs	datasource_name	Name of the data source
	dataserver_id	ID of data server
	datasource_type	Type of the data source
	cifs_server	CIFS server IP or host name
	cifs_username	CIFS server user name
	cifs_password	CIFS server password
	cifs_shared	CIFS server share directory
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceCm8	datasource_name	Name of the data source
	dataserver_id	ID of data server
	cm8_server	CM8 server IP or host name
	cm8_port	CM8 server port
	cm8_username	CM8 server user name
	cm8_password	CM8 server password
	cm8_repository	CM8 server repository
	cm8_dbtype	CM8 server remote DB type
	cm8_remotedb	CM8 server remote DB
	cm8_schema_name	CM8 server schema name
	cm8_connection_string	CM8 server connection string
	cm8_harvest_itemtypes	Item types to be harvested, separated by comma
	cm8_copy_to_itemtype	Copy to the item type: SiqDocument or leave empty
	index_options_container_metadata	Include metadata for contained objects
index_options_full_text_content	Include content tagging and full-text index	
createDatasourceCmis	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	cmis_server	CMIS server IP or host name
	cmis_port	CMIS server port
	cmis_username	CMIS server user name
	cmis_password	CMIS server password
	cmis_use_ssl	Boolean value to indicate ssl or not
	cmis_service	CMIS server service
	cmis_repository	CMIS server repository
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceDocumentum Note: Before you add a Documentum data source, see the Documentum data source setup procedures in the the topics about installing Documentum client jars to the data server and about adding a Documentum server as a data source in the IBM StoredIQ Data Server administration documentation.	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	documentum_docbase	Documentum server document base
	documentum_username	Documentum server user name
	documentum_password	Documentum server password
	documentum_include_versions	Documentum server flag to indicate whether to harvest all document versions
	index_options_container_metadata	Include metadata for contained objects
createDatasourceExchange20002003	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	exchange_server	Exchange server host name
	exchange_username	Exchange server user name
	exchange_password	Exchange server password
	exchange_mailbox_server	Exchange server mailbox server
	exchange_active_directory	Exchange server active directory server
	exchange_use_ssl	Exchange server use SSL flag
	exchange_folder	Exchange server folder
	exchange_virtual_root	Exchange server virtual root
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceExchange2007	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	exchange_server	Exchange server host name
	exchange_username	Exchange server user name
	exchange_password	Exchange server password
	exchange_mailbox_server	Exchange server mailbox server
	exchange_active_directory	Exchange server active directory server
	exchange_use_ssl	Exchange server use SSL flag
	exchange_folder	Exchange server folder
	exchange_virtual_root	Exchange server virtual root
	index_options_container_metadata	Include metadata for contained objects
index_options_full_text_content	Include content tagging and full-text index	
createDatasourceExchange20102013	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	exchange_server	Exchange server host name
	exchange_username	Exchange server user name
	exchange_password	Exchange server password
	exchange_mailbox_server	Exchange server mailbox server
	exchange_active_directory	Exchange server active directory server
	exchange_use_ssl	Exchange server use SSL flag
	exchange_folder	Exchange server folder
	exchange_virtual_root	Exchange server virtual root
	exchange_personal_archive	Exchange server personal archive
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceExchangeonline	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	exchange_server	Exchange server host name
	exchange_username	Exchange server user name
	exchange_password	Exchange server password
	exchange_folder	Exchange server folder
	exchange_virtual_root	Exchange server virtual root
	exchange_personal_archive	Exchange server personal archive
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index
createDatasourceFilenet	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	filenet_server	FileNet server IP or host name
	filenet_port	FileNet server port
	filenet_username	FileNet server user name
	filenet_password	FileNet server password
	filenet_connectionType	FileNet connection type
	filenet_path	FileNet server path
	filenet_stanza	FileNet server stanza
	filenet_object_store	FileNet server object store
	filenet_wheresql	FileNet sql where clause
	filenet_domain	FileNet server domain
	filenet_create_document	FileNet server create document flag
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceIBMConnections	connections_server	IBM Connections server IP or host name
	connections_username	IBM Connections server user name
	connections_password	IBM Connections server password
	connections_classname	IBM Connections class name
	connections_repository	IBM Connections repository
	connections_optionstring	IBM Connections additional options
createDatasourceLivelihood	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	livelihood_server	Livelihood server IP or host name
	livelihood_port	Livelihood server port
	livelihood_username	Livelihood server user name
	livelihood_password	Livelihood server password
	livelihood_database	Livelihood server database
	livelihood_search_slice	Livelihood server search slice
	index_options_container_metadata	Include metadata for contained objects
index_options_full_text_content	Include content tagging and full-text index	
createDatasourceNewsgator	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	newsgator_server	NewsGator server host name
	newsgator_username	NewsGator server user name
	newsgator_password	NewsGator server password
	newsgator_use_ssl	Boolean value to indicate ssl or not
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createDatasourceNfs	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	datasource_type	Type of the data source
	nfs_server	NFS server IP or host name
	nfs_export	NFS server export directory
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index
createDatasourceSharepoint	datasource_name	Name of the data source
	dataserver_id	ID of the data server
	sharepoint_server	SharePoint server host name
	sharepoint_port	SharePoint server port
	sharepoint_username	SharePoint server user name
	sharepoint_password	SharePoint server password
	sharepoint_use_ssl	SharePoint server use SSL flag
	sharepoint_active_directory_server	SharePoint server active directory server
	sharepoint_site_url	SharePoint server site URL
	sharepoint_version_type	SharePoint version
	index_options_container_metadata	Include metadata for contained objects
	index_options_full_text_content	Include content tagging and full-text index
createFilter	filter_name	Name of the filter
	description	Description of the filter
	expression	Expression of the filter
createSystemInfoset	infoset_name	Name of the infoset
	description	Description of the infoset
	datasource_ids	List of data source IDs separated by commas
	access_type	Access type of infoset, either public or private
	user_list	List of user IDs that can be accessed to the infoset
	waitForCompletion	Flag to indicate sync and async

