

Administering Veritas™ Information Studio

Information Studio 1.1

Information Studio 1.1

Documentation version: .1

PN:

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-party Programs"). Some of the Third-party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the Third-party Legal Notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. VERITAS TECHNOLOGIES LLC SHALL NOT BE LIABLE FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES IN CONNECTION WITH THE FURNISHING, PERFORMANCE, OR USE OF THIS DOCUMENTATION. THE INFORMATION CONTAINED IN THIS DOCUMENTATION IS SUBJECT TO CHANGE WITHOUT NOTICE.

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive.
Santa Clara, CA 95054

<http://www.veritas.com>

.

Technical Support

Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

www.veritas.com/support

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

Contacting Technical Support

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

www.veritas.com/support

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information

- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Technical Support
 - Recent software configuration changes and network changes

Licensing and registration

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

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

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

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Advice about technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs, DVDs, or manuals

Support agreement resources

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Contents

Technical Support	4
Chapter 1 Interacting with the Information Studio administration console	11
Overview	11
Logging into Veritas™ Information Studio	12
Veritas™ Information Studio administration console	14
Chapter 2 Application Switcher	16
About the application switcher	16
Chapter 3 Organizations	17
About Information Studio organizations	17
Role privileges of Veritas™ Information Studio users	18
Editing a customer account	19
Granting access to Information Studio users	21
Creating a custom role	21
Veritas™ Information Studio tenant management	22
Certificate management	23
Configuring SMTP (Simple Mail Transfer Protocol) services	30
Chapter 4 Licensing	31
Overview	31
Trial license model	32
Subscription license model	33
Perpetual license model	34
Chapter 5 Dashboard	36
About the dashboard	36
Chapter 6 Applications	39
About Information Studio application	39

Chapter 7	Data Engine	40
	About Information Studio data engine	40
Chapter 8	Connectors	41
	Overview of Connectors	41
	Supported connectors	42
	Requirements and prerequisites	43
	High-level workflow for setting up connections to content sources	
	from the Connections console	44
	Known limitations of Connections console	44
	Managing credentials	46
	Adding connections	47
	Managing connections	49
	Content sources and hierarchies	50
	Changing a pause schedule in Connections	51
	Assigning policies to Connectors	51
	Configuring Connections to cloud content sources	52
	Configuring data collection from Microsoft Azure	52
	Configuring Microsoft SharePoint Online connection	63
	Configuring data collection from Enterprise Box	66
	Configuring data collection from Google Cloud Storage	68
	Configuring Microsoft Exchange Online connection	71
	Configuring a Google Drive and Gmail Connection	74
	Configuring Microsoft OneDrive connection	77
	Configuring a Generic S3 Cloud Storage connection	80
	Configuring Connections to on-premises content sources	82
	Configuring data collection from Native File Server	82
	Configuring data collection from Microsoft SharePoint on-premises	
	97
	Configuring a Microsoft SQL Server connection	100
	Configuring an Oracle connection	102
	Configuring a NetBackup connection	104
	About OpenText Documentum connection	107
	About IBM FileNet connector	109
	About OpenText LiveLink connector	110
	Managing connections to Microsoft Exchange on-premises	
	connector	112
Chapter 9	Information Studio Policies	114
	Overview of Information Studio policies	114
	Creating a new policy	115

	Veritas™ Information Studio policy rules	117
	Managing policies	119
	Viewing policies	121
Chapter 10	Locations	124
	About locations in Veritas™ Information Studio	124
Chapter 11	Logs	127
	About logs	127
Chapter 12	Audit Logs	128
	About audit logs	128
Chapter 13	Monitoring Health	132
	About monitoring health in Veritas™ Information Studio	132
Chapter 14	Monitoring Jobs	136
	About Jobs	136
	Viewing jobs	136
Chapter 15	User Preferences	141
	About user preferences	141
Chapter 16	Troubleshooting Veritas™ Information Studio	143
	About troubleshooting	143
Chapter 17	Information Studio CLI	146
	Using the Information Studio CLI	146
Chapter 18	Back-up and restore Information Studio	149
	Backing up and restoring Information Studio	149
Appendix A	Reporting API	153
	Overview of reporting API	153
	Elasticsearch schema	161

Appendix B	Getting help	167
	Displaying the online help	167
	Using the Veritas™ Information Studio product documentation	167

Interacting with the Information Studio administration console

This chapter includes the following topics:

- [Overview](#)
- [Logging into Veritas™ Information Studio](#)
- [Veritas™ Information Studio administration console](#)

Overview

This section describes the user interface of the Veritas™ Information Studio administration console. The target audience of this guide includes administrator users from customer organizations.

Each Veritas™ Information Studio deployment comes with a default administration console component. The Information Studio administration console helps with the following:

- Managing Information Studio and setting up access for users
- Establishing policies for discovery and scanning
- Managing connections to enterprise content data assets residing on-premise

Web browsers supported by Veritas™ Information Studio

Veritas™ Information Studio supports the following web browsers for Windows and iOS users:

- Chrome version 72 or later
- Firefox version 65 or later

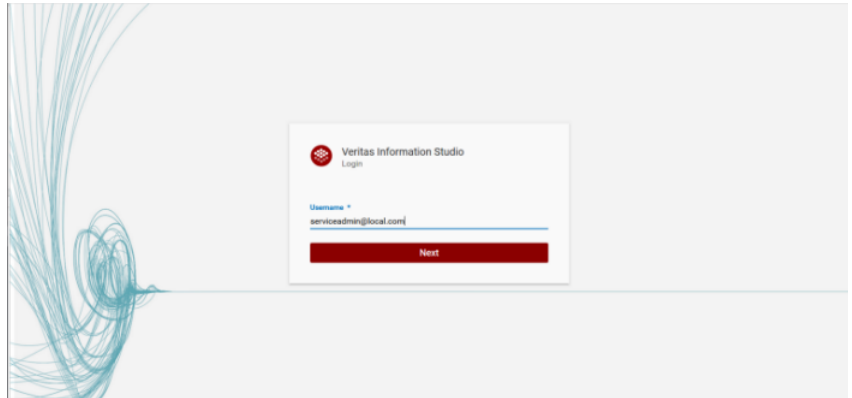
Logging into Veritas™ Information Studio

After you have deployed Veritas™ Information Studio, you can launch the UI in one of the recommended browsers using `https://<IP Address>`. Logging-in is a three-step process in Information Studio, and you can follow the procedure below to log into the Information Studio administration console.

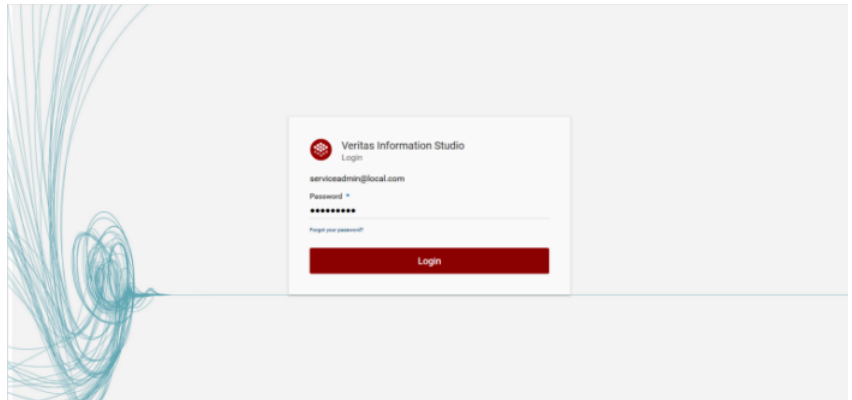
Note: If this is the first time you are logging into Information Studio, you need to log in with the credentials of the Customer Super Administrator account (Default_Customer). For more information on Information Studio roles and their access privileges, you can see the section on granting access to Information Studio users. See [“Granting access to Information Studio users”](#) on page 21.

To log into Information Studio administration console

- 1 Launch the UI using `https://<IP Address>`, enter the **Username**, **serviceadmin@local.com**, and click **Next**.

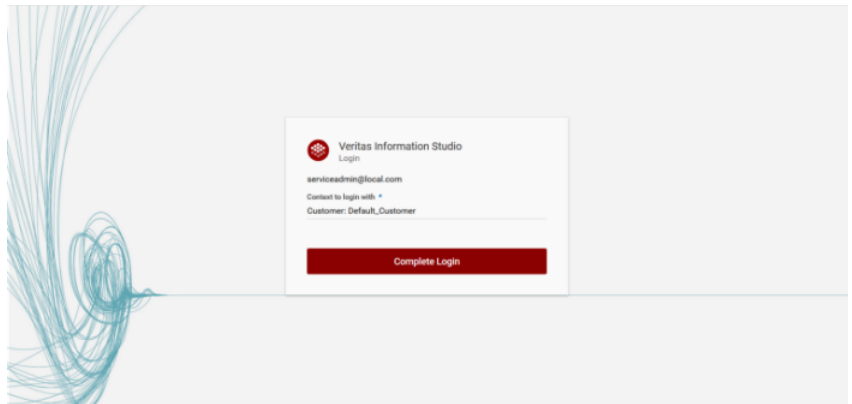


- 2 Enter the **Password** set during deployment, and click **Login**.



- 3 Complete the login process by setting the context with which you want to log in to Information Studio.

A role (for example, a Customer Admin, IT Admin and so on) defines permissions for a user. Each role that signs into Veritas™ Information Studio has the privilege to perform certain tasks. A context refers to the customer or tenant under which the role permissions are given. A context and a role are not the same. A user can have the same role across different contexts.



Veritas™ Information Studio administration console

After providing valid credentials, you can see the Information Studio administration console, which is divided into the following sections:

Title bar

The title bar appears at the top of the screen. Besides the product logo and name, the title bar provides a link to help topics and sign-out options. Clicking on the user icon in the title bar lets you change or update the user password.

Navigation bar

The navigation bar appears on the left-hand side of the Information Studio administration console. Depending upon the user privileges and role or persona, the items shown in the navigation bar may vary. For example, the navigation bar in the figure below is seen when a user signs with the role of a Tenant Admin.

Table 1-1 Navigation bar options in the Information Studio administration console

Area	Description
Dashboard	Monitor, manage, and gain insights from the various components of Veritas™ Information Studio administration console. See “About the dashboard” on page 36.
Applications	Access applications deployed on Veritas™ Information Studio. See “About Information Studio application” on page 39.
Organizations	View and manage Veritas™ Information Studio customers, users, and certificates. See “About Information Studio organizations” on page 17. Using the Organizations > Customers option, you can edit customer accounts and add Active Directory (AD) - related information.
Policy Manager	Create enterprise policies for various capabilities. See “Overview of Information Studio policies” on page 114.
Data Engines	Deploy on-premises Data Engine to monitor your on-premises content sources. See “About Information Studio data engine” on page 40.
Connectors	Configure credentials and add connections to content sources that you want to monitor. See “Overview of Connectors” on page 41.
Audit Logs	View, audit, and analyze Information Studio events. See “About audit logs” on page 128.
Monitoring	View the details and status of scheduled jobs. See “About Jobs” on page 136.
User Preferences	Configure user preferences for Information Studio and other applications. See “About user preferences” on page 141.

Main display

On the right-hand side of the navigation bar, you can see the main display area. Depending on the user privilege and the context set on the navigation bar, it displays varied information and the results of user actions.

Application Switcher

This chapter includes the following topics:

- [About the application switcher](#)

About the application switcher

Use the Application Switcher to switch between various screens on the Information Studio administration console and the Information Studio application console.

Click in the top-left corner of the Information Studio administration console to open and collapse the Application Switcher.

For information about the Information Studio user interface:

- See [“Veritas™ Information Studio administration console”](#) on page 14.
- See [“About Information Studio application”](#) on page 39.

Refer to the *Veritas™ Information Studio User Guide* for information on how to interact with the applications.

Organizations

This chapter includes the following topics:

- [About Information Studio organizations](#)
- [Role privileges of Veritas™ Information Studio users](#)
- [Editing a customer account](#)
- [Granting access to Information Studio users](#)
- [Creating a custom role](#)
- [Veritas™ Information Studio tenant management](#)
- [Certificate management](#)
- [Configuring SMTP \(Simple Mail Transfer Protocol\) services](#)

About Information Studio organizations

Organizations refers to enterprises or the customers that want to subscribe to Veritas™ Information Studio. The **Organizations** link is visible to the Customer Super Admin, Customer Admin, and Tenant Admin roles. The **Organizations** link is not visible to users with IT Admin and Tenant User roles.

About the Information Studio user roles

Roles specify access privileges to Information Studio and other applications. Roles control who has access to your data and the level of access they have. Information Studio allows access to users only on a need-to-know basis. For example, the Customer Super Admin cannot view or manage tenant-level users. They can only look at Customer-level users.

Information Studio provides the following user roles:

- Customer Super Admin
- Customer Admin
- Tenant Admin
- IT Admin
- Tenant User

A user cannot change the permissions of the above-mentioned user roles. However, permissions for custom roles are editable. See [“Creating a custom role”](#) on page 21.

The Customer Super Admin role is the first user logging into the Information Studio product. The Customer Super Admin role is responsible for configuring the customer authentication profile as well as for adding new users with Customer Admin role.

The Customer Admin is a role to which users can be added to give them specific abilities within the Information Studio application. This role is tasked with creation of tenants, users with appropriate roles, and other Customer Admin users for the same organization. It is also responsible for configuring the Active Directory authentication profile, for using their identity provider for authenticating, instead of using default Information Studio provided authentication.

Tenant Administrators are responsible for adding users and assigning various roles per their responsibilities in the organization. Tenant Admin users have permissions to monitor the deployment and the status of the jobs within the tenant. They are entitled to view audits of various actions taken by users under the tenant.

The IT Administrators are in charge of bringing up their organization-specific on-premise **Data Engines**. These **Data Engines** are set up to use the Veritas data connector framework, that connects various enterprise data sources (CIFS, NetBackup, and others) to the Information Studio deployment in the context of an organization and enables the end-users to access, visualize, and govern enterprise data assets. Information Studio users with IT Admin role also use the Information Studio application for visualizing data retrieved from connections. They are also capable of setting up policies related to data discovery, scan, and classification which are applicable enterprise wide.

Role privileges of Veritas™ Information Studio users

Depending upon their assigned role, users can add or assign other roles as well as perform various actions that are summarized in the following table. For example, a Customer Admin can only create other Customer Admins, Tenant Admins, IT Admins, and Tenant Users.

Note: A Customer Admin can create a Customer Admin for its own customers, and Tenant Admins for tenants under that customer. A Tenant Admin can create other Tenant Admins and tenant users under their tenant.

Table 3-1 Information Studio user role permissions

As a...	You can create...				
	Customer Super Admin	Customer Admin	Tenant Admin	Tenant User	IT Admin
Customer Super Admin	✓	✓	×	×	×
Customer Admin	×	✓	✓	✓	✓
Tenant Admin	×	×	✓	✓	✓
Tenant User	×	×	×	×	×
IT Admin	×	×	×	×	×

Editing a customer account

The identity provider of the customer account, by default, is set to **Local**. To configure Active Directory authentication, you need to edit the customer account details in the **Edit Customer Account** dialog box.

Note: You must be signed-in as a Customer Super Admin to edit customer accounts.

To configure Active Directory authentication

- 1 Sign in to Information Studio as the Customer Super Admin.
- 2 In the navigation pane on the left, click **Organizations > Customer**.
- 3 At the right-end of the **Default_Customer** entry row, click the vertical ellipses and select **Edit Customer**.
- 4 In the **Edit Customer Account** dialog box:
 - In the **Details** section, edit the **Account Name** if required, then click **Next**.
 - In the **Register Authentication Provider** section:

- Select **Add new Active Directory config** from the **Select Provider** drop-down, and enter or select the following:

Field	Description
Protocol Select LDAP or LDAPS	The connection protocol used between Information Studio and the Network Directory or Domain Controller.
Port	The network port.
New Profile Name	The name of the current Active Directory profile.
DNS Server	The DNS name of the domain controller or the IP of the domain to which the server computer responds to security authentication requests.
Domain Name	Active Directory domain name.
Username	Full name of the user account or service account you want to use to test the connection.
Password	The password of the account you want to use to test the connection.

Note: The **Username** and **Password** entries are used to test the connection and there is no requirement for an Active Directory service account.

Note: All the above entries are mandatory.

Note: If you select **LDAPS** as the **Protocol**, the **Enable certificate verification for LDAPS** check-box is selected by default, and you need to click **Choose File** to load the **Certificate**. See [“Certificate management”](#) on page 23.

- Click **Test Connection** to ensure that the credentials you have entered allow you to connect to the **DNS Server**.
After you receive a confirmation on the test connection is successful, click **Save**.

Note: After you have edited the customer account, proceed to adding a user with the Customer Admin role. See [“Granting access to Information Studio users”](#) on page 21.

Granting access to Information Studio users

The roles available for assignment depend on the role of the signed-in user. For example, if you are signed-in as a Customer Admin, you can create another Customer Admin, Tenant Admin, IT Admin, and Tenant Users.

Veritas™ Information Studio comes with a default customer account with the user role of Customer Super Admin. Once signed-in and configured using Active Directory authentication, a Customer Super Admin can edit the default customer account and add users with the role of Customer Admin. The Customer Admin can now add more users with the role of Customer Admin, Tenant Admin, IT Admin, or Tenant User. Tenant Admins, in turn, can add more users with the role of Tenant Admin, IT Admin, or Tenant User.

Information Studio by default allows customers to create and manage users for users other than local users, such as the Customer Super Admin.

To add a user

- 1 In the navigation pane on the left, click **Organizations > Users**.
- 2 Select the radio button for the **Context**.
 - Click **+ New User** to open the **Add New User** dialog box.
 - Enter the **First Name**, **Last Name**, **User Name**, **Domain**, and **Email** in the respective fields.
 - Select a role from the **Role** drop-down for the applicable role.

Note: The roles available for assigning to the new user depend on the role of the signed-in user.

- Click **Save** to save and add the user details.

Creating a custom role

You can create custom roles with minimum privileges for specific tasks in the Information Studio administration console. For example, you can create a custom role specifically for users to be able to initiate and approve a Delete request.

Note: A user with either of the roles, Customer Admin or Tenant Admin can create custom roles.

To create a role for the delete action

- 1 Log into the Information Studio administration console, and click **Organizations > Roles** in the left navigation pane.
 - 2 In the **Roles** screen, select a **Context**, then click **Add New Role**.
 - 3 In the **Permissions** pane on the right side of the **Roles** screen, enter/ select the following:
 - **Name** for the new role. For example, **Delete_Requester** or **Delete_Approver**.
 - **Description** for the new role.
 - (Optional) **Copy Permissions from Role** - select an existing role from the drop-down to copy permissions from that role to the new role or individually select each permission for a role.
 - Customer Super Admin
 - Customer Admin
- Click **Select All** to select all the permission check boxes or **Deselect All** to deselect the permissions.
- 4 Click **Save** to save the new role or **Cancel** to exit the Permissions panel without saving the new role.
 - 5 To edit or delete a role, select the vertical ellipses in the right-hand side of the **Roles** name and select **Edit** or **Delete**, respectively.

Veritas™ Information Studio tenant management

Note: This section describes how the Customer Admin can add, delete, or modify an Veritas™ Information Studio tenant.

Note: For Information Studio 1.0, Veritas supports 1 active tenant.

To add tenants

- 1 In the navigation pane on the left, click **Organizations > Tenant**.
- 2 Click **+ New Tenant** to open the **Add Tenant Account** dialog box.

- 3 Enter the **Account Name** for the new tenant account.
- 4 Enter the **First Name** and **Last Name**.
- 5 Select the radio button for the **Login Type**.
- 6 Select a role from the **Role** drop-down.
- 7 Click **Save**.

Once the tenant is activated, the newly created tenant administrator receives a welcome email to update the password for signing into Information Studio as the Tenant Admin.

Certificate management

About SSL client server certificates

Information Studio secures all data flowing between the Information Studio web server and the clients using the Secure Socket Layer/Transport Layer Security (SSL/TLS) protocol.

Information Studio implements authentication with the mandatory use of client and server-side certificates or keys. Connections between the Information Studio servers and clients use a single, self-signed or CA-signed certificate. A self-signed certificate is bundled with the OVF template and is installed with the application and is unique to your deployment.

The SSL certificates can be publically signed or self-signed certificates from your own CA. In the case of publically signed certificates, internet access may be required to verify the authenticity of the certificate against the CA.

Information Studio currently supports any certificates with the addition and replacement of PEM certificates. The certificates are stored as `tls-secret` in kubernetes cluster in each installation instance.

The following high-level steps are involved in the generation and installation of a valid SSL certificate:

- Generate a new certificate.
- Generate a Certificate Signing Request (CSR) and send to a signature authority to be signed.
- Create a PEM file.
- Upload the SSL certificate to the Information Studio server.

Generating a new SSL certificate

You access the Information Studio user interface with a web browser. The Information Studio server and browser communicate through an SSL connection. To ensure confidentiality, all communication between the server and the browser is encrypted using a symmetric key. To initiate a connection, the server and browser negotiate the encryption algorithm (algorithm, key size, and encoding) and encryption key to use.

This section describes how to generate a new certificate, how a Certificate Sign Request (CSR) and how to install the certificate on the Information Studio server. Certificates can be generated in many different ways. Veritas recommends using the KeyStore Explorer, a graphical UI for the Java Keytool. You can also choose to use the Java Keytool, similar to OpenSSL tools.

To generate a new certificate using the KeyStore Explorer

- 1 Open KeyStore Explorer.
- 2 Click **Create new KeyStore**, and select JKS.
- 3 Right-click in the KeyStore Explorer UI and click **Generate Key Pair**.
- 4 Select RSA with a size of 2048 bits (recommended), and click **OK**.
- 5 On the Generate Key Pair Certificate page, configure the following certificate settings.
 - Select **Version 3**.
 - It is recommended to use SHA 256 with RSA (or stronger) as the signature algorithm. MD5 or SHA-1 RC4 are no longer supported by modern browsers.
 - Enter a suitable validity period for the certificate, and click **Apply**.
 - The serial number field is automatically populated.
 - Select **RSA** with a size of 2048 bits (recommended), and click **OK**.
 - The serial number field is automatically populated.
 - In the Name field, set CN to the fully qualified domain name of the Information Studio server, for example, <server name>.<domain name>.com, OU to Organization Unit name, O to Organization name and E to email address of the person to be notified. Also set L to location name, ST to name of state and C to name of country.

Note: The fully qualified domain name must be the actual name of the server that is accessible by all the clients.

- 6 Click **Add Extensions**, click the + icon and select **Subject Alternative Name** (SAN) as the extension type. Click **OK**. In the following dialog add at least one DNS entry, and General Name Value enter the hostname of the Information Studio server, for example, `InformationStudio.com`, and click **OK** until you return to the Generate Key Pair Certificate dialog.
- 7 Enter an alias for the key pair to bind it to that server, and click **OK**.
- 8 Enter the key pair password, and click **OK**.

Note: Save this password for use as the password for the entire keystore.

The key pair generation successful message is displayed.

- 9 Save the keystore. From the KeyStore Explorer menu, click **File > Save**. You will be prompted to set a key store password.

Note: The password of the key pair generated in 8 and the keystore password must be the same.

- 10 Enter the name of the keystore, and click **Save** to save the keystore file to your computer.
- 11 The generated keystore is listed on the home page of the KeyStore Explorer UI. Double-click the keystore to check its properties.

Generating a Certificate Signature Request

Once you have generated a key pair and save it to a keystore file, you must create a Certificate Signature Request (CSR) file for an external Certification Authority (CA). A CSR file is the request that you submit to the CA to obtain a signed certificate.

To generate a CSR

- 1 Right-click the key pair entry that you have already generated, and select **Generate CSR**. You will be prompted to enter the keystore password. See [Generating a new SSL certificate](#).
- 2 On the Generate CSR pop-up, do the following:
 - Select the format of the certificate - PKCS#10.
 - In the Challenge phrase, enter the challenge phrase for the certificate. This challenge phrase is required when filling out the certificate signing form on the CA's website.
 - Select the **Add certificate extensions to request** check box.

- Click Browse to navigate to the location where you want to save the generated CSR file.
 - Click OK to create a CSR file. You must submit this file to the Certification Authority (your own or a third party, such as VeriSign).
- 3 To obtain a signed certificate from your internal CA, contact your system administrator for instructions.

To obtain signed certificates from a CA, go to their web site and follow the instructions to enroll and obtain a signed certificate. To purchase the signed certificate, you generally need the following information, in addition to the CSR.

- Your personal details, including your name, email address, department name and number, and the business owner
- The common name. This name is the host name and domain name, such as www.company.com or company.com
- The validity period for the certificate
- The number of servers that host a single domain (up to five servers)
- The server platform
- The organization, organizational unit, country, state, or locality (all spelled without abbreviations)
- The number of servers that will use the certificate
- Payment information and a billing contact
- An email where the CA can reach you to validate the information
- Documentation to demonstrate that your organization is legitimate

You must check with the organization to identify any additional environment information that may be needed for the certificate.

- 4 Select the Certificate Signature Algorithm as **SHA-256 with RSA and SHA-256 root**
- 5 Browse to the location where you have save the generated CSR file and upload the CSR file on the signature authority's web site.
- 6 Enter a challenge phrase that will be used to renew or revoke the certificate.
- 7 Click **Get Certificate**.

The certified CA sends you the signed certificate. This process might take 3-5 days.

Importing signed certificates to your keystore

Once you receive a response from the CA, download the signed certificate and import it using the KeyStore Explorer.

To import the signed certificate

- 1 Log in to the CA website and download your signed certificate.
- 2 Open a text editor and paste the entire body of the certificate into the text file. Make sure to include the beginning and end tags on each certificate. Include the top line and bottom line (-----Begin Certificate----- and -----End Certificate-----). Make sure that no extra lines, spaces, trailing carriage returns, or characters have been inadvertently added.
- 3 Save the file with P7B encoding.
- 4 Open the KeyTool Explorer.
- 5 Open the Keytool Explorer, Right-click the certificate, and select **View Details > View Certificate Chain Details**. Navigate to the saved .p7b file, and use the challenge phrase to open the file.
- 6 Click **PEM** to convert the .p7b file to PEM file format. Copy and paste the content in a text file.
- 7 Right-click the certificate, and select **View Details > View Private Key Details**. Click **PEM** to convert to PEM file format.
Copy and append the content in the text file.
- 8 Click **File > Save** to save the imported certificate to your keystore.

Installing the SSL certificate

To install the SSL certificate

- 1 Log in to the Information Studio administration console.
- 2 Click **Settings > Certificates**.
- 3 Click **Replace certificate**.
- 4 On the **Replace certificate** pop-up, enter the alias used when generating the key pair. See [Generating a new SSL certificate](#).
- 5 Browse to the location where the .pem file and .key files are saved, select the file, and click **Apply**.
 - The replaced certificate may take a while before it is reflected.
 - After the certificate change is reflected in the Information Studio administration console, all the users are logged out and re-authenticated with the new SSL keys.

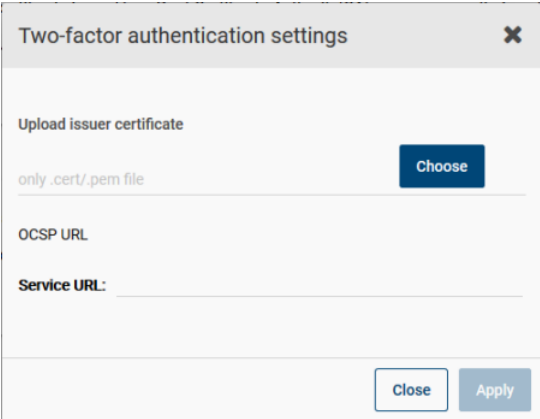
If the upload of the signed certificate fails for any reason, the original certificate is restored.

Multi-factor authentication support

Public sector organizations use digital certificates for user authentication. The digital user certificates are stored on smart cards protected by a user PIN. A combination of smart card and a PIN is used for two-factor authentication. The extraction of digital user certificate from a smart card is handled outside of the product. As part of supporting two-factor authentication (2FA), a browser presents the user certificate to be validated using an Online Certificate Status Protocol (OCSP). The OCSP URL is configured by the system administrator. The validation of the certificate confirms successful authentication of the user. The Distinguished Name (DN) value is extracted from the certificate and used to generate a JSON Web Token (JWT-based tokens).

To enable multi-factor authentication

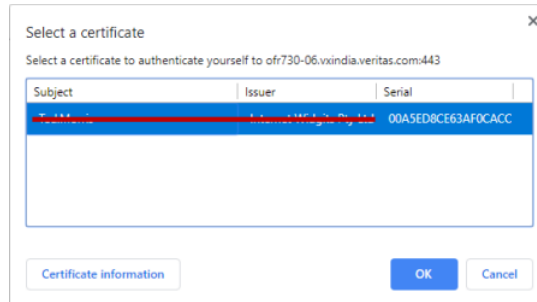
- 1 Log into Information Studio with the Customer-Super-Admin role.
- 2 Click **Organizations > Certificate Management**, and in the **Two-factor authentication** pane, click **Configure 2FA**.
- 3 In the **Two-factor authentication settings** dialog box, click **Choose** to navigate to the location and **Upload issuer certificate** enter the OCSP URL, and click **Apply**.



The image shows a dialog box titled "Two-factor authentication settings" with a close button (X) in the top right corner. The dialog contains the following elements:

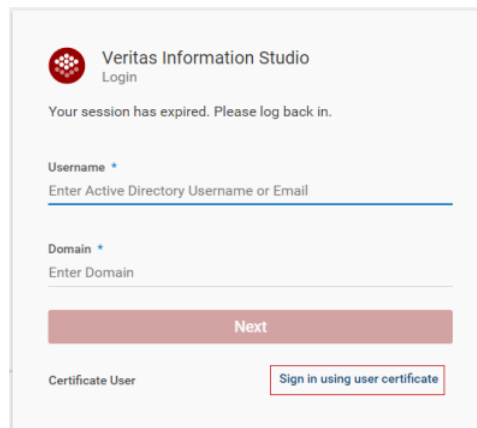
- A section labeled "Upload issuer certificate" with a text input field containing the placeholder "only .cert/.pem file" and a blue "Choose" button to its right.
- A section labeled "OCSP URL" with a text input field.
- A section labeled "Service URL:" with a text input field.
- At the bottom right, there are two buttons: "Close" and "Apply".

- 4 Import the user certificate into the browser. Take Chrome as an example, double-click the .pfx file, enter the password, and click **OK**. The user certificate is imported into your keychain.



- 5 To launch the Information Studio UI in the browser using `https://<host>`, and in the **Select a certificate** pop-up, select a certificate to authenticate yourself to log in, then click **OK**.

You can then see the Information Studio login screen with the 2FA certificate login option



- 6 Enter your **Username**, **Domain**, then click **Sign in using user certificate** to log in.

Configuring SMTP (Simple Mail Transfer Protocol) services

You must configure the SMTP service to enable Information Studio to send email notifications for various actions, such as Delete action settings, carried out in the Information Studio application:

If the SMTP settings are not configured, you cannot use email to send action-related notifications.

To configure the SMTP settings

- 1 Sign in to Information Studio as a Customer Admin.
- 2 Click **SMTP Configuration** in the left navigation panel of the Information Studio administrator console to view the SMTP details.
- 3 On the **SMTP Configuration** page, enter the following details:

Field	Description
Send emails from	Specify the email ID that is used to send notifications from Information Studio.
Mail server hostname/ IP address	Specify the mail server hostname or IP address used to send notifications.
Mail server port	Enter the SMTP server port number.
Enable TLS	Select the check box if the mail server you have specified requires the Transport Layer Security.
Enable authentication	Enable the check box if the mail server you have specified requires authentication, and enter the username and password of the user account.
<div><div>■ Username</div><div>■ Password</div></div>	<p>Note: A valid input in the username field can be the username or email ID of the user account configured on the SMTP server.</p>

- 4 Click **Send test email** and in the pop-up that appears enter a valid email id, then click **Send** to send the test email.
- 5 On successful validation of the email id, click **Save** to save the SMTP configuration details.

Licensing

This chapter includes the following topics:

- [Overview](#)
- [Trial license model](#)
- [Subscription license model](#)
- [Perpetual license model](#)

Overview

The licensing feature in Information Studio administration console helps you provide license-related information into the system. The licensing feature enables you to:

- View type of license.
- Check the expiry of the license.
- Check the available capacity.

Accessing License Manager and adding license information

As a Customer Admin you can view the license information like its expiry date and the available capacity. As a Customer Admin, you have also the privilege to update the license from the Information Studio administration console. IT Admins, Tenant Admins, and Tenant Users can only view the license package information and expiry date of the license, and need to contact a Customer Admin to add or update the license.

As a Customer Admin logging into the Information Studio administration console for the first time, you can see the **License** pop-up to add a license. Only after adding a valid licensing file you (Customer Admin) can access and use Information Studio.

License models available for using Information Studio include:

- Trial

Note: You can not apply trial license on any other type of license.

- Subscription
- Perpetual

These license models are applicable to the Information Studio administration console as well as the Information Studio application.

Trial license model

The Information Studio trial license model allows users (Customer Admins or Tenant Admins/ IT Admins/ Tenant Users) to try out Information Studio for a period of 30 days and then opt for the subscription model.

Trial Customer Admins with expired licenses can choose to update the license and continue using Information Studio. Customer Admins who do not update the license are locked out of the product.

Trial Tenant Admins/ IT Admins/ Tenant Users with expired licenses are prompted to contact the Customer Admin to update the license to be able to use Information Studio.

To add or update the license information as a trial user

- 1 Log into the Information Studio administration console with Customer Admin credentials.

Note: If you are a first-time user, you will see the **License** pop-up to add a license.

- 2 In the **License Details** pop-up, click **Add License** to upload a new license.

- 3 In the **Add License Details** dialog box, click **Choose file** to browse to the location of the license file (procured from Veritas Entitlement Management System (VEMS)) on your computer, verify the license details such as the license type and capacity displayed on the **Add License Details** modal, then click **Apply License**.

Note: Click **Log out** to exit Information Studio if you receive an error while uploading the license file.

- 4 Click **Apply License**.
The license is applied and you are directed to the **License Manager** view. License validity and the data limit for the applied license is listed on the **License Manager** screen.
- 5 Click **License Manager** in the left navigation panel to view available license details.
As a Customer Admin, you can choose to update the license at any time using the **Update License** option in the top-right corner or from the license sidebar in the left-bottom of the UI.

Subscription license model

The Information Studio subscription license model allows users (Customer Admins or Tenant Admins/ IT Admins/ Tenant Users) to subscribe to Information Studio for an extended period.

Customers Admins who are subscribing to Information Studio on license expiry are prompted to update the license.

Tenant Admins/ IT Admins/ Tenant Users subscribing to Information Studio, on license expiry are prompted to contact the Customer Admin to update the license.

To add or update the license information as a subscribed user

- 1 Log into the Information Studio administration console with Customer Admin credentials.

Note: If you do not have a valid license installed, you will see the **License** pop-up to add a license.

- 2 In the **License Details** pop-up, click **Add License** to upload a new license.

- 3 In the **Add License Details** dialog box, click **Choose file** to browse to the location of the license file (procured from Veritas Entitlement Management System (VEMS)) on your computer, verify the license details such as the license type and capacity displayed on the **Add License Details** modal, then click **Apply License**.

Note: Click **Log out** to exit Information Studio if you receive an error while uploading the license file.

- 4 Click **Apply License**.
The license is applied and you are directed to the **License Manager** view. License validity and the data limit for the applied license is listed on the **License Manager** screen.
- 5 Click **License Manager** in the left navigation panel to view available license details.
As a Customer Admin, you can choose to update the license at any time using the **Update License** option in the top-right corner or from the license sidebar in the left-bottom of the UI.

Perpetual license model

The Information Studio perpetual license model allows users (Customer Admins or Tenant Admins/ IT Admins/ Tenant Users) to use Information Studio for an extended period.

Customers Admins under the Perpetual model on license expiry are prompted to update the license and can, however, continue to use Information Studio without updating the license.

Tenant Admins/ IT Admins/ Tenant Users under the perpetual model, on license expiry are prompted to contact the Customer Admin to update the license.

To add or update the license information as a perpetual Information Studio user

- 1 Log into the Information Studio administration console with Customer Admin credentials.

Note: If you do not have a valid license installed, you will see the **License** pop-up to add a license.

- 2 In the **License Details** pop-up, click **Add License** to upload a new license.
- 3 In the **Add License Details** dialog box, click **Choose file** to browse to the location of the license file (procured from Veritas Entitlement Management System (VEMS)) on your computer, verify the license details such as the license type and capacity displayed on the **Add License Details** modal, then click **Apply License**.

Note: Click **Log out** to exit Information Studio if you receive an error while uploading the license file.

- 4 Click **Apply License**.

The license is applied and you are directed to the **License Manager** view. License validity and the data limit for the applied license is listed on the **License Manager** screen.

- 5 Click **License Manager** in the left navigation panel to view available license details.

As a Customer Admin, you can choose to update the license at any time using the **Update License** option in the top-right corner or from the license sidebar in the left-bottom of the UI.

Dashboard

This chapter includes the following topics:

- [About the dashboard](#)

About the dashboard

Once you log into Information Studio, the Information Studio administration **Dashboard** is displayed. The **Dashboard** provides an overall view of the system. You can view the total number of configurable connector types, content sources, Applications, Jobs, Connectors, the Asset Map (the world map showing the location of Information Studio Hub, Data Engines, and Content Sources) .

The **Dashboard** is divided into the following 3 areas:

[Summary bar](#)

[Asset Map](#)

[Widgets](#)

Summary bar

The summary bar shows the number of **Active Jobs**, **Content Sources**, **Applications**, and **Connectors** based on the signed-in user role.

In the top-right corner of the summary bar, there is a refresh option. You can click the refresh option to refresh/ reload the data in the widgets.

The table shows the details that each type of user role can view in the summary bar.

Table 5-1 Summary bar

As a...	You can view the number of			
	Content Sources	Active Jobs	Applications	Connectors
Customer Super Admin	×	×	✓	✓
Customer Admin	×	×	✓	✓
Tenant Admin	✓	✓	✓	✓
IT Admin	✓	×	✓	✓
Tenant User	×	×	×	×

Asset Map

The **Asset Map** shows the locations of **Information Studio Hub**, **Data Engine** (on-premise), or **Content sources** by marking them on the world map. You can filter the markers on the map by selecting a type of asset from the right-side panel. You can zoom-in, zoom-out, or reset the map using the controls in the bottom-left corner of the **Asset Map**.