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createTargetset	targetset_name	Name of the target set
	description	Description of the target set
	type	Type of the target set
	datasource_ids	Data source IDs separated by comma
createUserInfosetByOperations	operations	Infoset operations
	infoset_name	Name of the infoset
	description	Description of the infoset
	sources	List of source infoset IDs separated by commas
	overlays	Overlay filter IDs separated by commas
	notify	Notification email
	generate_datamap	Flag to indicate whether a datamap is generated
	waitForCompletion	Flag to indicate sync or async
createUserInfosetFromExceptions	infoset_name	Name of the infoset
	description	Description of the infoset
	infoset_id	ID of the infoset source
	event_ids	Event IDs separated by commas
	categories	Category names that are separated by commas
	overlays	Overlay filter IDs separated by commas
	notify	Notification email
	generate_datamap	Flag to indicate whether a datamap is generated
	waitForCompletion	Flag to indicate sync or async

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
createUserInfosetFromFilter	infoset_name	Name of the infoset
	description	Description of the infoset
	infoset_id	ID of the infoset source
	filter_id	ID of the filter
	overlays	Overlay filter IDs separated by comma
	notify	Notification mail
	generate_datamap	Flag to indicate whether a datamap is generated
	waitForCompletion	Flag to indicate sync or async
executeAction	execution_name	Name of the action
	user_infoset_id	ID of the user infoset
	action_id	ID of the action
	waitForCompletion	Boolean value to indicate sync or async
generateReport	report_name	Name of the report
	user_infoset_id	ID of the user infoset
	action_id	ID of the report
	terms	List of terms that are separated by commas for Term Hit Report only
	filter_ids	List of filter IDs that are separated by commas for Term Hit Report only
	execution_id	ID of execution for Audit Summary Report only
	notify	Notification email
	waitForCompletion	Boolean value to indicate sync or async
harvestDatasource	harvest_name	Name of harvest
	datasource_id	ID of the data source
	harvest_type	Type of harvest
	load_type	Type of load
	schedule	Harvest schedule
	waitForCompletion	Boolean value to indicate sync or async

Table 29. IBM StoredIQ Connector operations (continued)

Connector operation	Parameter	Description
search	resourceType	Type of the resource
	id	ID of the resource
	nameOptions	Name search options
	resource_name	Name of the resource
	status	Status of the resource
	max	Maximum number of the result to be returned
	delete	Flag to indicate whether the resource is to be deleted

Auditing

To enable you to track activities on the IBM StoredIQ application stack, specific events can be recorded.

In addition to the audit records that are written on the data server, audit records for a set of events on the AppStack can be tracked and written to an audit database. Optionally, you can have the audit records also be written to an audit file for easier access. While audit records on the data server provide details about harvests, volume imports, policy history, and log all user or system initiated actions that are performed on the data server, auditing on the application stack lets you track the user and system activities on the AppStack. You can review the records in the audit database or the audit file to find information about who triggered an action, what type of action was performed, and when this happened.

To manage the auditing scope, for example to limit space usage or audit overhead, you can configure whether audit records are written at all and if so, for which operations. Depending on the configuration, an audit record can be written for the following activities:

- Whenever a user logs in to (or attempts to) or logs out from IBM StoredIQ Administrator or IBM StoredIQ Data Workbench.
- All operations that are related to changing a password.
- For IBM StoredIQ Administrator users:
 - Creating, updating, or deleting users
 - Adding, editing, or deleting a volume
 - Triggering a harvest (immediate or scheduled)
 - Updating or deleting a scheduled harvest
 - Creating, editing, or deleting an infoset

An audit record is also written when the harvest is complete.