The following table describes the details seen in the **Asset Map** based on the signed-in user role.

Table 5-2 Details that each role views in the asset map

As a...	You can view the counts and the locations of...		
	Information Studio Hub	Data Engine	Content Sources
Customer Super Admin	✓	×	×
Customer Admin	✓	×	×
Tenant Admin	✓	✓	✓
IT Admin	✓	✓	✓
Tenant User	×	×	×

Widgets

Below the **Asset Map** in the **Dashboard**, you can see a set of widgets , you can see a set of widgets based on the role you have signed-in as.

You can drag and drop the widgets to re-arrange them. The re-arrangement persists on subsequent sign-in until you again change the arrangement of the widgets.

- **Health**

The **Health** widget is a chart that displays the health of each **Data Engine**. An inner pie shows the overall aggregated health of the **Data Engine** while an outer pie shows the health of each service in the **Data Engine**.

- **Job Types**

The **Job Types** widget is a bar chart that displays the total number of jobs sorted by the type of job created in the system

- **Recent Active Jobs**

The **Recent Active Jobs** widget is a table that displays the details of the recent jobs in active state.

- **Connectors by Deployment & Type**

The **Connectors by Deployment & Type** widget shows the connectors that are grouped by type of data (structured, semi-structured, and unstructured) and type of deployment (on-premises, cloud, or Veritas-integrated). Click on the **On-Premise Connectors** tab to display the on-premise connectors.

Applications

This chapter includes the following topics:

- [About Information Studio application](#)

About Information Studio application

You can access the Information Studio application from the **Applications** link in the navigation pane on the left. The Information Studio application also shows up in the top part of application switcher. See [“About the application switcher ”](#) on page 16.

If you have signed in as an IT Admin, click **Open** to launch the Information Studio application console from the **Applications** screen.

For details on the user interface and how to use this application, refer to the *Veritas™ Information Studio User Guide*.

Data Engine

This chapter includes the following topics:

- [About Information Studio data engine](#)

About Information Studio data engine

Veritas™ Information Studio constitutes one Information Studio Hub and one or more Data Engines that are required for monitoring your data. The Data Engine is made up of several micro-services which are critical for capabilities such as logging, monitoring, and data processing.

You can add one or more **Windows Connectors** nodes to the **Data Engines** that you have deployed.

IT Admins can access Data Engines from the **Data Engines** link in the navigation pane on the left. For more information about adding a new Data Engine to the deployment, refer to the *Veritas™ Information Studio Deployment Guide*.

Connectors

This chapter includes the following topics:

- [Overview of Connectors](#)
- [Managing credentials](#)
- [Adding connections](#)
- [Configuring Connections to cloud content sources](#)
- [Configuring Connections to on-premises content sources](#)

Overview of Connectors

The **Connectors** screen in the Information Studio administration console enables users with IT Admin role to configure and manage connections to on-premise, cloud and Veritas Integrated content sources

Details pane

The **Connectors** option comprises the following two tabs:

- **Credentials** - This displays the credentials used by Information Studio to discover and scan content sources.
- **Connections** - This displays all configured connectors (content sources) that Information Studio monitors.

The configured Credentials and Connections are listed alphabetically. You can search the list according to their names and types.

Supported connectors

Information Studio supports the following connectors. These supported connectors when configured in Information Studio enable you to discover data across various content sources.

Cloud connectors

- Box for Enterprise
- Generic S3
- Google Cloud Storage
- Google Drive
- Gmail
- Microsoft Azure
- Microsoft Exchange Online
- Microsoft OneDrive
- Microsoft SharePoint Online

On-premises connectors

Table 8-1 Supported versions

Connector	Supported version
IBM FileNet	Versions 5.1
OpenText Documentum	Versions 6.7
OpenText LiveLink	Versions 10.5
EMC Isilon	OneFS version 7.1.0.6 or higher
EMC Celerra	DART version 5.6.45 or higher
Hitachi NAS	Hitachi NAS 12.x
NetApp Cluster	
NetApp Standalone	7.3.5 or higher
Windows File Server (CIFS)	<ul style="list-style-type: none">■ Windows Server 2008 or 2008 R2 (32 bit and 64 bit)■ Windows Server 2012 or 2012 R2 (64 bit)■ Windows Server 2016 or 2016 R2 (64 bit)

Table 8-1 Supported versions (*continued*)

Connector	Supported version
Microsoft Exchange on-premises	Exchange 2013 and Exchange 2016 running the latest service pack.
Microsoft SQL Server	<ul style="list-style-type: none">■ SQL Server 2008■ SQL Server 2012■ SQL Server 2016
Microsoft SharePoint	<ul style="list-style-type: none">■ Microsoft SharePoint Server 2013■ Microsoft SharePoint Server 2016
Oracle Database	<ul style="list-style-type: none">■ Oracle Database 12c Enterprise Edition Release 12.2.0.1.0 64-bit Production■ Oracle Database 11g Standard Edition 64-bit Production■ Oracle Database 10g Express Edition Release 10.2.0.1.0 - Production
Veritas NetBackup	<ul style="list-style-type: none">■ Veritas NetBackup 7.6.x, 7.7.x, 8.0, 8.1, 8.1.1, 8.1.2, and 8.2■ Veritas NetBackup Appliance, 2.6 or later

Requirements and prerequisites

Before configuring a content source for monitoring, ensure that the following requirements are met:

- The connector machine and the on-premises content sources are part of the same domain or are in the trusted domains.
- The connector machine is able to successfully communicate with Information Studio Hub.
- Ensure that the Data Engines and Information Studio Windows Connectors are installed. Information Studio Windows Connectors should communicate with their respective Data Engines and Data Engines should communicate with the Information Studio Hub. If it is a single-node deployment, then the Information Studio Hub includes a default Data Engine and you need to install only Information Studio Windows Connectors.
- The connector service port is accessible and used by a single service.

- In case of on-premises content sources, make sure that the Data Engine and Information Studio Windows Connectors are geographically close to the content being monitored.
- The fully qualified domain name (FQDN) address that you specify when adding a connection must be resolvable from the Data Engine and Information Studio Windows Connectors.

In addition to the requirements mentioned in this topic, there may be configurations that must be completed before you can add a connection from the **Connectors**. Ensure that you complete the connection-specific requirements before adding a connection in **Connectors**.

High-level workflow for setting up connections to content sources from the Connections console

1. In case of on-premises content sources, the on-premises Data Engine and Information Studio Windows Connectors must be deployed. See [“Configuring Connections to on-premises content sources”](#) on page 82.
2. In case of cloud content sources, configure your cloud accounts to allow access to the Information Studio Cloud Connectors (for example, Enterprise Box, Microsoft OneDrive, and so on). See [“Configuring Connections to cloud content sources”](#) on page 52.
3. Launch the Information Studio administration console with IT Admin credentials and access the Connectors.
4. Configure credentials that Information Studio will use to discover and scan the content sources.
5. Set up a connection to the content source. See [Adding connections](#).
6. Assign a pre-configured Discover and Scan policy to the connection. You can also bind the connector with the Discover and Scan policy after adding the connections.
7. Visualize your data from the Information Studio application console. See *Veritas™ Information Studio User Guide*.

Known limitations of Connections console

Data collection limitations for other on-premises and cloud data stores are as follows:

CMIS (OpenText Documentum, Opentext LiveLink, IBM FileNet)

- The Connector does not fetch the last accessed user and date information.

- When a file has multiple versions, the **Map** in the Information Studio application console displays each file version as a separate entry with the file name as `<filename>_<version_number>.<extension>`.

OpenText Documentum

- Documentum permits creation of multiple files with same name within the same folder. However, the Connector does not collect metadata for such files.
- When a file is created but not checked in to the Documentum server, an empty template file gets created. Since a file is not present, the Connector does not fetch the extension of such template files.
- If you do not upload any files or choose to create a file using a template without adding any contents to it, Documentum creates only an empty template. The **Connections** console is, therefore, unable to fetch the size of the file and extension as there is no actual file present on the Documentum server.

IBM FileNet

- The Connector does not collect the last modified date for the root folder.

Opentext LiveLink

- The Connector collects metadata only for files with type Documents. It ignores other types such as Tasks, Polls, and Projects.
- The Connector discovers both, user's private space (user's repository) and public repositories. However, it scans only those repositories for which the user has sufficient access permissions.
- If a file has multiple versions and each version has a different extension, then the Connector does not collect the version details for such a file..

Microsoft SQL

- When the Microsoft SQL Server service is restarted, the Connector does not collect the last access and modified time information.

Microsoft SharePoint

- Sites that are locked are discovered, however, not scanned by the Connector.

Oracle Database

- The Connector does not collect the last access date and user details for Oracle data store.

Microsoft OneDrive

- For the Connector to discover the configured OneDrive account, the user of that account should log on to the account at least once.
- The Connector does not discover the accounts that have expired.
- When the Connector scans OneDrive accounts, it does not collect the location information. Consequently, the location for such data stores is not populated on the Map UI.

Microsoft SharePoint Online

- When scanning a SharePoint Online account, the Connector does not collect the location information. As a result, location for such data stores is not populated on the Map UI.
- The Connector discovers only the default document library (documents) and custom document libraries.

Managing credentials

You can store account credentials for the content sources (connections) that enable the Information Studio Windows Connectors to access data from the content source hierarchy.

The authentication credentials can be stored in a central credential store and referenced when configuring connections in the **Connections** console. Saving credentials simplifies the management of changes to the account user name and passwords.

For cloud content sources, such as Box and Microsoft OneDrive, you do not need to add credentials when configuring the content sources. The authentication for these data stores is controlled using the OAuth 2.0 workflow for each connector. For these content sources, you are redirected to the respective tenant accounts for authorization, where you must perform an interactive registration to acquire the authentication tokens.

Once registered the connector acquires two tokens that are used to authenticate itself to the Web APIs:

- The Access token is valid for 60 minutes; it can be exchanged for a fresh token when it expires.
- The Refresh token is valid for 60 days. In case of Microsoft OneDrive and Microsoft SharePoint Online, the Refresh token is valid for 90 days. Each time you request a new token from Microsoft, a new refresh token is returned as well.

You can add, edit, or delete the saved credentials from the **Connectors > Credentials** list page.

To add a credential

- 1 To add a credential, on the **Credentials** list page, click **+ Add Credential**.
- 2 Choose a Connector type, then in the **Enter Credential Details** panel, enter the **Name**, **Description**, **Domain**, **User Name**, **Confirm Password**, and **Confirm the password**.
- 3 Click **Save**.

You can view the newly-added Credential in the Credentials list screen.

To edit or delete a credential

- 1 To edit a credential, on the **Credentials** list page, select the credential that you want to edit.
- 2 Click the **Actions** drop-down, and click **Edit**. On the **Edit credential** page, make changes to the required parameter, and click **Save**.
- 3 To delete a credential, on the **Credentials** list page, select the credential that you want to edit. Select **Actions > Delete**.

Adding connections

Users with IT admin privileges can add content sources for monitoring from the Information Studio administration console.

To add a connection

- 1 Sign in to the Information Studio administration console URL as an IT Admin.
- 2 Click **Connectors**. The **Credentials** and **Connections** tabs are displayed.
- 3 The **Connections** tab displays information about the various content sources that Veritas™ Information Studio monitors for metadata. You can add new connections here.
- 4 To add a NetBackup connection, click **Add NBU Connection**. See [“Configuring a NetBackup connection”](#) on page 104.

To add a connection for a content source other than NetBackup, click **Add Other Connection**.

- 5 The workflow for adding a content source from the **Connections** console wizard is as follows:
 - If adding a connection other than NetBackup, select the content source.

- Add the credential that Veritas™ Information Studio uses to connect with the content source. You can choose to create use an existing credential or add a new one.
- Add details of the content source itself. For example, the server IP address or host name, the name of the domain to which the content source belongs, and the **Data Engine** responsible for collecting the metadata. For information about configuring specific cloud and on-premises content sources:
 - See [“Configuring Connections to cloud content sources”](#) on page 52.
 - See [“Configuring Connections to on-premises content sources”](#) on page 82.

Note: Information Studio validates the **Data Engine** selected for the **Connection** being added. If the **Data Engine** does not support the **Connection** being added, then you will not be able to add that **Connection**.

- Add a schedule to run the discovery and scan actions. See [“Changing a pause schedule in Connections”](#) on page 51.
- Associate an existing Discover and Scan administration policy with the connection. The policy defines the schedule and the scope of the discovery and scan jobs that run on the connector. You can choose to assign a policy when adding the connection or at a later time. See [“Assigning policies to Connectors”](#) on page 51.

Note: Make sure that the FQDN of the connection is resolvable from the Data Engine and from the Information Studio Windows Connectors.

- 6** Click **Test Connection** to validate the selected credential and connection details before you save the **Connection**.

Note: If you choose to save the connection without testing, you are prompted to validate the connection before saving it. If you choose to dismiss the operation, the testing of the connection does not take place.

You are notified once the **Connection** is successfully validated or has failed. The failure can be due to invalid credentials, invalid **Connection** details, or **Connection** not being reachable.

- 7** After validating the **Connection**, click **Save Connection**.

The **Test Connection** feature is supported for the following connectors:

- Veritas Netbackup
- EMC Celerra
- EMC Isilon
- NetApp Standalone
- Microsoft SharePoint on-premise
- NetApp Cluster
- Windows File Server
- Hitachi NAS
- Gmail
- Oracle

Managing connections

You can edit or delete an existing connection or credential. You can choose to edit a connection to either change the configuration parameters, such as the host name or domain of the connection, associate a different credential with the connection, change the pause schedule, or to assign a policy to the connection.

To edit or delete a connection or credentials

- 1 On the **Credentials** or **Connections** list page, select the connection or credentials that you want to edit.
- 2 Click **Actions > Edit**.
- 3 Edit the required parameter, click **Test Connection** to test the edited **Connection**, then click **Save and Close**.

To delete a connection

- 1 From the relevant list page, select the credential or connection that you want to delete.
- 2 Click **Actions > Delete**.

Note: After you delete a connection, the policies bound to the connection are unbound, the relevant jobs for those connections are aborted, and the data associated with the deleted connection also gets deleted.

Content sources and hierarchies

The different data levels discovered for each data store are represented as content sources and content repositories in the Information Studio application console.

Table 8-2 Content source and hierarchies

Datastore	Content store	Repositories
Box for Enterprise	Box for Enterprise account	User Box drives
Amazon S3	S3 buckets	S3 buckets
Google cloud Storage	Regional bucket	Bucket
Microsoft Azure	Azure storage account	<ul style="list-style-type: none"> ■ Blob container ■ File share ■ SQL Server
Microsoft OneDrive	OneDrive Enterprise Account (Office 365 Account)	User drives
Microsoft SharePoint online	Site collection	Filers, shares, and folders or files
OpenText Documentum	URL of the content management servers	Repository
OpenText Livelink	URL of the content management servers	Repository
IBM FileNet	URL of the content management servers	Repository
Native File Servers	Fully Qualified Domain Name (FQDN) of the file server	UNC share path
Microsoft Exchange Server	Mailbox name fetched from the Active Directory	User primary email address
Microsoft SQL Server	Server instance	Database
Microsoft SharePoint on-premise	Web application name/site collection	Document library
Oracle Database	FQDN of the server instance	Database
Backup Exec	FQDN of the Backup Exec server	File system volume or UNC share name
Veritas NetBackup	File Server	Drives or Shares

Changing a pause schedule in Connections

Changing a pause schedule in the **Connections** console lets you pause the Discover or Scan policy job for the selected hours of the day in a week. This setting ensures that the Discover or Scan action is not triggered for the stipulated time duration and you cannot fetch the metadata residing on this content source.

You might want to pause a job schedule when there is a high workload on Information Studio. When you configure a pause schedule for connections which have significantly large amount of data residing on them, Discovery and Scan jobs are halted which ensures that content sources aren't getting scanned during the peak business hours when users are accessing content sources.

You have the option of defining a pause schedule when configuring a connection in the **Connections** console.

To change a pause schedule

- 1 On the **Connection** tab of the configuration wizard, click **Change pause schedule**.
- 2 Select the day and time interval from the displayed calendar.
- 3 Click **Save Schedule**.

If you want to make changes to the pause schedule, select the **Edit Pause Schedule** check box.

For information on configuring connections:

- See [“Configuring Connections to cloud content sources”](#) on page 52.
- See [“Configuring Connections to on-premises content sources”](#) on page 82.

Assigning policies to Connectors

A policy is a group of conditions and rules that you can apply to different connectors to determine the schedule and scope of discovery or scan actions. You can only associate an existing policy of type Discovery or Scan. These policies can be created using the **Policy Manager** in the Information Studio administration console.

Following are the types of policies that you can associate with the connectors in **Connections** console:

- **Discovery policy:** Lets you identify and discover the content source, as well as append repositories.
- **Scan policy:** Lets you scan and analyze the data residing on the content sources.

For more information on policies, refer to the section on Policy management. See [“Overview of Information Studio policies”](#) on page 114.

When configuring a connection, you may choose to assign a default or custom policy to the content source. You can assign an existing policy to the connection when configuring the connection or you can bind the connection with the desired policy later by editing the configuration at a later time. See [“Managing connections”](#) on page 49.

To assign a policy to the connection

- 1 In the Add Connection wizard, click the **Policies** tab.
- 2 Select the **Discovery** or **Scan** policy.
- 3 From the drop-down, select the policy. You can associate a default policy with the connection, if a custom policy is not available.
- 4 Select the scheduler type from the drop-down list.
- 5 Select **Run Now** or **Recurrence** scheduler type.
- 6 If you select the **Recurrence** option, then specify the frequency of the schedule.
- 7 Click **Save & Close**.

When a policy is bound to a connection and a schedule is assigned to it, it results in the creation of a job definition. The Information Studio job scheduling framework uses these job definitions to create specific job instances that are based on the job definition schedule. You can view the progress of these jobs on the **Monitoring > Jobs** list page. These jobs are listed in a descending order based on their creation date.

For more information on jobs: See [“About Jobs”](#) on page 136.

Configuring Connections to cloud content sources

Cloud content sources, are highly-scalable content sources that are hosted in your cloud environment. It is a collaborative space optimized for storing, sharing, and editing files in the cloud environment. The Connectors console provides the ability to configure connections for cloud content sources and manage those connections.

Configuring data collection from Microsoft Azure

Microsoft Azure Storage provides two authentication mechanisms; Storage Account Keys (SAK) or Shared Access Signatures (SAS). SAKs provide full unrestricted access to a Storage Account and should be kept private and secure at all times; they should also be regularly rotated in line with your organization's information security policy. SAS tokens provide restricted access to a storage account, allowing read and list only permissions to be set; they are derived using a SAK. SAS tokens

are designed to expire, meaning a third party using the SAS token must regularly request a new one.

Prerequisites for configuring data collection from Microsoft Azure Storage

To discover and scan items stored within Azure Storage, while also keeping your data secure, Information Studio leverages Azure Active Directory, Automation Accounts, and Key Vault alongside SAKs and SAS tokens. This mechanism allows a centralized role-based access control (RBAC) system to control the production of SAS tokens securely in your Azure environment, and allows access to these tokens to the **Connections** console through an Azure Active Directory service account.