- For users of the IBM StoredIQ data workbench:
 - Creating, editing, or deleting a filter
 - Creating, editing, or deleting an infoset
 - Viewing the data objects in an infoset
 - Viewing the content of a data object
 - Previewing a filter result
 - Creating or deleting a report

If you enable user action auditing, audit records are by default written only to the audit database. However, you can select to have audit records also written to an audit file (dual logging). In this case, each line in the audit file represents a single audit record in JSON format as described in [“Viewing audit entries in the database”](#) on page 92.

Audit events

Auditing records information about specific types of user activity by storing audit events in a separate database on the AppStack. This database is accessible only to specific audit users. In addition to capturing audit events in the database, the same or a smaller set of audit events can optionally be written to an audit file.

The following events are captured depending on the configured auditing scope. Operation types with the suffix `_REQUEST` are logged at the start of long running operations. The success or failure result of such an operation is written to an audit record for the actual operation type.

Event category	Operation
AUTHENTICATION	<ul style="list-style-type: none"> • LOGIN • LOGOUT
USER	<ul style="list-style-type: none"> • USER_CREATE • USER_UPDATE • USER_PASSWORD_UPDATE • USER_PASSWORD_RESET_REQUEST • USER_PASSWORD_RESET • USER_DELETE
FILTER	<ul style="list-style-type: none"> • FILTER_CREATE • FILTER_UPDATE • FILTER_DELETE
INFOSET	<ul style="list-style-type: none"> • INFOSET_CREATE_REQUEST • INFOSET_CREATE_FINISH • INFOSET_CREATE_OR_UPDATE_FINISH • INFOSET_UPDATE_REQUEST • INFOSET_UPDATE • INFOSET_DELETE_REQUEST • INFOSET_DELETE • INFOSET_VIEW_FILES • INFOSET_VIEW_FILE_CONTENT
VOLUME	<p>Volume operations:</p> <ul style="list-style-type: none"> • VOLUME_CREATE_REQUEST • VOLUME_CREATE • VOLUME_UPDATE_REQUEST • VOLUME_UPDATE • VOLUME_DELETE <p>Harvest operations:</p> <ul style="list-style-type: none"> • VOLUME_HARVEST_REQUEST • VOLUME_HARVEST_UPDATE • VOLUME_HARVEST_FINISHED • VOLUME_HARVEST_DELETE
REPORT	<ul style="list-style-type: none"> • REPORT_CREATE_REQUEST • REPORT_CREATE • REPORT_UPDATE • REPORT_DELETE

Configuring auditing on the AppStack

Configure whether and to which extent audit records are written on the AppStack.

By default, auditing on the AppStack is turned on with the full set of possible operations to be tracked. To manage the amount of audit records being written, you can pick the operations which to you are the most important to track. You can also turn off auditing for those operations completely.

With auditing enabled, audit records can optionally also be written to an audit file.

To change the audit event configuration:

1. Using an SSH tool, log in to the AppStack as `siqadmin` user.
2. Launch the Appstack Audit Event Configuration utility by running the following command: `auditcfg`
3. To turn off auditing, clear the **Enable auditing for application stack operations** checkbox. To limit the scope of auditing, leave the option selected and clear the checkboxes for operations that you do not want to audit.

Depending on your selections, audit records are written as follows:

- **Login and logout:** Whenever a user logs in to (or attempts to) or logs out from IBM StoredIQ Administrator or IBM StoredIQ Data Workbench
- **User management:** Whenever an IBM StoredIQ Administrator user adds, edits, or deletes a user or otherwise changes a user profile. An audit record is also written when a IBM StoredIQ user changes the password or requests a password reset.
- **Volume operations:** Whenever an IBM StoredIQ Administrator user adds, edits, or deletes a volume.
- **Harvest executions:** Whenever an IBM StoredIQ Administrator user triggers an immediate or a scheduled harvest, or updates or deletes a scheduled harvest. An audit record is also written when the harvest is complete.
- **Infoset operations:** Whenever an IBM StoredIQ Administrator user or an IBM StoredIQ data workbench user creates, edits, or deletes an infoset.
- **Filter operations:** Whenever an IBM StoredIQ data workbench user creates, edits, or deletes a filter.
- **Data object listing and viewing:** Whenever an IBM StoredIQ data workbench user views the data objects in an infoset, the content of a specific data object, or previews a filter result.
- **Report generations:** Whenever an IBM StoredIQ data workbench user creates or deletes a report.

If auditing is enabled, you can select to have an audit file written in addition to the database records. For more information about the default settings and customization options for the audit file, see [“Configuration settings for the audit file” on page 91](#).

4. Select **Restart appstack services**.

Restarting the services is required for any configuration changes to take effect. Selecting this option is an alternative to restarting the application-stack services from the command line.