Microsoft Azure Active Directory provides a centralized user and service account repository with RBAC to control access to Azure resources. Here you can create a service account for Veritas™ Information Studio to use to unlock dark data stored within Azure. See [Create an application account](#) for more information.

Microsoft Automation Accounts provide the ability to automate tasks within your Organization's Azure environment using PowerShell. Using an automation account you can set up scripts to rotate your SAKs and to generate SAS tokens which get stored in Azure Key Vault. Microsoft Azure Key Vault can securely store your cryptographic private keys and secrets, such as shared keys and passwords. You can grant privileges to Azure Active Directory users or services to access the key vault, perform encryption or decryption operations, or read and write secrets. See [Set up key vault with automation](#) and [Assign permissions to storage and SQL server account](#) for more information.

Veritas provides an example functioning PowerShell script to generate and store SAS tokens into Key Vault. You can install, modify, and configure this script as per your requirements. See [Install the automation script](#) and [Record the Azure account details](#) for more information.

Before you configure Microsoft Azure Storage from the Information Studio administration console, you must complete the following prerequisite tasks within your Azure environment. The configuration procedure is split into five high-level steps:

1. [Create an application account](#)
2. [Set up key vault with automation](#)
3. [Assign permissions to storage and SQL server account](#)
4. [Install the automation script](#)
5. [Record the Azure account details](#)

Create an application account

The following table shows the steps to create an application account for configuring data collection from Microsoft Azure connector:

Table 8-3 Steps to create an application account

Step (in the order of sequence)	How to?
Create (or reuse) an Azure-application Service Principal for the automation account in Azure Active Directory. For future reference, this will be called the Azure-application Service Principal.	<ol style="list-style-type: none"> 1 Log on to the Azure portal. 2 In the left navigation panel, click Azure Active Directory. 3 From the Azure Active Directory pane, click App registrations. 4 In the App registrations pane, click New registration. 5 On the Create pane, enter an appropriate name, select Web app / API as the Application type, and enter a sign-on URL of https://[name] where [name] is the name you choose. 6 Click Register. 7 From the App registrations pane, record the Application ID of the Azure-application Service Principal you just created. This value is input into the Automation as the AzureSASKeyGeneratorDaemon-ClientId variable. 8 From the App registrations pane, select the Azure-application Service Principal you created. 9 From the resulting right-hand pane, select Certificates & secrets. 10 Enter a description and select an appropriate expiry date that conforms to your information security policy, and click Save. 11 Record the displayed value, which is the value of the Application Secret. This value is input into the Automation as the AzureSASKeyGeneratorDaemon-ClientSecret variable.

Set up key vault with automation

Below are the steps to set up the Azure key vault for configuring data collection from Microsoft Azure connector:

Table 8-4 Steps to set up the Azure key vault

Step (in the order of sequence)	How to?
<p>Create or use an existing key vault. Veritas recommends using a new key vault to keep Information Studio access to SAS Tokens isolated from other keys and secrets.</p> <p>Note: You can use multiple key vaults to segregate different business functions (for example Scanning Pre-production Storage Accounts, and Scanning Production Storage Accounts).</p>	<ol style="list-style-type: none"> 1 In the left navigation panel of the Azure portal, click Key vaults. 2 From the Key vaults pane, click Add. 3 In the Create key vault pane, input a name, select a subscription, create or use an existing resource group and select a location for the Key vault. 4 Click Create.
<p>Provide the Azure-application Service Principal with Secret Management Rights on the Key vault(s) being used.</p>	<ol style="list-style-type: none"> 1 In the left navigation panel of the Azure portal, click Key vaults. 2 From the Key vaults pane, select the key vault you want to use to store the read-only SAS tokens for Information Studio. 3 Select Access control (IAM) for your key vault. 4 From the Access control (IAM) pane, click Add. 5 In the Add permissions pane, select Contributor as the role. Search for and select the name of the Azure-application Service Principal you created in step 1 (see Create an application account). 6 Click Save. 7 In the Add access policy pane, select the following: <ul style="list-style-type: none"> ■ Secret Management as the Configure from template option. ■ Azure-applicationService Principal as the Principal. ■ Under Secret Permissions ensure that only Set, Get, and List are selected. 8 Click OK.

Assign permissions to storage and SQL server account

To assign permissions to the storage and SQL Server accounts in Microsoft Azure connector, do the following:

Table 8-5 Steps to assign permissions to storage account

Step (in the order of sequence)	How to?
Provide the Azure-application Service Principal with Reader and Storage Account Key Operator Service roles on the Storage Account(s) you want to provide Information Studio access to.	<ol style="list-style-type: none"> 1 In the left navigation panel of the Azure portal, click Storage accounts. 2 From the Storage accounts pane, repeat the following for all storage accounts you want Information Studio to scan. <ul style="list-style-type: none"> ■ Select the storage account. ■ Select Access control (IAM) for your Storage account. ■ From the Access control (IAM) pane, click Add. ■ In the Add permissions pane, select Reader as the role, and then search for and select the name of the Azure-application Service Principal you created in step 1 (see Create an application account). ■ In the Add permissions pane, select Storage Account Key Operator Service as the role, and then search for and select the name of the Azure-application Service Principal you created in step 1 (see Create an application account).

Table 8-6 Steps to assign permissions to SQL Server account

Step (in the order of sequence)	How to?
Provide the Azure-application Service Principal with Reader and Storage Account Key Operator Service roles on the SQL Server Account(s) you want to provide Information Studio access to.	<ol style="list-style-type: none"> 1 In the left navigation panel of the Azure portal, click SQL servers. 2 From the SQL servers pane, repeat the following for all SQL servers you want Information Studio to scan. <ul style="list-style-type: none"> ■ Select the SQL Server account. ■ Select Access control (IAM) for your SQL Server. ■ From the Access control (IAM) pane, click Add. ■ In the Add permissions pane, select Reader as the role, and then search for and select the name of the Azure-application Service Principal you created in step 1 (see Create an application account). ■ In the Add permissions pane, select Storage Account Key Operator Service as the role, and then search for and select the name of the Azure-application Service Principal you created in step 1 (see Create an application account).

Install the automation script

To install the automation script in Microsoft Azure connector, do the following:

Table 8-7 Steps to install the automation script

Step (in the order of sequence)	How to?
Create or use an existing Automation account to run a PowerShell script that generates SAS Tokens periodically.	<div><div>1</div><div>In the left navigation panel of the Azure portal, click More Services, and search for Automation Accounts.</div></div> <div><div>2</div><div>In the Automation Accounts pane, click Add.</div></div> <div><div>3</div><div>In the Add Automation Account pane, input a name, select a subscription, create or use an existing resource group, and select a location for the Automation account.</div></div> <div><div>4</div><div>Click Create.</div></div>

Table 8-7 Steps to install the automation script (*continued*)

Step (in the order of sequence)	How to?
Create new Variable Type Assets to hold the parameters for running the automation.	<ol style="list-style-type: none"> 1 In the left navigation pane of the Azure portal, click More Services, and search for Automation Accounts. 2 In the Automation Accounts pane, select the automation account that you want to use. 3 In your automation account pane, select Variables. 4 Click Add to create the variables. <ul style="list-style-type: none"> ■ DirectoryDomainName - The domain name of the directory. To obtain the domain name of the directory, mouse over the top right-hand option in the Azure Portal to display a tool tip which contains an entry for the Directory Domain name. Type: String Encrypted: No ■ SubscriptionId - The ID of the Azure subscription. Type: String Encrypted: No ■ AzureSASKeyGeneratorDaemon-ClientId - The client ID of the Azure-application Service Principal. Type: String Encrypted: No ■ AzureSASKeyGeneratorDaemon-ClientSecret - The client secret of the Azure-application Service Principal. Type: String Encrypted: Yes ■ AzureSASKeyGeneratorDaemon-StorageAccountKeyVaultMapping - The variable holds the information about the storage accounts for which the SAS tokens should be regenerated and stored in what key vaults. Type: String Encrypted: No The format of the string is as follows: <pre><*storage_account_name*>,<*storage_account_name*>: <*key_vault_name*>;<*storage_account_name*>, <*storage_account_name*> <*key_vault_name*></pre> For example, <pre>prodservice,prodservicefiles,prodservicelogs: PROD-Service-SASTokens;prodinfrastorage: Infra-SASTokens</pre> The string tells the automation script job instance (when running) to do the following: <ul style="list-style-type: none"> ■ Generate SAS tokens for prodservice, prodservicefiles, and prodservicelogs storage accounts and store them in the PROD-Service-SASTokens key vault. ■ Generate SAS Tokens for prodinfrastorage storage and it in the Infra-SASTokens key vault.account.

Table 8-7 Steps to install the automation script (*continued*)

Step (in the order of sequence)	How to?
Create a Runbook	<ol style="list-style-type: none"> 1 Log on to the Azure portal. 2 In the left navigation panel of the Azure portal, click All Services, and search for Automation Accounts. 3 In the Automation Accounts pane, select the Automation account you have previously configured. 4 In your Automation Account pane, select Runbooks. 5 Click Add a Runbook. 6 In the Add Runbook pane, select Create a new Runbook. 7 In the Create Runbook pane, input an appropriate name for the Runbook, for example, AzureSASTokenGenerator, and select PowerShell Workflow Runbook as the Runbook type. Note: Ensure that you create a RunBook of Type <code>PowerShell Workflow Runbook</code>. 8 Click Create. An editor is displayed. Note: Ensure that you import the latest versions of the modules used in the automation script. If the modules are not up-to-date, some of the commandlets used in the automation script generating the SAS tokens will not be found and the script can fail.
Configure the SAS token generation script.	<ol style="list-style-type: none"> 1 In the editor, paste the provided example script file. Note that Veritas provides the script. It is distributed freely and can be modified appropriately. It is designed to function as an initial way to populate Key Vault with SAS tokens. It can be freely modified, the headers should be kept intact. The example script is available at the Veritas Support site. 2 If you want to test the script, select Test Pane from the menu bar. Note: Ensure that you test the automation script from Test Pane and proceed with the subsequent steps only after successful execution of the script. In the absence of the testing, Discovery and Scan in Information Studio can fail. 3 Click Save, and then click Publish to make it operational.

Table 8-7 Steps to install the automation script (*continued*)

Step (in the order of sequence)	How to?
Create or use an existing schedule to run the script periodically (for example, every hour).	<ol style="list-style-type: none"> On the Azure portal, in the left navigation panel, click All services and search for Automation Accounts. In the Automation Accounts pane, select the Automation account you have previously configured. In the Automation Accounts pane, select Runbooks. In the Runbooks pane, select the Runbook previously created. In your Runbooks pane, click Schedules. Click Add a Schedule. In the resulting pane, select Link a Schedule to your Runbook. Select an existing schedule or if you are creating new schedule, do the following: <ul style="list-style-type: none"> Click Create a new schedule. Enter an appropriate name and interval that suits your organization's information security policy. Click Create. After you configure the schedule, select Configure parameters and runsettings to set the duration of the validity of SAS tokens (otherwise it defaults to 525600 seconds which is 1 year). <ul style="list-style-type: none"> Enter a value in seconds for TOKENDURATION in accordance with your organization's information security policy. Click OK.

Record the Azure account details

Record the information about the Microsoft Azure account that is required when configuring your Azure Storage account in the **Connections** console as described in the table below.

Table 8-8 Data required for configuring Azure Storage account

Step	How to?
Record the Tenant ID, which is the ID of the Azure Active Directory in which you created the Azure-application Service Principal.	<ol style="list-style-type: none"> On the Azure portal, in the left navigation panel, click Azure Active Directory. Select Properties for your Azure AD tenant. From the Properties pane, copy the Directory ID. This is the value of the Tenant ID.

Table 8-8 Data required for configuring Azure Storage account (*continued*)

Step	How to?
Client ID which is the ID of the Azure Active Directory in which you created the Azure-application Service Principal.	<ol style="list-style-type: none"> 1 On the Azure portal, in the left navigation panel, click Azure Active Directory. 2 From the App registrations pane, record the Application ID of the Azure-application Service Principal you created previously. This is the value of the Client ID.
Create and record the Client Secret, which is the login secret of the Azure-application Service Principal which you create. The Application secret is passed along with the authentication request when Information Studio logs in to Azure. See Step 1 (see Create an application account).	<ol style="list-style-type: none"> 1 On the Azure portal, in the left navigation panel, click Azure Active Directory. 2 Select App registrations for your Azure Directory Instance. 3 From the App registrations pane, record the Application ID of the Azure-application Service Principal you created previously. 4 From the resulting right-hand pane, select Keys. 5 Enter a description and select an appropriate expiry that confirms to your information security policy, and click Save. 6 Record the value displayed. This is the value of the Application Secret. This is the value of the Client Secret. <p>Note: The client secret is not displayed again. Record it before continuing.</p>
Key Vault URL - URL of the Azure Key Vault in which you are storing read-only SAS Tokens. The key vault URL is used to retrieve the SAS tokens before authenticating Information Studio to Azure Storage.	<ol style="list-style-type: none"> 1 On the Azure portal, in the left navigation panel, click Key Vaults. 2 From the Key Vaults pane, select the key vault you are using to store read-only SAS tokens for Information Studio. 3 Select Properties for your key vault. 4 From the Properties pane, record the DNS name. This is the value of the key vault URL.
Azure subscription ID. The subscription ID is a GUID that uniquely identifies subscription to use Azure services.	<ol style="list-style-type: none"> 1 Log on to the Azure portal. 2 In the left navigation panel, click Subscriptions. 3 From the Subscriptions pane, record the subscription ID of the subscription you choose to be scanned.

Table 8-8 Data required for configuring Azure Storage account (*continued*)

Step	How to?
Azure environment - Identify your Azure environment. This ensures that the correct URLs are used when access Azure resources.	You can discover your Azure environment using the DNS suffix of the key vault URL:
Information about the Azure environment default to Azure. The available options are; Azure, Azure China, Azure Germany, or Azure US Government.	<ul style="list-style-type: none">■ .vault.azure.cn – Azure China■ .vault.usgovcloudapi.net – Azure US Government■ .vault.azure.net – Azure

Configuring credentials for Microsoft Azure in Information Studio

In the Connections console, add the credentials that Veritas™ Information Studio uses to connect to Microsoft Azure.

Table 8-9 New credential for Microsoft Azure

Field	Description
Display Name	Enter a logical name for the credential. It can be your tenant ID or the Azure Active Directory ID. The name you specify here helps you select the relevant credential when configuring the Microsoft Azure connection.
Description	Enter a description for the credential. This field is optional.
Tenant ID	Enter the ID of the Azure Active Directory in which you created the application. The details of the tenant ID, client ID, secret key, and the key vault URI are recorded as a part of completing the prerequisite steps.
Client ID	Enter the ID of the Azure Active Directory Application that you created.
Secret Key	Enter the login secret of the Azure Active Directory Application account which you created.
Key vault URI	Enter URL of the Azure Key Vault in which you have stored the read-only SAS Tokens. The key vault URL is used to retrieve the SAS tokens before authenticating Information Studio to Azure Storage.

Adding a Microsoft Azure connection in Information Studio

Before you set up a connection to allow applications to discover the Microsoft Azure storage accounts, you must complete all the prerequisites to authorize the access.

Table 8-10 Adding a connection for a Microsoft storage account

Field	Description
Display Name	This is a free-form field. Enter a name that Connector Framework uses to identify your Microsoft Azure account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description that you can use to identify the Azure storage account.
Subscription ID	Enter the Azure subscription ID, which is a GUID that uniquely identifies your subscription to use Azure services. You must record the subscription ID as a part of the prerequisite steps for configuring Azure from the Connections console.
Azure environment	Select your Azure environment. The available options include Azure Global, Azure China, Azure Germany, or Azure US Government.
Max Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Microsoft Azure connection is 2.
Credential	From the drop-down, select the respective credential.
Data Engine	Select the Data Engine which is responsible for scanning the Azure storage account.

Configuring Microsoft SharePoint Online connection

Microsoft SharePoint Online is a cloud-hosted connector that connects to your Office 365 account to discover the site collections, sub-sites, document libraries, files, and folders. It scans each site collection to collect summary metadata and uploads the metadata to Information Studio for analysis and geographical visualization.

Note that you do not need to add credentials in Connectors for the Microsoft SharePoint Online datastore. You are redirected to the Microsoft login page for authorization at the time of adding the connection. SharePoint Online uses the Open Authorization 2 (OAuth2) protocol to permit access to a third-party application, which is done using a Microsoft-registered application.

In the Connectors wizard, do the following:

1. Specify the **Client ID** and **Client Secret**.
2. Click **Authorize on SharePoint Online** to display the Microsoft authorization page. To scan the site collections under the mentioned SharePoint Online account, Information Studio uses the Microsoft registered application pre-created by the user(the credentials of which were mentioned in Step1).
Sharepoint Online account, Information Studio uses the
3. Log in using Microsoft global administrator credentials that the Data engine uses to gain access to the SharePoint Online account, and click **Sign In**.

Create and register an application with Microsoft

To authorize Information Studio to access the Microsoft SharePoint Online account, you must create an application for every Information Studio installation and register it with Microsoft Azure Active Directory. This step involves associating a set of credentials with the application and providing the application with the required permissions, which enables communication between Information Studio and Microsoft. This step also creates an authorization token that is stored as a named credential in the Information Studio configuration.

To create and register an application with Microsoft

- 1 Log on to [Azure portal](#).
- 2 In the left navigation panel, click **Azure Active Directory**.
- 3 From the **Azure Active Directory** pane, click **App registrations**.
- 4 In the **App registrations** pane, click **New application registration**.
- 5 On the **Create** pane, enter an appropriate name, supported account type as **Accounts in any organizational directory** (any Azure AD directory - Multi-tenant) and personal Microsoft accounts (for example, Skype, Xbox, and so on).
- 6 Select **Web** in redirect URL and specify the URL of the Information Studio Management Server. For example, <https://10.209.91.6/vcc/oauth2/callback>.
- 7 Click **Register**.
The portal assigns your application a unique Application ID. Make a note of the Application ID. You need it when configuring the Microsoft SharePoint Online account monitoring in Information Studio.
- 8 Create a secret key.
 - Navigate to **Certificates and secrets** in the left navigation pane.
 - Click **Client secrets** > **New client secret**.
Provide an appropriate description and click **Expires** > **Never**.

- Record the displayed value, which is the value of the Application Secret. The Application Secret is required when configuring the Microsoft SharePoint Online account monitoring in Information Studio.
- 9** Assign permissions to the app.
- Click **API Permissions** in the left navigation pane.
 - Click **Add permissions**, select the **Microsoft graph permissions** tab, and provide the following set of permissions under delegated and application permissions respectively.

Delegated Permissions [Add](#) [About delegated permissions](#)

Files.Read	Files.Read.All	Files.Read.Selected	Files.ReadWrite	Files.ReadWrite.All
Files.ReadWrite.AppFolder	Files.ReadWrite.Selected	offline_access	openid	People.Read
Sites.Read.All	Sites.ReadWrite.All	User.Read	User.ReadBasic.All	User.ReadWrite
Directory.AccessAsUser.All (Admin Only)	Directory.Read.All (Admin Only)	Directory.ReadWrite.All (Admin Only)		
Group.Read.All (Admin Only)	Group.ReadWrite.All (Admin Only)	People.Read.All (Admin Only)		
User.Read.All (Admin Only)	User.ReadWrite.All (Admin Only)			

Application Permissions [Add](#) [About application permissions](#)

Application.ReadWrite.All (Admin Only)	Application.ReadWrite.OwnedBy (Admin Only)	AuditLog.Read.All (Admin Only)
Directory.Read.All (Admin Only)	Directory.ReadWrite.All (Admin Only)	Domain.ReadWrite.All (Admin Only)
Files.Read.All (Admin Only)	Files.ReadWrite.All (Admin Only)	Group.Read.All (Admin Only)
Group.ReadWrite.All (Admin Only)	Sites.FullControl.All (Admin Only)	Sites.Manage.All (Admin Only)
Sites.Read.All (Admin Only)	Sites.ReadWrite.All (Admin Only)	User.Read.All (Admin Only)
User.ReadWrite.All (Admin Only)		

10 Click **Save**.

You return to the Connectors console once the authorization is complete.

Adding a Microsoft SharePoint Online connection in Information Studio

You now need to add the Microsoft SharePoint Online connector in the **Connections** console to discover SharePoint site collections.

Table 8-11 Adding a SharePoint Online account in Information Studio

Field	Description
Display Name	This is a free-form field. Enter a name that Connector Framework uses to identify your SharePoint Online account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the SharePoint Online account.
Client ID	Enter the application id of the application that you created.
Client Secret	Enter the password that you generated for the application.
Maximum Concurrent Scans	<p>Enter the maximum number of scans you want to run in parallel on this connection.</p> <p>The feature of maximum concurrent scans lets you scan multiple site collections in parallel.</p> <p>The default available value for the Microsoft SharePoint Online connection is 2.</p>
Data Engine	Select the Data Engine which is responsible for scanning the Microsoft SharePoint account.
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Microsoft SharePoint Online content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

You can now go on to assigning a policy to the connection.

Configuring data collection from Enterprise Box

For cloud content sources, such as Box and Microsoft OneDrive, you do not need to add credentials in the **Connections** console. The authentication for these content sources is controlled using the OAuth 2.0 workflow for each connector. For these content sources you are redirected to the respective tenant accounts for authorization, where you must perform an interactive registration to acquire the authentication tokens.

Configuring credentials for Box

Box uses the Open Authorization 2 (OAuth2) protocol to permit access to a third-party application. The **Connections** console uses Box Enterprise administrator

credentials to scan a Box Administrator account. The Administrator credentials authorize the **Connections** console to access the Box Enterprise account. The authorization received from Box is encrypted and saved in the **Connections** console configuration.

The credentials are used by the Information Studio Box Connector to impersonate a user account to query Box for metadata.

Note: You do not need to add credentials in the **Connections** console for the Box content source. You are redirected to the respective tenant accounts for authorization at the time of adding the connection.

In the Connectors wizard, do the following to authorize Information Studio to access Box for Enterprise account:

1. Go to **Connectors > Connections > Add Other Connection > Box for Enterprise**, and click **Authorize on Box** to display the Box authorization page.
2. Specify the Box administrator credentials that Connector Framework uses to gain access to the Box account, and click **Authorize**.
3. Click **Grant Access to Box**. This step creates an authorization token that is stored as a named credential in the **Connections** console configuration. The Information Studio application can now access the user, folder, and file metadata.

You can now add the connection details on the **Connection** tab.

Adding a Box for Enterprise Connection

Table 8-12 Adding a Box for Enterprise Connection

Field	Description
Display Name	This is a free-form field. Enter a name that the Connector framework uses to identify this connection of your Box account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the Box account.

Table 8-12 Adding a Box for Enterprise Connection (*continued*)

Field	Description
Maximum Concurrent Scans	<p>Enter the maximum number of scans you want to run in parallel on this connection.</p> <p>The feature of maximum concurrent scans lets you scan multiple Box user accounts in parallel.</p> <p>The default available value for the Box for Enterprise connection is 2.</p>
Data Engine	Select the remote Data Engine which is responsible for scanning the Box account.
Edit pause schedule	<p>Specify the hours when Connector Framework is allowed to perform a full scan of the configured Microsoft SharePoint online content source. By default, the scan is allowed all hours of the day. See "Changing a pause schedule in Connections" on page 51.</p>

You can now go on to assigning a policy to the connection.

Configuring data collection from Google Cloud Storage

The Google Cloud Storage platform lets you store your data into regional and multi-regional buckets (content sources) within your Google Cloud Storage projects. It connects to your Google Cloud Storage project, discovers and records the data locations and buckets, scans each bucket to collect metadata on the items that are discovered, and uploads the metadata to Information Studio for analysis and geographical visualization.

Prerequisites for configuring data collection from Google Cloud Storage

Before you can create a connection for Google Cloud Storage from the Connectors console, you must complete the following prerequisite steps:

1. Create a service account for your project on the Google Cloud Storage platform. This service account is used by Information Studio to gain access to the Google buckets that you want to discover. See [Creating a service account](#) for more information.
2. Generate a key for the service account. See [Creating a key for the service account](#) for more information.

Creating a service account

To create a service account

- 1 Sign in to <https://console.cloud.google.com>.
- 2 Select the project for which you want to create a service account. For example, InfoStudioTest. This is the project that contains the Google buckets that you want to discover.
- 3 Select **IAM & Admin > Service Accounts**.
- 4 Select **Add a new service account**.
- 5 On the **Create Service account** page, enter a service account name.
- 6 Set the role for the service account to **Storage > Storage Admin**.
- 7 Click **Create**.

Creating a key for the service account

To create a service account key

- 1 In the Google Cloud Platform, select the relevant service account from the list of configured service accounts. Click **Create key**.
- 2 On the **Create private key** page for *<Name of service account>*, select **JSON**, and click **Create**. This step downloads the key file to your computer. Make a note of the location where the key file is saved. The contents of the key are required when you configure the credentials to your Google Cloud Storage account.

Configuring credentials for Google Cloud Storage in Information Studio

You can either select existing credentials to configure credentials for Google Cloud Storage or configure new credentials.

The following section describes how you can configure the credentials required for Google Cloud Storage content source. Ensure that you complete the prerequisite steps before you add the credentials for Google Cloud Storage.

Enter the following details to add the Google Cloud Storage credentials:

Table 8-13 New credential for Google Cloud Storage

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Google Cloud Storage connection.

Table 8-13 New credential for Google Cloud Storage (*continued*)

Field	Description
Description	Enter a description for the credential. This field is optional.
Service Account Key	Copy and paste all the contents from the file that you downloaded while creating the service account key in creating a key for the service account.
Project ID	This field is auto-populated if the Project ID is included in the service account key JSON.

Adding a Google Cloud Storage connection in Information Studio

Before you set up a connection to allow applications to discover the Google Cloud Storage buckets and objects, you must configure the required permissions and roles to authorize the access. See [Prerequisites for configuring data collection from Google Cloud Storage](#) for more information.

Note: Currently, the Google Cloud Storage connector only supports one Google cloud Storage project per the **Connectors** console connection. Service accounts that have access to multiple projects are not supported.

Enter the following details to add a connection for Google Cloud Storage:

Table 8-14 Adding a connection for a Google Cloud Storage account

Field	Description
Name	This is a free-form field. Enter a name that Connector Framework uses to identify your Google Cloud Storage account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description that helps you identify the Google Cloud Storage account.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Google Cloud Storage connection is 2.
Data Engine	Select the Data Engine which is responsible for scanning the Google Cloud Storage bucket.

Table 8-14 Adding a connection for a Google Cloud Storage account
(continued)

Field	Description
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Google Cloud Storage content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Configuring Microsoft Exchange Online connection

Exchange Online is a cloud-hosted connector that connects to your Office 365 account or your Microsoft Exchange Online account to discover the Exchange mailboxes of users. It scans each mailbox to collect summary metadata, such as email addresses of all users for a tenant account, and uploads the metadata for analysis and geographical visualization.

Prerequisites for configuring data collection from Microsoft Exchange Online

Before you can create a connection for Microsoft Exchange Online from the Connectors console, you must complete the following prerequisite steps:

1. Perform an Azure Active Directory app registration and give appropriate permissions to the application account. . See [Create an application account](#) below for more information.
2. Assign the ApplicationImpersonation role to the user account to be used for discovery and scanning. Refer to Microsoft documentation for more details.

Create an application account

The following procedure shows the steps to create an application account for configuring data collection from Exchange Online connector:

To create and register an application with Microsoft Exchange Online

- 1 Log on to the [Azure portal](#).
- 2 In the left navigation panel, click **Azure Active Directory**.
- 3 From the **Azure Active Directory** pane, click **App registrations**.
- 4 In the **App registrations** pane, click **New application registration**.
- 5 On the **Create** pane, enter an appropriate name, select Web app / API as the Application type, and enter a sign-on URL of https://[name] where [name] is the name you choose.

- 6 Click **Create**.
- 7 From the **App registrations** pane, record the Application ID of the Exchange Online -application Service Principal you just created. This value is input into the Automation as the AzureSASKeyGeneratorDaemon-ClientId variable.
- 8 In **Settings**, click **Keys**, add **Key description = SecretKey**, then click **Save**.
Record the displayed value, which is the value of the Application Secret. This value is input into the Automation as the AzureSASKeyGeneratorDaemon-ClientSecret variable.
- 9 Go to **Settings > Required Permissions > Add**, and add the following service principals.
 - Office 365 Exchange Online
 - Microsoft Graph
 - Office 365 SharePoint Online
 - Windows Azure Service Management API
 - Office 365 Management APIs
 - Microsoft Rights Management Services

Configuring credentials for Microsoft Exchange Online in Information Studio

The Exchange Online connector requires credentials with appropriate roles and permissions to be able to connect to Exchange and discover the account metadata.

You can either use the existing credentials by selecting existing credentials from the drop-down or create new credentials.

Table 8-15 New credential for Microsoft Exchange Online

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Exchange Online connection.
Description	Enter a description for the credential. This field is optional.
User Name	Enter the user name of the service account with the ApplicationImpersonation role to be able to impersonate other users and access their mailbox details.
Password/ Confirm Password	Enter the relevant password for the service account.

Adding a connection for Microsoft Exchange Online in Information Studio

Add a connection of type Microsoft Exchange Online in Connectors console to discover Exchange Online accounts. Exchange Online uses the Open Authorization 2 (OAuth2) protocol to permit access to a third-party application.

1. Specify the **Client ID** and **Client Secret**.
2. Click **Authorize on ExchangeOnline** to display the Microsoft authorization page. To scan the Exchange Online account and the underlying site collections, Information Studio uses the pre-created Microsoft-registered applications.
3. Log in using Microsoft global administrator credentials that the Data engine uses to gain access to the Exchange Online account, and click **Sign In**.

Table 8-16 Adding an Exchange Online account in Information Studio

Field	Description
Display Name	This is a free-form field. Enter a name that Connector Framework uses to identify your Exchange Online account. The name that you enter in this field represents a content source in analysis.
Description	Enter a logical description to help you to uniquely identify the Exchange Online account.
Server	Enter the URL for your Exchange Online account.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Microsoft Exchange Online connection is 2.
Data Engine	Select the Data Engine which is responsible for scanning the Exchange account.
Client ID	The ID of the Microsoft Exchange Online account.
Client Secret	The Client secret of the Microsoft Exchange Online account.
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Microsoft Exchange Online content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

You can now go on to assigning a policy to the connection.

Configuring a Google Drive and Gmail Connection

G Suite is a collection of enterprise cloud applications such as Google Drive and Gmail accounts from Google Cloud. G Suite users can store their data within the projects residing in these applications. The Gmail account predominantly stores message and attachments, and Google Drive preserves common file types such as data files, archive files, audio, videos, and images.

The G Suite connector connects with the enterprise accounts, discovers and records the data locations, scans each user account to collect metadata on the items that are discovered, and uploads the metadata to Information Studio for analysis and geographical visualization.

However, for Gmail connectors only the total size of all emails in the Gmail account is collected after scanning and is uploaded to Information Studio for analysis.

Prerequisites for configuring data collection from Gmail and Google Drive accounts

To create a connection for Google Drive and Gmail account from the Connectors console, you must complete the following prerequisite steps:

1. Create a service account for your project on the Google Cloud Storage platform. This service account is used by Information Studio to gain access to the Google buckets that you want to discover. See [Creating a service account](#) for more information.
2. Generate a key for the service account.

Creating a service account

To create a service account

- 1 Sign in to <https://console.developers.google.com/flows/enableapi?apiid=drive> using your Google credentials.
- 2 Select the project for which you want to create a service account. For example, InfoStudioSample. If the project is not listed, then create a new project.

This project must contain the user accounts that you want to discover.
- 3 On the left pane, select **APIs & Services > Credentials**.
- 4 On the **Credentials** page, select **OAuth consent screen**.
- 5 Enter the email address and product name for OAuth protocol to work. The other fields are optional. Click **Save**.
- 6 On the **Credentials** page, click the **Credentials** tab and select **Create credentials > Service account key**.

- 7 On the **Create service account key** page, select the **service account** or **New service account** option from the drop-down. If you chose to create a new service account, make sure to enter a service account name and select the role as Owner, then click **Create**.
- 8 On the **Credentials** page, click **Manage service accounts**. The **Service accounts** page opens. This page lists the service accounts configured for your project.
- 9 Select the service account that you want to use. Under the **Options** column, click the more options icon to view the operations that you can perform. Select **Edit**.
- 10 In the **Edit service account** dialog box, verify the service account name, click **SHOW DOMAIN-WIDE DELEGATION**, and select the **Enable G Suite Domain-wide Delegation** check box. This ensures that the specified service account is authorized.
- 11 Click **Save**.
- 12 Select **JSON** to download the key file to your computer and click **Create**.

One account user can download credential JSON file for any number of times, but should not edit this file. If edited, it may result in failing discover and scan jobs.

This step downloads the key file to your computer. Make a note of the location where the key file is saved. The contents of the key are required when you configure the credentials to your G Suite account.

Configuring Client Key and Defining the Access Scope in the G Suite Console

To add a client key to Google Admin

- 1 Sign in to <https://gsuite.google.com/products/admin/> using the **Administrator** credentials.
- 2 On the Google Admin page, click the Hamburger Menu and select **Security**. Alternatively, on the **Admin Console** page, select **Security**.
- 3 On the **Security** page, click **Advanced settings > Manage API client access**.
- 4 Enter the **Client ID** in the **Client Name** field. The Client ID information is available in the JSON file that you downloaded in the previous step.
- 5 In the One or More API Scopes field, enter <https://www.googleapis.com/auth/drive.readonly>.

The definition of this API scope is to gain read-only access to file metadata and file content.

- 6 Click **Authorize**.

Configuring credentials for Google Drive / Gmail connection

You can either select existing credentials to configure credentials for Google Drive Storage or create new credentials.

Ensure that you complete the prerequisite steps before you add the credentials for Google Drive storage or Gmail.

On the Enter Credential Details pane enter the following details:

Table 8-17 New GSuite credentials for Google Drive/ Gmail

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Google Drive connection.
Description	Enter a description for the credential. This field is optional.
Domain Name	The name of the domain to which the user belongs.
Admin Email Id	Enter admin email Id for this account.
Service Account Key	Enter service account key in JSON format.

Adding a Google Drive / Gmail Connection

You can configure the discovery and scan of a Google Drive/ Gmail connection by adding the instance in Connectors console. The configuration requires you to specify certain parameters and privileges for the consumption of the connector.

Table 8-18 Add a Google Drive connection

Field	Description
Name	This is a free- form field. Enter a unique name that Connector Framework uses to identify Google drive data.
Description	Enter a logical description for the credential. The description can contain up to 1024 characters. This field is optional.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Google Drive / Gmail connection is 2.

Table 8-18 Add a Google Drive connection (*continued*)

Field	Description
Data Engine	Select the Data Engine which is responsible for scanning the Gmail or Google Drive account.
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Google Drive content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

You can now go on to assigning a policy to the connection.

Configuring Microsoft OneDrive connection

For cloud connectors, such as Box and Microsoft OneDrive, you need not add credentials in Connectors console. The authentication for these connectors is controlled using the OAuth 2.0 workflow for each connector. For these connectors, you are redirected to the respective tenant accounts for authorization, where you must perform an interactive registration to acquire the authentication tokens.

Note: Information Studio only supports business accounts for a Microsoft OneDrive content source. It does not support private OneDrive accounts.

Prerequisites for discovering OneDrive user accounts

The Information Studio Windows Connectors is able to discover only those OneDrive user accounts for which the Access files privilege is enabled in the Office365 Admin Center. You need to log into your OneDrive account at least once to be discovered.

To enable Access files permission for user accounts

- 1 Sign in to the Office 365 Admin Center with global administrator credentials.
- 2 Navigate to **Users > Active Users**.
- 3 Select the **User** that you want to discover.
- 4 Expand the OneDrive **Settings** section.
- 5 Click **Access Files**.
- 6 Repeat steps 3-5 for every user that you want to discover.

Note: You can also use a script to enable the Access files permission for all configured users at once.

Create and register an application with Microsoft

To authorize Information Studio to access the Microsoft OneDrive account, you must create an application for every Information Studio installation and register it with Microsoft Azure Active Directory. This step involves associating a set of credentials with the application and providing the application with the required permissions, which enables communication between Information Studio and Microsoft. This step also creates an authorization token that is stored as a named credential in the Information Studio configuration.

Note: You need a Microsoft account to add an application.

To create and register an application with Microsoft

- 1 Log on to [Azure portal](#).
- 2 In the left navigation panel, click **Azure Active Directory**.
- 3 From the **Azure Active Directory** pane, click **App registrations**.
- 4 In the **App registrations** pane, click **New application registration**.
- 5 On the **Create** pane, enter an appropriate name, supported account type as **Accounts in any organizational directory** (any Azure AD directory - Multi-tenant) and personal Microsoft accounts (for example, Skype, Xbox, and so on).
- 6 Select **Web** in redirect URL and specify the URL of the Information Studio Management Server. For example, <https://10.209.91.6/vcc/oauth2/callback>.
- 7 Click **Register**.

The portal assigns your application a unique Application ID. Make a note of the Application ID. You need it when configuring the Microsoft OneDrive account monitoring in Information Studio.

- 8 Create a secret key.
 - Navigate to **Certificates and secrets** in the left navigation pane.
 - Click **Client secrets > New client secret**.
Provide an appropriate description and click **Expires > Never**.
 - Record the displayed value, which is the value of the Application Secret.
The Application Secret is required when configuring the Microsoft OneDrive account monitoring in Information Studio.
- 9 Assign permissions to the app.
 - Click **API Permissions** in the left navigation pane.

- Click **Add permissions**, select the **Microsoft graph permissions** tab, and provide the following set of permissions under delegated and application permissions respectively.