5. Select **Save and Exit** and press Enter.

6. Optional: If writing to audit file is enabled and you customized any audit file settings, you must also restart the tomcat service.

Run the following command as `siqadmin` user:

```
systemctl restart appstack-tomcat.service
```

Audit records are written as configured. At any time, you can broaden or further limit the auditing scope.

Configuration settings for the audit file

You can have audit records written to file without further configuration. In this case, the default settings are used.

The audit file is a rolling log file. A new file is written whenever the configured time interval passes. The old audit file is compressed and renamed following the convention *file_name-time_pattern.log.gz*. Each line in the audit file represents an audit record in JSON format. For more information, see the description of the [audit record format](#).

To customize the configuration, provide a configuration file in JSON format with your settings. The file must be named `audit-file-service-config.user.json` and must be stored in the `siq/env/tomcat` directory on the AppStack.

For any configuration changes to take effect, you must restart the tomcat service. Log in to the AppStack as `siqadmin` user and run the following command: **systemctl restart appstack-tomcat.service**

Option	Default value	Description
<code>app_log_level</code>	WARN	Controls the level of audit file service messages that are written to the <code>tomcat.log</code> file. To have messages logged to the <code>tomcat.log</code> file, set the level to DEBUG. Then, a message is logged for each audit event that is written to the audit file.
<code>operation_types</code>	All available events as configured in the audit event configuration	Can be used to further limit the set of event operations types to be written to the audit file. Specify the operation types as an array in this format: <code>["event_type", "event_type"]</code> If you do not specify this option at all, all event types that are selected in the audit event configuration are tracked. If you specify this option with an empty array as value, nothing is logged.
<code>file_dir_name</code>	No value is initially set, which means that the file is written to the working directory, in this case, the <code>/siq/env/tomcat</code> directory.	Specifies the directory to which the audit file is written.
<code>file_name</code>	<code>storediq-audit</code>	Specifies the file name of the audit file. The file extension is always <code>log</code> . Thus, the default name of the audit file is <code>storediq-audit.log</code> .

Option	Default value	Description
file_pattern	{yyyy-MM-dd}	<p>Sets the time pattern for the time unit used for audit file rollover. With the default setting, files are rolled over at intervals defined in days. The size of the interval is determined by the rollover_interval option.</p> <p>Possible pattern are:</p> <p>{yyyy} The time unit is year.</p> <p>{yyyy-MM} The time unit is month.</p> <p>{yyyy-MM-dd} The time unit is day.</p> <p>{yyyy-MM-dd-HH} The time unit is hour.</p> <p>{yyyy-MM-dd-HH-mm} The time unit is minute.</p> <p>This pattern also controls the naming of old log files: <i>file_name-time_pattern.log.gz</i>. For example, with the pattern {yyyy-MM-dd-HH}, an old file could be named storediq-audit-2018-11-30-06.log.gz.</p>
rollover_interval	1	Defines after the number of time units after which the audit file is rolled over. With the default configuration, a new file is written every day.

Example:

```
{
  "app_log_level": "WARN",
  "operation_types": [
    "LOGIN",
    "LOGOUT",
    "USER_CREATE",
    "USER_PASSWORD_UPDATE",
    "USER_UPDATE",
    "USER_DELETE",
    "USER_PASSWORD_RESET",
    "USER_PASSWORD_RESET_REQUEST"
  ],
  "file_dir_name": "",
  "file_name": "storediq-audit",
  "file_pattern": "{yyyy-MM-dd}",
  "rollover_interval": "1"
}
```

Viewing audit entries in the database

To view the data in the audit database, you must use SQL SELECT statement queries.

For remote access to the audit database, port 5432 must be opened. Port 5432 is the port for communication with the PostgreSQL instance. Thus, opening this port allows for communication with the entire instance in general, which means with all database tables. The `audituser` account should be the only account that has access from the outside and is limited to accessing the "audit_events" view. However, you should use caution and close the port again after retrieving the audit records because any change to the `hba_conf` file can impact the access rights to the databases and tables.

To open or close the port, run the `postgres_port_open` or `postgres_port_close` script as `siqadmin` user.

You cannot access the audit database directly. Instead, access audit events by means of SQL `SELECT` statements on the "audit_events" database view. You must run such queries with the `audituser` account.

The default password for the `audituser` account is `Passw0rd!`. For security purposes, you should change the password. You can do so at any time by running the `change_audituser_password` script as `siqadmin` user on the AppStack. You are asked to enter and confirm the new password. Your input is not visible on the screen.