Delegated Permissions
Add
About delegated permissions

Files.Read ×
Files.Read.All ×
Files.Read.Selected ×
Files.ReadWrite ×
Files.ReadWrite.All ×
Files.ReadWrite.AppFolder ×

Files.ReadWrite.Selected ×
User.Read ×
User.ReadBasic.All ×
User.ReadWrite ×

Directory.AccessAsUser.All (Admin Only) ×
Directory.Read.All (Admin Only) ×
Directory.ReadWrite.All (Admin Only) ×

Group.Read.All (Admin Only) ×
Group.ReadWrite.All (Admin Only) ×
User.Read.All (Admin Only) ×
User.ReadWrite.All (Admin Only) ×

Application Permissions
Add
About application permissions

Directory.Read.All (Admin Only) ×
Directory.ReadWrite.All (Admin Only) ×
Domain.ReadWrite.All (Admin Only) ×

Files.Read.All (Admin Only) ×
Files.ReadWrite.All (Admin Only) ×
Group.Read.All (Admin Only) ×
Group.ReadWrite.All (Admin Only) ×

User.Invite.All (Admin Only) ×
User.Read.All (Admin Only) ×
User.ReadWrite.All (Admin Only) ×

10 Click **Save**.

Configuring credentials for a Microsoft OneDrive connection

Add a connection of type Microsoft OneDrive in Connectors console to discover OneDrive accounts. OneDrive uses the Open Authorization 2 (OAuth2) protocol to permit access to a third-party application.

You need not add credentials in Connectors console for the Microsoft OneDrive connector. You are redirected to the respective tenant accounts for authorization at the time of adding the connection.

For every OneDrive account, Information Studio fetches metadata for files and folders residing on the user drives of a OneDrive account. In the Information Studio configuration, the OneDrive tenant account (your organization's OneDrive account) corresponds to a content source and the individual user accounts correspond to containers.

Adding a Microsoft OneDrive connection

Enter the following details to add a connection for OneDrive account:

Table 8-19 Adding a OneDrive Account in Connectors console

Field	Description
Name	This is a free-form field. Enter a name that Connector Framework uses to identify your OneDrive account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the OneDrive account.
Client ID	Enter the application id of the application that you created.
Client Secret	Enter the password that you generated for the application.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Microsoft OneDrive connection is 2.
Data Engine	Select the Data Engine which is responsible for scanning the OneDrive account.
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Microsoft OneDrive content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

You can now go on to assigning a policy to the connection.

Configuring a Generic S3 Cloud Storage connection

Note: The Generic S3 Cloud Storage connection is specific to AWS S3.

Information Studio visualizes the unstructured data that resides in your Amazon S3 content sources. After you set up a connection, the Connectors console gains the ability to discover the buckets and objects residing in your S3 account.

Prerequisites for adding a Generic S3 connection

Before you can create a connection for Generic S3 content source from the Connectors console, you must complete the following pre-requisite steps:

1. Go to <https://aws.amazon.com/console/>.

2. If you already have an account, click **Sign In to the Console**. Or, create a new account. Ensure that you sign in as an AWS account root user.
3. On the **AWS Management Console**, choose your account and then select **My Security Credentials**.
4. Expand the **Access keys (access key ID and secret access key)** section. Click **Create New Access Key**.
5. Click **Download Key File** to save the access key ID and secret access key to a file. This step downloads the key file to your computer. Make a note of the location where the key file is saved. The contents of the key are required when you configure the credentials to your Amazon S3 account.

Configuring credentials for Generic S3 Cloud Storage

You can either select the existing credentials to configure credentials for Generic S3 cloud storage or create new credentials.

The following section describes how you can configure the credentials required for Generic S3 cloud storage.

Ensure that you complete the prerequisite steps before you add the credentials for Generic S3 cloud storage.

Enter the following details to add the credentials for Generic S3 cloud storage:

Table 8-20 New credential for Generic S3 cloud storage

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Generic S3 cloud connection.
Description	Enter a description for the credential. This field is optional.
Provider Name	Select Amazon S3 content resource.
Identity	Enter the project associated with the credential.
Secret Key	Enter the sign-in secret of the secret access key for an AWS account.

Adding a Generic S3 connection

Enter the following details to add a connection for Generic S3 cloud storage:

Table 8-21 Adding a connection for a Generic S3 Cloud Storage account

Field	Description
Name	This is a free-form field. Enter a name that Connector Framework uses to identify your Generic S3 cloud storage account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the Generic S3 cloud storage account.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Generic S3 connection is 2.
Data Engine	Select the Data Engine which is responsible for scanning the Generic S3 cloud storage account.
Edit pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Generic S3 Cloud Storage content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Configuring Connections to on-premises content sources

The on-premises content sources are the content sources that are located locally in the organizations infrastructure. The data resides on the local hardware that is available with the organization. The Connectors console provides you the ability to add these content sources for fetching metadata that is required.

Configuring data collection from Native File Server

Native File Servers are physical storage platform used for high-volume storage, backup, and archiving of data. The Connectors console monitors the Native File Server connections for collecting the metadata.

Prerequisites for configuring data collection from Native File Server Connections

You can add these native file servers to the Connectors console for collecting the data. For Information Studio to monitor the Native File Server connections, make sure that the following requirements are completed for different device types:

NetApp Standalone

- The DNS lookup and reverse-lookup for host name of the Information Studio Windows Connectors node from the filer must work fine.

NetApp Cluster

- The DNS lookup and reverse-lookup for host name of the Information Studio Windows Connectors node from the filer must work fine.
- ONTAP version 8.2.1 or higher cluster is configured in accordance with NetApp documentation.

EMC Isilon

- Microsoft .Net Framework version 4.5 is installed on the on-premise Information Studio Windows Connectors node.
- Note the port number from the URL used to access the Isilon OneFS management console. This port number is used by connectors for discovery purposes. The default port is 8080. Ensure that this port is not blocked by the Windows firewall in the Information Studio Windows Connectors node.

Configuring credentials for Native File Servers

When you configure a connection in Connectors console, the connector accesses the connections on behalf of a user account associated with the connection. The account must have permissions to discover and scan the Native File Servers content source. The connector can effectively use the account only when the corresponding credentials are configured on Connectors console.

Table 8-22 Adding credential for Native File Servers connections

Field	Field
Display Name	This is a free-form field. Enter a unique and logical name which you can identify when configuring a Native File Server content source in Connectors console.
Description	Enter a brief description for the Native File Server datastore. You can use this description to distinguish between multiple Native File Servers content source. This is an optional field.

Table 8-22
 Adding credential for Native File Servers connections (*continued*)

Field	Field
Domain	The name of the domain to which the user belongs. This is an optional field.
Username	Enter the username for authentication. The username should belong to the user who has certain privileges on the Native File Servers connections.
Password	Enter the password.
Confirm Password	Re-enter the password for verification.

Native File Server credentials for content source type EMC Celerra or EMC VNX

For discovery of shares: The credential must belong to the EMC filer Control Station user who has administrative rights including XMLAPI v2 privilege (for example, nasadmin).

For scanning the shares: The credential must belong to the user in the domain of which the EMC filer is a part.

To scan CIFS shares successfully, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes

EMC Isilon

For discovery of shares: Requires a user account on Isilon to perform automatic discovery of CIFS shares and to list all local groups, group memberships, and local users. The connector can use a non-administrator account for this purpose.

For scanning the shares: Required for scanning of shares from the Isilon cluster. This credential belongs to the user in the domain of which the Isilon is a part.

You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes

- Read Extended Attributes
- Read Permissions

Hitachi NAS

Note: The credentials for discovery and scanning can be same.

For discovery and scan of shares: The credential must belong to a domain user with share-level READ on Hitachi NAS EVS to perform the following tasks:

- To discover shares.
- To scan the shares for metadata.

NetApp Cluster

For discovery of shares: The credential must belong to the NetApp ONTAP cluster root user who is a local user on the ONTAP cluster. Or, this credential must belong to the ONTAP cluster non-administrator user with specific privileges.

This account can be a local account or a domain account. For scanning the shares: When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part. The connector can use a non-administrator account for this purpose.

You must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

NetApp Standalone

For discovery of shares: The credential should belong to a user in the domain of which the NetApp filers are a part of. The credential can also be an account of a local user on the NetApp filer with the required privileges for performing discovery.

For scanning the shares: When scanning CIFS shares, this credential belongs to the user in the domain of which the NetApp filer is a part of. Typically, to scan CIFS shares, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs enabled for the scan credential:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

Windows File Server (CIFS)

For discovery of shares: This credential belongs to a user who has share-level READ permissions on the file server.

For scanning the shares: This credential must belong to a user with necessary share-level permissions on a Windows File Server share.

To be able to scan a Windows File Server share successfully, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

Note: To enable the connector to successfully scan the shares on a clustered Windows File Server, ensure that the scanning user has domain level permissions of Allow logon locally.

Note: For more information on the system requirements and installation of the Information Studio Windows Connectors, refer to the sections System requirements for Information Studio Hub, Remote Data Engine, and Information Studio Windows Connectors and Installing the Information Studio Windows Connectors, respectively, in the *Veritas™ Information Studio Deployment Guide*.

Preparing a non-administrator domain user on the NetApp filer

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

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

To create a non-administrator user

- 1
- Create a new role, for example testrole, using the useradmin utility on the filer.
- 2
- Add the login and API capabilities to the role.

For example,

```
useradmin role add testrole -a login-*,api-*
```

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

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

For example,

```
useradmin group add testgroup -r testrole.
```

- 4
- Add the user testdomaintestuser to testgroup.

For example,

```
useradmin domainuser add testdomain\testuser -g testgroup
```

Additional capabilities for adding a non-administrator user account

Table 8-23

Capability	Description
login-http-admin	Enables you to log into the NetApp filer and run commands. With this capability, you can discover shares.
api-system-get-ontapi-version api-system-get-version	Enables you to get the ONTAPI version number and the system version number respectively. These are required to set the login handle context properly. Also, if these capabilities are absent, you cannot execute any APIs including those required to discover shares.
api-cifs-share-list-iter-start api-cifs-share-list-iter-next api-cifs-share-list-iter-end	Used to discover shares on the NetApp filer. Absence of these capabilities can result in a failure to discover the shares. Optionally, you can add shares manually from the Connectors console.
api-volume-list-info	Used to periodically fetch size information for NetApp volumes.

Table 8-23 *(continued)*

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

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

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

To create a non-administrator user

- 1 Launch a Telnet session with the NetApp Cluster Management host.
- 2 Create a new role, for example testrole, using the `useradmin` utility on the filer.

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

```
security login role create -role testrole -cmddirname "version"  
-access all  
  
security login role create -role testrole -cmddirname "vserver  
cifs" -access readonly  
  
security login role create -role testrole -cmddirname "vserver"  
-access readonly  
  
security login role create -role testrole -cmddirname "vserver  
cifs share" -access readonly  
  
security login role create -role testrole -cmddirname "vserver  
nfs" -access all  
  
security login role create -role testrole -cmddirname "volume"  
-access readonly  
  
security login role create -role testrole  
-cmddirname "network interface show" -access readonly
```

The `network interface show` privilege automatically assigns the following privileges to the role:

```
network interface create  
network interface delete  
network interface modify
```

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

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

Preparing a non-administrator domain user on a NetApp cluster

To configure a NetApp cluster from the Connectors console, you can use an account which is not in the administrators group on the NetApp filer, but has some specific privileges. You can use the credentials of a domain user to configure Connector to monitor a NetApp cluster. These credentials are required to discover shares on the NetApp cluster.

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

To use domain user credentials to configure a NetApp cluster

- 1 Sign in using SSH to the NetApp cluster with administrator credentials. Do one of the following:

- If the NetApp cluster has a data SVM with a CIFS server that is already created, you can use that data SVM as an authentication tunnel. Use the following command:

```
security login domain-tunnel create -vserver name of data SVM
```

The following security command displays the specified authentication tunnel:

```
login domain-tunnel show
```

- If the cluster does not have a data Storage Virtual Machine (SVM) with a CIFS server created, you can use any data SVM in the cluster and join it to a domain by using the `vserver active-directory create` command. Set the `--vserver` parameter to the data SVM. Joining a data SVM to a domain does not create a CIFS server or require a CIFS license. However, it enables the authentication of users and groups at the SVM or cluster-level.

- 2 Grant a user or a group access to the cluster or SVM with the `-authmethod` parameter set to `domain`. Also, create a new role, for example `testrole`, using the `useradmin` utility on the filer.

The following command enables `<testuser>` in the `<DOMAIN1>` domain to access the cluster through SSH:

```
cluster1::> security login create -vserver cluster1  
<-user-or-group-name> <DOMAIN1\testuser> -application ontapi  
-authmethod domain -role testrole
```

Where, `<cluster1>` is the name of Admin Vserver.

Note the following:

- The value of the `-user` and `-group-name` parameter must be specified in the format `domainnameusername`, where `<domain name>` is the name of the CIFS domain server and user name is the name of the user or group you want to grant access to.
- The user group authentication supports only SSH and ONTAPI for the `-application` parameter.
- If the authentication tunnel is deleted, the directory service logon sessions cannot be authenticated by the cluster, and users and groups cannot access

the cluster. The open sessions that were authenticated before the deletion of the authentication tunnel remain unaffected.

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

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

Enables you to sign into the NetApp filer and run commands. With this capability, you can discover shares.

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

```
security login role create -role testrole -cmddirname "version"  
-access all  
  
security login role create -role testrole -cmddirname "vserver  
cifs" -access readonly  
  
security login role create -role testrole -cmddirname "vserver  
nfs" -access all  
  
security login role create -role testrole -cmddirname "vserver"  
-access readonly  
  
security login role create -role testrole -cmddirname "volume"  
-access readonly  
  
security login role create -role testrole -cmddirname "statistics"  
-access readonly  
  
security login role create -role testrole -cmddirname "network  
interface show" -access readonly
```

The network interface show privilege automatically assigns the following privileges to the role:

```
network interface create  
  
network interface modify  
  
network interface modify
```

You can optionally specify a default role such as admin/vsadmin which already has these privileges.

Creating a non-administrator user for an EMC Isilon cluster

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

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

- 1 Sign in as an Isilon administrator to the Isilon cluster CLI using SSH or Telnet.
- 2 Run the following commands:

- To create a role named imrole:

```
isi auth roles create --name imrole --description Read-only  
role for Connector
```
- To give the user the privileges to sign in to the REST API platform framework to get a list of CIFS shares and to list users and groups:

```
isi auth roles modify imrole --add-user=username@domain  
--add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI  
--add-priv-ro=ISI_PRIV_AUTH --add-priv-ro=ISI_PRIV_NETWORK
```

To configure a local user for discovery of CIFS shares

- 1 Sign in as an Isilon administrator to the Isilon cluster CLI using SSH or Telnet.
- 2 Run the following commands:

- To create a new local user called imuser

```
isi auth users create imuser --enabled yes --password xxxxxx
```
- To create a role named imrole

```
isi auth roles create --name imrole --description Read-only  
role for Connector
```
- To grant the user the privileges to Sign in to the REST API platform framework to get a list of CIFS shares and to list users and groups

```
isi auth roles modify imrole --add-user=imuser  
--add-priv-ro=ISI_PRIV_SMB --add-priv-ro=ISI_PRIV_LOGIN_PAPI  
--add-priv-ro=ISI_PRIV_AUTH --add-priv-ro=ISI_PRIV_NETWORK
```

Adding a domain user to a local group on a Hitachi NAS file server

The Connector needs a domain user with administrative privileges on Hitachi NAS EVS to perform the following tasks:

- To discover shares - This credential belongs to a user who has share-level READ permissions on the file server.
- To scan the shares for metadata - This credential must belong to a user with necessary share-level permissions on a Hitachi NAS share.

To be able to scan a Hitachi NAS file server share successfully, you must have the share-level READ permission. Additionally, the folder within the share must have the following file system ACLs:

- Traverse Folder/Execute File
- List Folder/Read Data
- Read Attributes
- Read Extended Attributes
- Read Permissions

To enable the Connector to successfully scan the shares on a Hitachi NAS file server, ensure that the scanning user has domain level permissions of Allow logon locally.

To add a domain user to a local group on the Hitachi NAS file server

- 1 Sign in using SSH to the Hitachi NAS Admin Services EVS using the manager (administrator) credentials.
- 2 Execute the following command:

```
localgroup add Administrators <domain name> /<username>
```

Adding Native File Server connections

An administrator can configure the Native File Server content sources in Connectors console. You can visualize the analysis on the Map in the Information Studio application console.

Information Studio supports the following native file server types:

- EMC Celerra or EMC VNX
- EMC Isilon
- Hitachi NAS
- NetApp Cluster
- NetApp Standalone
- Windows File Server
- Veritas File Server

Note that the fields are common for most of the Native File Server content sources. The fields that are specific to certain content sources have been exclusively mentioned.

Enter the following details to add a connection for the respective content sources:

Table 8-24 Fields in Native File Server connections panel

Field	Description
Name	This is a free-form field. Enter a unique name that the Connector console uses to identify your Native File Servers datastore. The name can contain up to 260 characters. Note that as the field name implies, this name is only for display purpose.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters.
Server	<p>Based on the datastore that you want to configure, specify the server details as appropriate:</p> <p>EMC Celerra or VNX: Enter the fully qualified domain name (FQDN) address of the CIFS server that is exported by the filer. Do not enter the IP address of the CIFS server.</p> <p>EMC Isilon: Enter the FQDN address of the Isilon cluster. It can be EMC Isilon SmartConnect Cluster name.</p> <p>Enter the EMC Isilon server Port number and select the Configuration type.</p> <p>Hitachi NAS: Enter the FQDN address of the HNAS file system EVS that you want Information Studio to monitor.</p> <p>NetApp Cluster: Enter the FQDN address of the NetApp Cluster Management host interface that is used to manage the nodes in the cluster.</p> <p>NetApp Standalone : Enter the FQDN address of the filer that you want Information Studio to monitor.</p> <p>Windows File Server: Enter the FQDN address of the filer that you want Information Studio to monitor. In case of a clustered Windows File Server, enter the cluster name of the cluster. In case of a clustered file system, make sure to select the This is a clustered windows file server check box.</p> <p>Veritas Files Server: Enter the FQDN of the filer that you want Information Studio to monitor. Ensure that the specified FQDN address is resolvable.</p>
Maximum Concurrent Scans	<p>Enter the maximum number of scans you want to run in parallel on this connection.</p> <p>The default available value for the Native File Server connection is 2.</p>

Table 8-24 Fields in Native File Server connections panel (*continued*)

Field	Description
Data Engine	<p>From the drop-down, select the Data Engine that you want to associate the connection with.</p> <p>Note: Ensure that a Windows Connectors server is registered with the Data Engine.</p>

Configuring NetApp Cluster connection

Veritas Connectors console supports CIFS shares on clustered NetApp cluster connector. In Data ONTAP Cluster-Mode (C-mode), a Storage Virtual Machine (SVM) is a logical unit within an ONTAP cluster which can contain a CIFS server and an SVM with NFS protocol enabled. An SVM facilitates data access within a cluster. It contains data volumes and one or more LIFs through which they serve data to the clients. A CIFS server within an SVM represents a filer.

Configuring credentials for NetApp Cluster connector

The below table lists the details of the parameters required to configure credentials for NetApp Cluster connector.

Table 8-25 Configuring credential for NetApp Cluster connector

Field	Description
Display Name	This is a free-form field. Enter a unique and logical name which you can identify when configuring a NetApp Cluster content source in Connectors console.
Description	Enter a brief description for the NetApp Cluster credential. This is an optional field.
Domain	<p>The name of the domain to which the user belongs. This is an optional field.</p> <p>You can add either a local NetApp filer user for Discovery or a domain account. If you are using a local user account for Discovery, then this field is optional. If you are creating credentials for scanning, then it has to be a domain account.</p>
Username	Enter the username for authentication. The username should belong to the user who has certain privileges on the NetApp Cluster connections.