1. Connect to the audit database where the database name is "audit" and the user must be `audituser`. You can use a tool such as `pgAdmin`.
2. To retrieve audit records, submit SQL `SELECT` queries similar to the following examples:

```
select * from audit_events;
select * from audit_events where actor_name = 'johndoe';
select * from audit_events where event_date > '2018-06-20' order by id desc;
```

Each record provides the following information:

Field name	Description
<code>id</code>	The unique ID of the audit record.
<code>origin</code>	The host name of the server where the operation causing an audit request was triggered.
<code>correlation_id</code>	A unique identifier that links related audit records, for example, in a long running process where the trigger of the operation is audited as well as the completion of the operation.
<code>event_date</code>	The date and time of the event. The time stamp is in UTC format and includes the time zone, for example, <code>2018-06-20 11:33:09+00</code> .
<code>event_category</code>	The category of the audit event, which can be <code>AUTHENTICATION</code> , <code>VOLUME</code> , or <code>INFOSET</code> .
<code>operation_type</code>	The specific event operation type within an event category.
<code>actor</code>	A static JSON structure containing actor information like the actor's ID, first and last name, the email address and the IBM StoredIQ user name. For audit events that are not triggered by a user but internally, for example by a scheduled task, the values are set as follows: <ul style="list-style-type: none"> • The <code>name</code> attribute is set to <code>#{SYSTEM#}</code>. • The <code>first_name</code> attribute is set to <code>automated-task</code>. • The <code>last_name</code> attribute reflects the actual task that triggered the event.
<code>actor_name</code>	The IBM StoredIQ user name of the user who triggered the operation being audited or <code>#{SYSTEM#}</code> for internally triggered events.
<code>result</code>	The status of the audited operation, which can be <code>NONE</code> , <code>SUCCESS</code> , or <code>FAILURE</code> .

Field name	Description
reason	The stated business justification. This field is currently not used.
event_details	A JSON object containing key value pairs that are specific to the event_category. The structure of the JSON is either static for the event_type or the event_category.
event_payload	The event payload as JSON with details about the specific event.

Administering IBM StoredIQ components

For a secure IBM StoredIQ environment, a specific administration framework allows for controlling and maintaining the IBM StoredIQ AppStack server without direct root access.

The system administrator account

Processes running under the `root` account can be vulnerable. To restrict root access and limit the processes and services that run under the root account, IBM StoredIQ provides a separate administrative account for managing the AppStack server.

The `siqadmin` user account is available in the IBM StoredIQ application stack starting with version 7.6.0.17. Use the `siqadmin` account for administrative tasks on the AppStack, such as starting, stopping, and maintaining the AppStack server.

In new deployments, the `siqadmin` user account is created during installation. You set the password during the configuration of the AppStack. You can change it any time later by using the Linux `passwd` command.

The `siqadmin` user is a regular user with login capabilities. This account neither has root privileges nor is it a member of the root group. This user's designated home directory is the `/home/siqadmin` directory. All scripts that the `siqadmin` user can run to manage the application stack are stored in the `/home/siqadmin/bin` directory. The `siqadmin` user is also allowed to manage a defined set of services.

As `siqadmin` user, you can create and store your own command files in the `/home/siqadmin/bin` directory without interfering with the IBM StoredIQ provided command setup.

Administration commands

All administration commands for use by the `siqadmin` user are stored in the `/home/siqadmin/bin/` directory.

These administration commands serve as resources that allow the `siqadmin` user to run specific commands or scripts that would otherwise require root authority. These general administration commands are currently available:

- appstackcfg** - AppStack configuration utility.
- auditcfg** - utility for configuring AppStack auditing.
- bake** - internally used.
- bastioncfg** - internally used.
- bootstrap** - command-line interface for provisioning, updating, and testing the server.
- certcfg** - certificate configuration utility.
- change_atlas_password** - utility for changing the `atlas` password.
- change_audituser_password** - utility for changing the `audituser` password.
- change_cda_db_password** - utility for changing the password for the CDA service database. Run this command at least once after enabling the CDA service to overwrite the default password (`Passw0rd!`).
- change_reportuser_password** - utility for changing the `reportuser` password.
- change_superuser_password** - utility for changing the `superadmin` password.
- create_cda_service_user** - internally used.
- dbcheck** - internally used.
- fipscfg** - FIPS configuration utility.
- force_sync** - command to trigger an immediate synchronization with the governance catalog.
- ldapcfg** - LDAP configuration utility.

netcfg - Network configuration utility.

pg_dumpall - internally used.

postgres_port_close - command to close the port for communication with the PostgreSQL instance.

postgres_port_open - command to close the port for communication with the PostgreSQL instance.

postupgrade - command for applying security updates as part of upgrading the AppStack.

psql - interactive interface for working with Postgres.

smtpcfg - command-line interface for setting up the SMTP account. For UI based changes, use the **appstackcfg** utility.

storediqapi - enable or disable the IBM StoredIQ APIs.

sync_ldap - command to trigger an immediate synchronization of user details in the IBM StoredIQ database with the directory server.

systemctl - command to manage services, for example, to start and stop the services or to query the status.