Table 8-25 Configuring credential for NetApp Cluster connector (*continued*)

Field	Description
Password	Enter the password. This field allows lowercase letters, uppercase letters, numerals, and special characters (@, #, &, etc.).
Confirm Password	Re-enter the password for verification.

Enter the following details to add a connection for the respective content sources:

Table 8-26 Creating NetApp Cluster connection

Field	Description
Name	This is a free-form field. Enter a unique name that the Connectors console uses to identify your Native File Servers datastore. The name can contain up to 260 characters. Note that as the field name implies, this name is only for display purpose.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters.
Server	NetApp Cluster: Enter the FQDN address of the NetApp Cluster Management host interface that is used to manage the nodes in the cluster.
SVM Name	Storage Virtual Machine (SVM) name which is a logical unit within an ONTAP cluster which can contain a CIFS server and an SVM with NFS protocol enabled.
CIFS LIF	Interface that connects to CIFS server within an SVM which represents a filer. This can be the FQDN or IP address.
Max Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the NetApp Cluster connection is 2.
Data Engine	From the drop-down, select the group that you want to associate the connection with.

Note: CIFS LIF and NFS LIF are optional and should only be provided when SVM name and LIF names (FQDN) are different.

Configuring data collection from Microsoft SharePoint on-premises

Prerequisites for configuring Microsoft SharePoint

1. Information Studio Windows Connectors and SharePoint Server configuration:
 - Ensure that the Information Studio Windows Connectors and the SharePoint server are in the same domain.
2. Credentials used for configuration:
 - Use domain user account.
 - Make sure that the user account is a part of the farm administrator group on SharePoint server, else the discovery policy will not work.
 - Add the same account and credentials, that were used while configuring SharePoint connection in Information Studio.
3. While configuring Microsoft SharePoint, ports 80, 443, and 8331 need to be open and unused.

Complete the following steps on the SharePoint Server:

Configure the SharePoint web application policy for discovery and scan of the site collections.

1. Sign in to the SharePoint Central Administration console with Administrator credentials, and click **Application Management**.
2. Under the web applications section, click **Manage Web Applications**.
3. In the table displaying web application details, select the appropriate web application.
4. Click **User Policy**.
5. In the Policy for web application pop-up, click **Add Users**.
6. Select the appropriate zone. You can select (All Zones) if you want the user to be given permissions on all zones for the web application.
7. Click **Next**. Select the user which you want that to have full permission on the web application. Make sure to use this account as a credential to configure SharePoint in Information Studio.
8. In the **Choose Permissions** section, select **Full Control - Has full control**.
9. Specify whether this account operates as SharePoint System account. If you select the Account operates as System check box, all accesses made by this user account are recorded with the user name, SharePoint System.
10. Click **Finish**.

Configuring credentials for Microsoft SharePoint

You can use Information Studio to gain visibility into the unstructured data residing on servers running Microsoft SharePoint. The on-premises data source connector discovers and scans the following SharePoint repositories:

- Site collections
- Sites and sub-sites
- Document library - Stores documents in the .pdf, .doc, .xls, .txt, and other such file extensions.
For a given document library, the on-premises connector fetches metadata of files and folders.

Before you add a connection for SharePoint server, you must ensure the following:

- The SharePoint server 2013 or 2016 on premise is installed and running.
- InformationStudioSharePoint service is running on a Information Studio Windows Connectors machine. Make sure that Connector service port does not conflict with any other service. The default port is 8331.
- In case you want to configure visibility into a SharePoint web application, the SharePoint server machine that is added as a connection must be running with the Application Server role.
- The Information Studio Windows Connectors machine and the SharePoint server machine should be in the same domain or a trusted domain.
- The SharePoint Server is geographically close to the Connector machine.
- For successful discovery and scanning of web application, the credentials must have full control for the target web applications under the web application permission policies.

Enter the following details to add a credential for Microsoft SharePoint:

Table 8-27 Add a credential for Microsoft SharePoint

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Microsoft SharePoint connection.
Description	Enter a description for the credential. This field is optional.
Domain	Enter the name of the domain that the Microsoft SharePoint server is a part of.
User Name	Enter the account ID for the user account.

Table 8-27 Add a credential for Microsoft SharePoint (*continued*)

Field	Description
Password	Enter the password for the user account.
Confirm Password	Re-enter the password for verification.

Adding a Microsoft SharePoint connection

You can configure the Microsoft SharePoint on-premises connection by adding the instance in the Connectors console. The configuration requires you to specify certain parameters and privileges for the consumption of the connector.

Table 8-28 Adding a connection for Microsoft SharePoint

Field	Description
Name	This is a free-form field. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the SharePoint web application.
Application Server	Enter the FQDN of the Microsoft SharePoint server.
Web Application URL	Enter the SharePoint web application URL
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Microsoft SharePoint Server connection is 2.
Data Engine	Select the Data Engine corresponding to the Information Studio Windows Connectors, which will perform discovery and scanning.
Change pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Microsoft SharePoint content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Note: If the connector encounters an I18N character, the scan of the SharePoint server fails. In such a case, edit the connector_context.property file located in the Install folder.

The Microsoft SharePoint connection is now listed under the **Connections** tab on the navigation pane.

Configuring a Microsoft SQL Server connection

Before configuring the credentials for Microsoft SQL instances, make sure that the following requirements are accomplished:

- Ensure that the credential that you specify belongs to a user with system administration role, dbcreator, and server administrator role. For discovering SQL instances, the system administration role is required. To initiate scanning of the SQL instances, the dbcreator and server administrator roles are required.
- If you choose to use the Mixed Mode (SQL Server authentication and Windows authentication), ensure that you enable the mode when installing the SQL server.

Note: Since an SQL server may include multiple SQL instances, you must configure the credential that has the required privileges for each instance. If the credential does not have access to a specific instance, then the connector cannot collect the size and other information for that instance.

- The Windows Authentication mode is supported. If you opt to use this mode, ensure that the SeAssignPrimaryTokenPrivilege, SeImpersonatePrivilege, and SeIncreaseQuotaPrivilege privileges are configured on Information Studio Windows Connectors node.
- The Connector service uses certain default TCP/IP ports for internal process communication. Thus, ensure that the default ports are open.
- Ensure that the SQL Server Browser Service is running on the SQL Server machine.
- For a successful scan, all the discovered instances should be up and running.

Configuring credentials for Microsoft SQL Server

Table 8-29 Add credential details for Microsoft SQL Server

Field	Description
Name	This is a free-form field. Enter a unique and logical name which you can identify when configuring a Microsoft SQL Server instance in Connectors console. The user name and password is associated with the display name.

Table 8-29 Add credential details for Microsoft SQL Server (*continued*)

Field	Description
Description	Enter a brief description for the Microsoft SQL Server data source. It can be used to distinguish between multiple SQL servers, not instances running on the same SQL server. This is an optional field.
User Name	Enter the user name for authentication. In case you are using Mixed Mode authentication, the username should belong to the Microsoft SQL Server administration user having access to the SQL Server instances that you want to monitor For Windows authentication mode, specify the user name in the domain/ username or username@domain format. These credentials can later be referenced when you add a SQL Server connection on Connectors console.
Password	Enter the password.
Confirm Password	Re-enter the password for verification.

Adding a Microsoft SQL Server Connection

Table 8-30 Add connection details for Microsoft SQL Server

Field	Description
Name	This is a free-form field. Enter a unique and logical name Connectors console uses to identify your Microsoft SQL Server content source. The name can contain up to 260 characters.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters. This field is optional.
Server	Enter the fully qualified domain name of the Microsoft SQL Server where the instance resides. The path can contain up to 255 characters.
Instance	Specify an instance through which Connector Framework discovers the other instances within the SQL Server. For example: SQLEXPRESS1, SQLEXPRESS2.
Port	Specify the port number through which the Connectors console will connect to the instance specified above. This field is optional.

Table 8-30 Add connection details for Microsoft SQL Server (*continued*)

Field	Description
Maximum Concurrent Scans	<p>Enter the maximum number of scans you want to run in parallel on this connection.</p> <p>The default available value for the Microsoft SQL Server connection is 2.</p>
Data Engine	From the drop-down, select the Data Engine that you want to associate the content source with.
Change pause Schedule	<p>Select to add a pause schedule to allow pausing of the Discovery or Scan policy jobs for a selected time period. See “Changing a pause schedule in Connections” on page 51.</p>

Configuring an Oracle connection

The connector service obtains access to an Oracle instance by using an account that has privileges to discover and scan the underlying databases. During the scan, the connector collects the name of the database, ID, size, creation date, and modification date.

Configuring credentials for Oracle database

To enable access between the Connector and Oracle connection, make sure that the following requirements are completed:

- The InformationStudioAgentOracle service is running on the Information Studio Windows Connectors machine.
- For discovery and scanning of Oracle instances, ensure that the user has CREATE SESSION and SELECT ANY DICTIONARY privileges.

Table 8-31 Table: Add credentials for Oracle database

Field	Description
Name	This is a free-form field. Enter a unique and logical name which you can identify when configuring an Oracle instance in Connectors console. The user name and password is associated with the display name.
Description	Enter a brief description for the Oracle instance. You can use this description to distinguish between multiple instances. This is an optional field.

Table 8-31 Table: Add credentials for Oracle database (*continued*)

Field	Description
Username	Enter the username for authentication. The user name should belong to the Oracle Server administration user having access to the Oracle instances that you want to monitor.
Password	Enter the password. This field allows lowercase letters, uppercase letters, numerals, and special characters (@, #, &, etc.).
Confirm Password	Re-enter the password for verification.

Adding an Oracle database connection

You can configure the discovery and scan of an Oracle database by adding the instance in Connectors console. The configuration requires you to specify certain parameters and privileges for the consumption of the connector.

Table 8-32 Adding an Oracle database connection

Field	Description
Name	This is a free-form field. Enter a unique name that Connectors console uses to identify your Oracle content source. The name can contain up to 260 characters.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters. This field is optional.
Server	Enter the fully qualified domain name of the Oracle database server where the instance resides. The path can contain up to 255 characters.
Instance	Specify the Oracle instance that you want to monitor. For example: orc1, orcl1, orcl2, and orcl3
Port	Specify the port number for the Oracle instance. The default port number is 1521.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Oracle database connection is 2.

Table 8-32 Adding an Oracle database connection (*continued*)

Field	Description
Data Engine	From the drop-down, select the Data engine that you want to associate the content source with.
Change pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Oracle database content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Configuring a NetBackup connection

In your data centers, the NetBackup software or appliances protect the files and folders on your file servers. The metadata pertaining to these files and folders is stored in the catalog of your NetBackup master server(s). Using NetBackup as the source of metadata has the benefit of not requiring additional scans of your file servers and allows Veritas™ Information Studio to collect data efficiently, with minimum impact on your front-end application. It also provides an efficient way of gathering information from an organization’s primary content sources.

Veritas™ Information Studio facilitates the collection of metadata from NetBackup catalogs. NetBackup stores data in terms of policies and their corresponding backup images. This data is stored under one master catalog. A policy defines the backup selection and schedule information. On every run it generates a new backup image. The backup image contains the actual metadata. It contains all the files which are being backed up in one policy run plus the contents for that run. When you add a NetBackup connection in the Connectors console, the Information Studio Data Engine stores the policy name and expanded backup selection information.

During discovery, the policy backup selection is expanded and stored as content sources on the Data Engine. The Veritas™ Information Studio Data Engine periodically fetches the metadata from the most recent full backup and subsequent incremental and full backups from the content sources. It then securely uploads a subset of the collected metadata to Information Studio.

The Data Engine fetches this information for the following type of NetBackup policies:

- Standard
- MS-Windows
- NDMP
- VMware
- Hyper-V

For more information about NetBackup and NetBackup policies, see the Veritas NetBackup documentation.

Information Studio supports collection of metadata for the following NetBackup versions:

- Veritas NetBackup, 7.6.x, 7.7.x, 8.0, 8.1, 8.1.2, and 8.2.

Note: There is no support for protection plan from 8.1.2 release.

- Veritas NetBackup Appliance, 2.6 or later.

Prerequisites for adding a NetBackup connection

Before you can add a NetBackup connection, the following prerequisites steps must be complete:

- Ensure that the Data Engine is deployed for the tenant.
- Configure Data Engine's access to the master server.
- Ensure that policies have back-up client associated with them and the back-up has taken place.

Configure access to a NetBackup Master Server

Information Studio pulls metadata from the NetBackup master server. To enable Information Studio to access the master server, you must add the IP address or FQDN of the **Data Engine** to the master server. The network name for the Data Engine server must be resolvable in DNS for reverse name lookup (for example, foo.company.com resolving to 192.168.0.150).

Note: Ensure that the firewall connection on the NetBackup master server allows the Data Engine to connect via port 1566 or the firewall should be switched off.

Note that you must complete the procedure for every NetBackup master server that you want the Data Engine to connect with.

To configure communication between Information Studio and the master server

- 1 On the NetBackup Administration Console, click **NetBackup Management > Host Properties > Master Servers**.
- 2 On the **Master Servers Properties** window, click **Servers**.

- On the **Servers** pop-up, click **Additional Servers**.
- Click **Add**, and enter the IP address or FQDN of the Data Engine (default or remote) on which you want to add the connection.

This step adds the Data Engine servers to the list of servers that can access the NetBackup master server.

Adding a NetBackup connection

You can add the NetBackup Connector in the following ways:

- After selecting **Connections** in the Information Studio administration console navigation pane, the **Add NBU Connection** option is available next to **Add Other Connection**. Select the tab and proceed with the NetBackup configurations.
- The **Add NBU Connection** option is also listed in the Application Switcher under Workflows.

Following table describes the fields that are used for adding a NetBackup connection:

Table 8-33 Adding a NetBackup connection

Field	Description
Name	This is a free-form field. Enter a name that Information Studio uses to identify your NetBackup connection. The name that you enter in this field represents a data store in Information Studio application console.
Description	Enter a logical description to help you to uniquely identify the NetBackup connection.
Server	Enter the FQDN of the NetBackup master server.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for a NetBackup master connection is 1.
Days Ago	Enter the number of days you want to retract to search for valid backups for starting connector operations.
Data Engine	Select the Data Engine that is responsible for scanning the NetBackup server.

Table 8-33 Adding a NetBackup connection (*continued*)

Field	Description
Change pause schedule	Specify the hours when Information Studio is allowed to perform a full scan of the configured NetBackup data store. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Note: If catalog compression setting in NetBackup is not enabled, the Information Studio process responsible for scanning images in the NetBackup catalog causes the NetBackup filesystem to fill up before the backups are complete.

To prevent filling up of the NetBackup master server, ensure that the catalog compression setting in NetBackup is enabled. Set the **Days Ago** property for the NetBackup connection to 1 week (7 days).

About OpenText Documentum connection

OpenText Documentum is a centralized content management system that offers ease of management and flexible control over the content. Information Studio supports monitoring of Documentum data sources by providing the discovery and scanning capabilities.

Adding credential for Documentum

Table 8-34 Table: Add credentials for Documentum

Field	Description
Name	This is a free-form field. Enter a unique and logical name which you can identify the Documentum content source in Connectors console. The username and password is associated with the display name.
Description	Enter a brief description for the Documentum content source. You can use this description to distinguish between multiple Documentum datastores. This is an optional field.
User Name	Enter the username for authentication. The username should belong to the user who has administrative rights on the content source or a domain user who is part of the Administrators group on the device.

Table 8-34 Table: Add credentials for Documentum (*continued*)

Field	Description
Password/ Confirm Password	Enter the password. This field allows lowercase letters, uppercase letters, numerals, and special characters (@, #, &, etc.). Re-enter the password for verification.

Next step: Add a connection to a Documentum content source.

Adding Documentum connection

To add a Documentum content source, enter the following details on the Connectors console wizard.

Table 8-35 Adding a Documentum connection

Field	Description
Name	This is a free-form field. Enter a unique name that Connectors console uses to identify your Documentum content source. The name can contain up to 260 characters.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters.
Server	Enter the URL of the Documentum connection that you want to monitor. Typically the URL is made up of the IP address or host name of the device. For example: <a href="http://<hostname>/emc-cmis/resources">http://<hostname>/emc-cmis/resources .
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the OpenText Documentum connection is 2.
Change pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured CMIS content source. See “Changing a pause schedule in Connections” on page 51. To reset the default schedule, click the tiles in the schedule chart. The scan runs adhere to the active schedules.
Data Engine	From the drop-down, select the Data Engine that you want to use for discovering and scanning the content source.

About IBM FileNet connector

IBM FileNet is a centralized content management system that offers ease of management and flexible control over the content. Information Studio supports monitoring of FileNet content sources by providing the discovery and scanning capabilities.

Adding credential for IBM FileNet connection

Table 8-36 Table: Add credentials for IBM FileNet

Field	Description
Name	This is a free-form field. Enter a unique and logical name which you can identify the FileNet connection in Connectors console. The username and password is associated with the display name.
Description	Enter a brief description for the FileNet content source. You can use this description to distinguish between multiple FileNet content sources. This is an optional field.
User Name	Enter the username for authentication. The username should belong to the user who has administrative rights on the content source or a domain user who is part of the Administrators group on the device.
Password/ Confirm Password	Enter the password. This field allows lowercase letters, uppercase letters, numerals, and special characters (@, #, &, etc.). Re-enter the password for verification.

Next step: Add a connection to a FileNet content source.

Adding an IBM FileNet connection

To add a connection for a FileNet content source, enter the following details on the Connectors console wizard.

Table 8-37 Adding a FileNet connection

Field	Description
Name	This is a free-form field. Enter a unique name that Connectors console uses to identify your FileNet content source. The name can contain up to 260 characters.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters.

Table 8-37 Adding a FileNet connection (*continued*)

Field	Description
Server	<p>Enter the URL of the FileNet connection that you want to monitor. Typically the URL is made up of the IP address or host name of the device.</p> <p>For example: <a href="http://<hostname>/emc-cmis/resources">http://<hostname>/emc-cmis/resources.</p>
Maximum Concurrent Scans	<p>Enter the maximum number of scans you want to run in parallel on this connection.</p> <p>The default available value for the FileNet connection is 2.</p>
Change pause schedule	<p>Specify the hours when Connector Framework is allowed to perform a full scan of the configured FileNet content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51. To reset the default schedule, click the tiles in the schedule chart. The scan runs adhere to the active schedules.</p>
Data Engine	<p>From the drop-down, select the Data Engine that you want to associate the content source with.</p>

About OpenText LiveLink connector

OpenText LiveLink is a centralized content management system that offers ease of management and flexible control over the content. Information Studio supports monitoring of LiveLink content sources by providing the discovery and scanning capabilities.

Adding credential for OpenText LiveLink connection

Table 8-38 Add credentials for LiveLink

Field	Description
Name	<p>This is a free-form field. Enter a unique and logical name which you can identify when configuring a LiveLink connection in Connectors console. The username and password is associated with the display name.</p>
Description	<p>Enter a brief description for the LiveLink credential. You can use this description to distinguish multiple credentials of a specific type in Connectors console. This is an optional field.</p>

Table 8-38 Add credentials for LiveLink (*continued*)

Field	Description
User Name	Enter the username for authentication. The username should belong to the user who has administrative rights on the LiveLink content source or a domain user who is part of the Administrators group on the device.
Password/ Confirm Password	Enter the password. This field allows lowercase letters, uppercase letters, numerals, and special characters (@, #, &, etc.). Re-enter the password for verification.