The set of services that can be managed with the **siqadmin** account are defined in the `/siqsec/siqadmin.syslst` file. The default set consists of the **appstack** and **vault** services.

upgrade_appstack - internally used.

upload_report - utility for uploading custom report designs.

vault_check - utility for identifying and eventually fixing issues with the vault service.

vault_status - command to display the vault status.

Important: Run these commands as shown. Do not try to run them from the `/siq/bin` directory; this will result in an error. For example, instead of running `/siq/bin/appstackcfg`, simply run **appstackcfg**.

Additional administration commands

The following set of commands is available in addition to the general administration commands.

All edit commands listed here run the Linux **sudoedit** command to edit the respective system files. For detailed information about the **sudoedit** command, see the Linux man pages.

edit_etc_hosts

Edits the `/etc/hosts` system file.

edit_etc_resolv_conf

Edits the `/etc/resolv.conf` system file.

edit_etc_siq_gateway_yaml

Edits the `/etc/siq/gateway.yaml` system file.

edit_etc_siq_local_yaml

Edits the `/etc/siq/local.yaml` system file.

edit_etc_siq_oauth_conf

Edits the `/etc/siq/oauth.conf` system file.

edit_etc_siq_proxy_conf

Edits the `/etc/siq/proxy.conf` system file.

edit_etc_siq_report_conf

Edits the `/etc/siq/report.conf` system file.

siq_files_access

Changes the group ownership for the IBM StoredIQ log files and other IBM StoredIQ related files and directories to **siqadmin** and sets the permissions for these files and directories to read/write for the group.

By default, the group ownership is **siqadmin** for this set of files and directories:

```
/siq/env/tomcat
/siq/env/websphere-liberty/usr/servers/ediscovery_connector/etc/
atlas_setup.properties
/var/bootstrap
/var/siq/download
```


`/var/siq/log`

Because the `siqadmin` user is a member of the `siqadmin` group, this user is granted read and write access to these files.

By default, this command is run each time you log in as `siqadmin` user to make sure the appropriate permissions are set for any new files that were created in the listed directories and that do not have the required access permissions for group `siqadmin` set initially. However, you can run this command manually at any time if required.

If the `siqadmin` user needs access to additional files, the default set of files and directories can be extended by the root administrator. For more information, see [“Administration files”](#) on page 97.

vault_check

Runs the vault check tool (`/siqsec/bin/vaultCheck.sh`) for troubleshooting the vault service.

The tool helps you identify and eventually fix issues with the vault service. It checks whether the vault is functional and its data is consistent; it provides a detailed summary including a list of the detected problems and provides instructions for fixing those problems. The output of the tool can be used for debugging and analysis and also contributes to the information that you must gather before contacting IBM Support.

vault_status

Runs the vault status tool (`/siqsec/bin/vaultStatus.sh`)

Extending the set of administration commands

As `siqadmin` user, you can create and store your own command files in the `/home/siqadmin/bin` directory without interfering with the IBM StoredIQ provided command setup.

Administration files

The scope of the administration commands and system files that the `siqadmin` user has access to is defined in several list files.

These files are stored in the `siqsec` directory. Usually, you don't need to touch these list files because, by default, they are set up to cover all commands and files that the `siqadmin` user needs access to. Only if this user requires further permissions to run additional commands, to edit additional system files, or to access additional files or directories, the root user must extend these list files:

List of commands: `/siqsec/siqadmin.cmdlst`

This file lists the file names of those executable programs on the AppStack that the `siqadmin` user is allowed to run.

For each listed file name, a command with the respective name must exist in either the `/siq/bin` directory or the directory given in the path specification. Empty lines in the file and lines starting with `#` are ignored.

For each listed file name, a command file referencing the respective command in either the `/siq/bin` directory or the directory given in the path specification is created in the `/home/siqadmin/bin` directory, initially during installation and triggered by a command any time after installation. The administration commands in this directory serve as resources that allow the `siqadmin` user to run specific commands or scripts that would otherwise require root authority.

Only the root user can edit the `/siqsec/siqadmin.cmdlst` file and must run the `/siqsec/bin/vaultSetup.sh` script after updating this file to reflect the changes in the `/home/siqadmin/bin` directory.

List of editable files: `/siqsec/siqadmin.editlst`

This file lists the file names of those system files on the AppStack that the `siqadmin` user is allowed to edit by using designated edit commands. See the list of [“Additional administration commands”](#) on

page 96. These commands run the Linux **sudoedit** command to edit the respective system files. For detailed information about the **sudoedit** command, see the Linux man pages.

For each listed file name, a command file named in the format `edit_directory_file` is created in the `/home/siqadmin/bin` directory, initially during installation and triggered by a command any time after installation. Empty lines in the file and lines starting with `#` are ignored.

Only the root user can edit the `/siqsec/siqadmin.editlst` file and must run the `/siqsec/bin/vaultSetup.sh` script after updating this file to reflect the changes in the `/home/siqadmin/bin` directory.

List of services: `/siqsec/siqadmin.syslst`

This file lists the services on the AppStack that the `siqadmin` user is allowed to manage. The default set consists of the `appstack` and `vault` services. Any services not listed in the `/siqsec/siqadmin.svclst` file must be managed by using the root account.

Only the root user can edit the `/siqsec/siqadmin.syslst` file.

Group access list: `/siqsec/siqadmin.grplst`

This file lists the names of files or directories on the AppStack for which access permission and group ownership is set such that the `siqadmin` user and other members of the `siqadmin` group can access and modify these files.

For each listed file or directory, ownership and privileges are initially set during installation and by running a command at any time after installation.

If the `siqadmin` user or other members of the `siqadmin` group require read/write access to further files or directories, these files and directories can be added to the `/siqsec/siqadmin.grplst` file at any time. However, only the root administrator can edit the file. After the root user updates the list, ownership and permission settings must also be updated, which can be done in either of these ways:

- The root user runs the `/siqsec/bin/vaultSetup.sh` script.
- The `siqadmin` user runs the `siq_files_access` script.

Notices

This information was developed for products and services offered in the U.S.A. This material may be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing
IBM Corporation
North Castle Drive, MD-NC119
Armonk, NY 10504-1785
US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp. Sample Programs.

© Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" <http://www.ibm.com/legal/copytrade.shtml>.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Microsoft and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

VMware, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

Red Hat and OpenShift are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek

your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Index

A

- access restrictions
 - data [41](#)
 - infosets [41](#)
 - users [41](#)
 - volumes [41](#)
- action
 - clone [52](#)
 - delete [52](#)
 - edit [52](#)
- action types
 - copy [49](#)
 - copy to retention [49](#)
 - delete [49](#)
 - discovery export DAT [49](#)
 - discovery export DAT Light [49](#)
 - discovery export EDRM XML [49](#)
 - modify attribute [49](#)
 - move [49](#)
 - Step-up Full-Text [49](#)
 - Step-up Snippet [49](#)
- actions
 - create [49](#)
 - types [48](#)
- Active Directory [26](#)
- active user status [43](#)
- administration
 - application stack [95](#)
 - AppStack [95](#), [97](#)
 - commands [95](#), [97](#)
 - components [95](#)
 - root access [95](#)
 - siqadmin [95](#), [97](#)
 - system files [95](#), [97](#)
- All Data Servers tab [4](#)
- All Volumes tab [4](#)
- AppStack
 - administration [95](#), [97](#)
- audit
 - records [92](#)
 - reviewing [92](#)
- audit file
 - configuration settings [91](#)
- auditing
 - AppStack [88](#), [90](#)
 - configuring [88](#), [90](#)
 - events [88](#)
 - filter operations [88](#)
 - harvest operations [88](#)
 - infoset operations [88](#)
 - report generations [88](#)
 - scope [90](#)
 - user actions [88](#)
 - volume operations [88](#)
- Auto-classification model [67](#)
- auto-classification

- auto-classification (*continued*)
 - use [62](#)
- auto-classification feature [62](#)
- auto-classification model
 - building [63](#)
 - classification model
 - details [64](#)
 - create [64](#)
 - import [64](#)
 - select [64](#)
- Auto-classification model [67](#), [68](#)
- automated document categorization [61](#)

B

- best practices [67](#), [68](#)
- BIRT
 - connection profile [57](#)
 - JDBC driver [57](#)
 - report design [57](#)
- Box
 - volume configuration [26](#)

C

- cartridge
 - types [70](#)
- CIFS [7](#), [26](#)
- classification export utility
 - download [63](#)
- classification model
 - search [65](#)
- classification models [66](#)
- clone
 - action [52](#)
- CMIS [7](#)
- commands
 - administration [95](#), [97](#)
- Compliance Report [55](#)
- Concordance DAT [48](#)
- Concordance DAT Light [48](#)
- configuration settings
 - audit file [91](#)
- connection profile
 - BIRT [57](#)
- Content Collector Manifest File CSV Export [55](#)
- copy action [49](#)
- Copy action [48](#)
- copy to retention action [49](#)
- Copy to Retention action [48](#)
- create a Step-up Analytics action [51](#)
- CSV All Audited Object Exports [55](#)
- CSV Attribute Summary [55](#)
- CSV Exception List Export [55](#)
- CSV Infoset Data Object Export [55](#)
- CSV Term Hit Details Export [55](#)
- custom reports [57](#)

D

- Data Assessment Report [55](#)
- Data Dashboard [4](#)
- Data Privacy Report [55](#)
- data server
 - configuration guidelines [62](#)
 - rebooting [35](#)
 - restarting [35](#)
 - sizing guidelines [62](#)
 - viewing details of [35](#)
- data servers
 - managing [4](#)
- data sources
 - limiting access [41](#)
- Data Topology Report [55](#)
- DataServer - Classic [6](#)
- DataServer - Distributed [6](#)
- delete
 - action [52](#)
- delete action [49](#)
- Delete action [48](#)
- delete cartridges [73](#)
- discovery export DAT [49](#)
- Discovery Export DAT action [48](#)
- discovery export EDRM XML action [49](#)
- Discovery Export EDRM XML action [48](#)
- discovery export volume
 - adding [33](#)
- Duplication Summary Report [55](#)

E

- edit
 - action [52](#)
- Elasticsearch [6](#)
- exceptions [1](#), [38](#)
- Exchange [7](#)

F

- feedback
 - providing [66](#)
- FileNet [7](#)
- fonts
 - deploy [60](#)
- full harvest [37](#)

G

- General Data Protection Regulation (GDPR) [70](#)

H

- harvest
 - details [38](#)
 - review [38](#)
- harvests
 - delete [38](#)
 - edit [38](#)
 - searching for a specific harvest [38](#)
- high_sensitive [55](#)

I

- IBM Content Classification [61](#)
- IBM Content Classification model
 - prepare for import [63](#)
- IBM Content Classification Workbench [67](#)
- IBM Content Manager [7](#)
- inactive user status [43](#)
- incremental harvest [37](#)
- infoset
 - private [40](#)
 - public [40](#)
 - system [40](#)
- infosets
 - limiting access [41](#)
 - private [41](#)
 - public [41](#)

J

- JDBC driver
 - BIRT [57](#)

K

- Knowledge Base [67](#)

L

- LDAP
 - application stack [45](#)
 - server [45](#)
 - synchronizing [45](#)
- legal
 - notices [99](#)
- List of Harvests page [37](#)
- Livelink [7](#)
- low_sensitive [55](#)

M

- medium_sensitive [55](#)
- model
 - delete [66](#)
 - edit [65](#)
 - retrain [66](#)
- model accuracy [66](#)
- modify attribute action [49](#)
- move action [49](#)
- Move action [48](#)

N

- NewsGator [7](#)
- NFS [7](#)
- node ID [55](#)
- notices
 - legal [99](#)

O

- OneDrive [7](#)

- OneDrive configuration [30](#)
- OneDrive for Business [30](#)
- OneDrive volume [30](#)
- OneDrive volume prerequisites [30](#)
- Overlay Hit Report [55](#)

P

- password
 - change [46](#)
 - local user [46](#)
- prerequisites
 - Exchange Online [29](#)
- primary volume [7](#)

R

- report
 - create [56](#)
 - data set [57](#)
 - data source [57](#)
 - design [57](#)
 - parameter [57](#)
 - reportuser [57](#)
- report configuration file [55](#)
- reports
 - Compliance Report [55](#)
 - Content Collector Manifest File CSV Export [55](#)
 - CSV All Audited Object Exports [55](#)
 - CSV Attribute Summary [55](#)
 - CSV Exception List Export [55](#)
 - CSV InfoSet Data Object Export [55](#)
 - Data Assessment Report [55](#)
 - Data Topology Report [55](#)
 - Duplication Summary Report [55](#)
 - Overlay Hit Report [55](#)
 - Term Hit Report [55](#)
- reportuser [57](#)
- retention volume
 - adding [32](#)
- roles
 - user [43](#)
- root access [95](#)

S

- SharePoint [7](#), [26](#)
- single sign-on [26](#)
- sqadmin [95](#), [97](#)
- Step-up Analytics [70](#)
- Step-up Analytics action [48](#)
- Step-up Full-Text action [48](#)
- Step-up Snippet action [48](#)
- stopping a harvest [38](#)
- StoredIQ Connector [74](#)
- StoredIQ Connector operation [74](#)
- System Activity [4](#)
- system files
 - administration [97](#)
- system infosets
 - create [40](#)
- System Status [4](#)

T

- target set
 - cloning [53](#)
 - creating [53](#)
 - deleting [53](#)
 - editing [53](#)
- target sets [53](#)
- Term Hit Report [55](#)
- train [67](#)

U

- update cartridges [72](#)
- uploading cartridges [71](#)
- use_node_id [55](#)
- user [43](#)
- user account
 - manage [43](#)
- user password [46](#)
- user roles [43](#)
- user status [43](#)
- using cartridges [70](#)

V

- volume
 - adding a primary [7](#)
 - discovery export [33](#)
 - removing [36](#)
 - retention [32](#)
 - viewing details of [35](#)
- volumes
 - limiting access [41](#)
 - managing [4](#)