Adding a LiveLink connection

To add a LiveLink content source, enter the following details on the Connectors console wizard.

Table 8-39 Adding a LiveLink connection

Field	Description
Name	This is a free-form field. Enter a unique name that Connectors console uses to identify your LiveLink connection. The name can contain up to 260 characters.
Description	Enter a logical description for the content source. The description can contain up to 1024 characters.
Server	Enter the URL of the LiveLink connection that you want to monitor. Typically the URL is made up of the IP address or host name of the device. For example: <a href="http://<hostname>/emc-cmis/resources">http://<hostname>/emc-cmis/resources .
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the LiveLink connection is 2.
Change pause Schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured LiveLink connection. By default, the scan is allowed. See “Changing a pause schedule in Connections” on page 51. To reset the default schedule, click the tiles in the schedule chart. The scan runs adhere to the active schedules.
Data Engine	From the drop-down, select the Data Engine that you want to use for discovering and scanning the content source.

Managing connections to Microsoft Exchange on-premises connector

The Microsoft Exchange connector scans the Microsoft Exchange mailbox database (content sources) and the mailboxes (repositories) to collect information about the size and count of the user, archived, and shared messages. The on-premises connector then uploads the metadata in Information Studio for analysis and visualization. Before you can configure the discovery and scan of Microsoft Exchange mailboxes, you must configure credentials (the details of the account) that the connectors use to access the Exchange Server.

As a prerequisite:

- You must provide the credentials of a service account, which is assigned the Application Impersonation role to enable the Exchange on-premises connector to impersonate users and to accesses their mailbox details.
- Ensure network connectivity between Information Studio Windows Connectors and Microsoft Exchange on-premises Server.

Configuring credentials for Microsoft Exchange

The following table describes the fields that are used for configuring credentials for Microsoft Exchange on-premises.

Table 8-40 Add a credential for Microsoft SharePoint

Field	Description
Name	Enter a logical name for the credential. The name you specify here helps you select the relevant credential when configuring the Microsoft Exchange connection.
Description	Enter a description for the credential. This field is optional.
Domain	Enter the fully qualified name of the domain that the user account is a part of.
User Name	Enter the name of the user account that will be used for discovery and scanning.
Password/ Confirm Password	Enter and confirm the password for the user account.

Adding a Microsoft Exchange connection

To enable the discovery and scan of the Microsoft Exchange mailboxes, you must add the Microsoft Exchange server using the Connectors console. Ensure that a Information Studio Windows Connectors is installed on a Windows machine. Enter the following details to add a Microsoft Exchange connection:

Table 8-41 Adding a connection for Microsoft Exchange

Field	Description
Name	This is a free-form field. Enter a name to identify your Exchange account. The name that you enter in this field represents a content source in Information Studio.
Description	Enter a logical description to help you to uniquely identify the on-premises Exchange account.
Server	Enter the FQDN of the host/server machine.
Use Kerberos authentication	Select the check box if Kerberos authentication is enabled in the Exchange environment being configured.
Domain Controller	Enter the details of the Active Directory domain controller being used by the Exchange Server. The domain controller should be a Global Catalog Server.
Maximum Concurrent Scans	Enter the maximum number of scans you want to run in parallel on this connection. The default available value for the Microsoft Exchange on-premises connection is 2.
Data Engine	From the drop-down, select the Data Engine that you want to associate the Exchange connection with.
Change pause schedule	Specify the hours when Connector Framework is allowed to perform a full scan of the configured Exchange server content source. By default, the scan is allowed all hours of the day. See “Changing a pause schedule in Connections” on page 51.

Information Studio Policies

This chapter includes the following topics:

- [Overview of Information Studio policies](#)
- [Creating a new policy](#)
- [Veritas™ Information Studio policy rules](#)
- [Managing policies](#)
- [Viewing policies](#)

Overview of Information Studio policies

Veritas™ Information Studio provides different capabilities, such as discovery, scan, and classification of an organization's on-premises content sources. The Information Studio **Policy Manager** provides an extensible framework that lets Information Studio and the deployed applications define various capabilities and expose a consistent user interface for managing those. A policy is a collection of conditions or rules that determine the scope and the behavior of the policy. A policy can specify the inclusion or exclusion rules for connectors and rules that allow the end user to filter the target data on which the action supported by the policy can be executed. For example, you may create a policy that applies to the Box for Enterprise connector, with a rule to include or exclude all PDF files that are owned by specific users.

Note: Only users with the IT Admin role can view and interact with the **Policy Manager** screen or manage the policies' permissions.

The following table describes each type of policy and lists the asset types associated with the policy.

Table 9-1 Policy type descriptions and corresponding asset type

Policy type	Description	Associated with data asset type
Discovery	Contains rules for the discovery of content sources and their hierarchy.	Connections
Scan	Contains rules for scanning metadata from connections configured in Information Studio.	Connections and Repositories

Users with IT Admin role can create, edit, and publish a policy from the **Policy Manager**.

The following table displays the different actions supported for the various policies.

Table 9-2 Policy type and corresponding supported actions

Policy type	Policy					Published jobs	
	Create	View	Publish	Edit	Delete	Unbind	Reschedule
Discovery	✓	✓	✓	✓	✓	✓	✓
Scan	✓	✓	✓	✓	✓	✓	✓

Creating a new policy

The **Policy Manager** lets you create data management policies to enable discovery and scan functions. Every policy defines a set of rules that govern the behavior of applications that are deployed on Information Studio, in terms of how each application can execute specific actions. You can configure multiple rules within a policy and apply them in the context of selected data connectors.

To create a new policy in the Information Studio administration console

- 1 In the **Policy Manager** screen, click **New Policy**.
- 2 From the **Select Policy type** drop-down, select **Discovery** or **Scan** as the case may be.

Note: Discovery and Scan are Information Studio administration policies, while Classification policy type is available in the Information Studio application.

- 3 Enter a **Policy Name** and **Policy Description** in the respective fields. Ensure that you do not enter a duplicate name for a policy. A duplicate name throws an error while saving the policy.

Policy Description is an optional field.

- 4 Select a connector.

Note: You must select at least one connector while creating a policy.

- 5 In the **Create Rules for Policy** pane, create the rules that define the application behavior for discovering or scanning the configured connectors. See [“Veritas™ Information Studio policy rules”](#) on page 117.

Note: You can either create new rules for the policy or select rules from an existing (already configured) policy. Click the **Add existing rules to policy** drop-down to select a preconfigured rule.

The following tables list the connectors that support creation of custom policies in the Information Studio administration console and the Information Studio application console.

Table 9-3 Connectors supporting creation of custom policies in the Information Studio administration console

Connector	Policy type	
	Discovery	Scan
SQL	✓	×
Microsoft Exchange	×	×
Microsoft SharePoint On-Premise	✓	×
Microsoft SharePoint Online	✓	×
Box for Enterprise	✓	×
Gmail	✓	×
Google Drive	✓	×
Microsoft Exchange Online	✓	×
Microsoft OneDrive	✓	✓

Table 9-3 Connectors supporting creation of custom policies in the Information Studio administration console (*continued*)

Connector	Policy type	
	Discovery	Scan
Google Cloud Storage	✓	×
Microsoft Azure	✓	×
Windows File Server (CIFS)	✓	✓
Native File Server EMC Celerra	✓	✓
Native File Server EMC Isilon	✓	✓
Hitachi NAS	✓	✓
IBM FileNet	✓	×
OpenText Documentum	✓	×
OpenText LiveLink	✓	×
Generic S3	✓	×
NetApp Standalone	✓	✓
NetApp Cluster	✓	✓
Oracle	×	×
NetBackup	✓	×

Veritas™ Information Studio policy rules

Policy rules are essentially predefined conditions that an application uses for executing actions on the underlying entities in a connector. You can select multiple connectors for a policy and each connector can contain multiple rules.

Basic components of a rule

A rule has this basic form:

rule type, attribute, operator, value

For example, in the rule type, "exclusion", "name" is the attribute, "Eq" is the operator, and "C\$" is the value.

The attribute is an entity to which a rule applies. Every rule type supports one or more attributes.

Some examples of attributes include, fetch native tags, site collection URL, share name, email ID, size, bucket name, instance name, etc. The attributes available for selection, however, vary with the connector applicable to a policy type.

The entries for value depend on data type. For example, if the data type is a string, you can enter any string in the values section of the rule; if the data type is a number, you can enter only numbers. For Boolean data type, you can select true or false with the help of a radio button.

To create a new rule

Note: The rules available for selection depend on the type of the policy you are creating.

- 1 Select whether you want to create an **Inclusion** or **Exclusion** rule.
- 2 Select a predefined attribute type.
- 3 Click a condition.
- 4 Enter a pattern or value.

For example, when configuring a scan policy, you can create a rule of type **Inclusion** with an attribute to **Fetch Permissions** and choose a value of **True** if you want the permissions to be fetched during the scan.
- 5 Click the plus icon to enter more patterns/ values or click the minus icon to delete a pattern/ value.
- 6 Click **Apply rule for selected connector**.
- 7 Repeat steps 1 through 6 to create another rule.
- 8 Click **Save Policy** to save the configured policy or click **Save Policy & Publish** to schedule a job for the saved policy.
- 9 View the newly created policy in the **Policies** screen.

Note: The created policy remains inactive and is not applied till it is published. Publishing a policy binds it to assets, and triggers a job at scheduled intervals to apply the policy rules to the selected connectors. See [“Managing policies”](#) on page 119.

Managing policies

A policy is a definition of the rules that applications use when executing actions. Creating a policy does not automatically trigger any action for applying these rules. To trigger an action, you must bind the policy instance to respective assets. This binding operation is also referred to as publishing a policy. When you publish a policy, a job is created as per schedule (or for immediate execution if you select **Run Now**). An application picks up this job for execution.

Note: Veritas recommends that you avoid binding multiple policies of the same type to any assets, as this may result in an unintended outcome.

In addition, you can edit or delete an existing policy.

To manage a policy (publish, edit, or delete), select the check box to the left of a listed policy, then click the vertical ellipsis at the right-end of a policy row.

Publishing a policy

Note: You can publish an inactive policy as well as re-publish an active policy.

To publish a policy

- 1 In the navigation pane on the left, click **Policy Manager**.
- 2 Select the check box of a policy listed on the **Policies** page.
Use the **Status** filter to select active or inactive policies.
- 3 At the right-end of a policy row, click the vertical ellipsis and select **Publish** from the drop-down.
- 4 On the **Publish Policy** screen:
 - Specify a name for the Job Definition in the **Specify Job Definition Name** field.
The job definition name helps in identifying your job. You can view it in the **Monitoring > Jobs** screen.
 - Select the **Assets** from the assets list under **Choose Asset(s) to Publish** and filtered by type and location. For example, infrastructure assets type or information assets (based on policy type).
- 5 In the **Add Schedule** pane, choose a frequency and start time for evaluating the policy rules.

- 6 From the **Select Scheduler Type** drop-down, select either **Run Now** or **Recurrence**.

If you select **Recurrence**, select the time span from the **Repeats Every** drop-down and the corresponding values from the respective drop-downs.

- 7 Click **Publish & Close** to publish the policy.

Note: It can take up to 24 hours for the data to appear in Information Studio. You can view the published jobs on the **Monitoring > Jobs** screen.

Editing policies

You cannot edit the two default policies, Discovery and Scan, in Information Studio, but you can edit other policies you may have created. If you edit an active policy, then the changes only apply to those jobs that are created after the policy change. All the jobs that were scheduled before this change continue as per the previous setting for that policy.

To edit a policy

- 1 In the navigation pane on the left, click **Policy Manager**.
- 2 Select a check box of a policy listed on the **Policies** screen.
- 3 At the right-end of a policy row, click the vertical ellipsis and select **Edit** from the drop-down.
- 4 Edit the rules or the connector types to which the policy is applicable.

Deleting a Policy

You can choose to delete single or multiple policies. You cannot, however, delete an active policy. To delete an active policy, you must first unbind the policy from all the assets and then proceed to delete it.

To delete a policy

- 1 On the **Policies** screen, select the check box of a listed policy.
- 2 From the **Actions** drop-down, select **Delete**.
- 3 Click **Confirm**.

The deleted policy will no longer be visible on the **Policies** page.

Viewing policies

Use the **Policy Manager** screen to view the list of configured policies, the policy details, and to publish the policies.

You can also filter the listed policies by their state or by the type of the policy. For example, if you are interested in all active scan policies, use the drop-down at the top of the **Policy Manager** screen to select the filters.

To view policies:

- 1 In the navigation pane on the left, click **Policy Manager**.
- 2 To view policy details, at the right end of the policies row, click the vertical ellipsis and select **Show details**. The policy details are displayed on the right.

The **Policy Details** pane displays the following policy attributes:

- **State** - Whether the policy is active or inactive.
- **Created by** - The application context for which the policy is created.
- **Policy Type** - The type of policy.
- **Connectors** - The list of Connectors to which the policy is applied. See [“Configuring Connections to on-premises content sources”](#) on page 82.

Note: The rules applied to a connector are listed below each connector name.

The policy details differ for active and inactive policies. An active policy displays a **Published Jobs** tab along with **Policy Details**.

Next to the **Policy Details** tab, if you click the **Published Jobs** tab, you can see the following information:

- **Job Definition Name**
- **Asset(s)**

Editing a schedule and unbinding assets

In the **Published Jobs** tab, you can unbind assets from a policy as well as edit a schedule.

To edit a schedule

- 1 In the **Published Jobs** tab, at the right-end of a job row, click the vertical ellipsis and select **Edit Schedule** from the drop-down.
- 2 From the **Select Scheduler Type** drop-down, select either **Run Now** or **Recurrence**.

Note: If you select **Recurrence**, select the time span from the **Repeats Every** drop-down and the corresponding values from the respective drop-downs.

- 3 Click **Publish**.

To unbind an asset

- 1 In the **Published Jobs** tab, at the right-end of a job row, click the vertical ellipsis and select **Unbind** from the drop-down.
- 2 In the **Confirm** modal, click **Yes** to unbind the asset from the corresponding policy.

Viewing asset details

In the **Published Jobs** tab, you can view the details of the assets that have been bound to a policy.

To view asset details

- 1 In the **Published Jobs** tab, click on the asset number square in the **Assets** column.
- 2 In the pop-up, view the **AssetName**, **Type**, and **SubType**.

Viewing policy metrics

You can view the **Policy Metrics** on the **Policy Manager** screen using a toggle option. The **Policy Metrics** provides an overview of Information Studio policies along with the jobs associated with them. The default toggle option setting can be configured in the **User Preferences** screen. You can click the refresh option to refresh or reload the metrics data. The metrics include:

- Total number of policies
- Number of **Active Policies**
- Number of **Inactive Policies**
- Number of **Published Jobs**

To view the Policy Metrics

- ◆ Use the toggle option on the top-right corner of the **Policy Manager** screen to display or hide the policies summary bar.

See [“Managing policies”](#) on page 119. and [Viewing policies](#).

Locations

This chapter includes the following topics:

- [About locations in Veritas™ Information Studio](#)

About locations in Veritas™ Information Studio

As an IT Admin, you can add a location (with its subnet information) in the **Locations** screen. A connection maps to the location using this subnet information; however, certain connectors report their own location while some inherit the location of the Data Engine. Connectors that report their own location include, EMC Celerra Cluster-Mode, EMC Isilon, NetApp, Winnas, Winnas-Cluster, Veritas NetBackup, Microsoft SharePoint on-premises, Oracle Database, Microsoft SQL Server, OpenText Documentum, Opentext LiveLink, IBM FileNet, and Microsoft Azure.

You can view the list of locations in the **Locations** screen, as well as on the **Asset Map** in the **Dashboard** screen in Information Studio administration console or the **Map** in Information Studio application console. In the Information Studio application console, IT Admins can use the **All Locations** filter at the bottom of the **Map** to filter out information for that location.

Note: The **Locations** screen is visible only to users with the IT Admin role.

To add a location

- 1 In the Veritas™ Information Studio navigation pane on the left, click **Locations** to display configured locations.
- 2 Click **Add Location**, and in the **Add Location** pane, enter the following details.

Field	Description
Display Name	A user-defined name for the location of content sources.
Country	The name of the country where content sources reside. For example, France.
Latitude/Longitude	The coordinates for the latitude and longitude specifying the exact geographical location as represented on the map display.
Address	The address of the location.
Region	The region of the location. For example, Europe.
Contact Name	Name of the executive/ IT Admin responsible for maintaining and securing the data present on content sources.
Contact Telephone	Personal/ desk number of the executive/ IT Admin responsible for maintaining and securing the data present on content sources.
Contact Email	Email ID of the executive/ IT Admin responsible for maintaining and securing the data present on content sources.
Subnets	<p>The IP subnet(s) of content sources within the organization's network.</p> <p>Note: In the absence of Subnet values, during scanning, the content source maps to the location of the data engine (default value).</p>
Notes	Extra information to help you identify content sources.

- 3 Click **Save**.

You can now view the newly-added location in the **Locations** screen.

To edit a location

- 1 In the **Locations** screen, select the check-box next to the **Name** of location you want to edit.
- 2 From the **Actions** drop-down in the far-right of the **Locations** screen, click **Edit**.
- 3 In the **Edit Location** panel, edit the required fields, and click **Save**.

You can view the edited locations in the **Locations** screen.

Within 24 hours of a successful scan, a connection (content source) maps to the appropriate location and the information for that connection (content source) shows up against that particular location in the **Map** in the Information Studio application console.

To delete a location

- 1 In the **Locations** screen, select the check-box next to the **Name** of location you want to delete.
- 2 From the **Actions** drop-down in the far-right of the **Locations** screen, click **Delete**.
- 3 To confirm the delete action, click **Okay**.

The deleted location is no more listed in the **Locations** screen or in the Asset Map on the **Dashboard** screen.

Note: The content sources mapped to the deleted location get remapped to the default location.

Logs

This chapter includes the following topics:

- [About logs](#)

About logs

The **Logs** screen allows you to download the logs of Information Studio service. Users with the role of Customer Admin can access the **Logs** screen in the Information Studio administration console.

To download and view the logs

- 1 Select a start date from when you want to view the logs from the **Choose a start date** drop-down.
- 2 Select an end date till when you want to view the logs from the **Choose an end date** drop-down.
- 3 Click **Download Logs**.

A zipped folder containing all the logs between the selected date range is downloaded on your computer. You can then extract the files and view the information in them.

Audit Logs

This chapter includes the following topics:

- [About audit logs](#)

About audit logs

Veritas™ Information Studio consists of various services on both, the **Information Studio Hub** and the **Data Engine** layer. The audit-logs services are used for logging important actions that are carried out on Information Studio. You can use the **Audit Logs** screen to view and access all the audit logs applicable to the role that you have signed-in as.

Audit logging is a critical functionality of Information Studio, which deals with business critical and sensitive customer data. Logging information for the Information Studio events like logging in and out of users, who carried out the event, creation of customers or tenants, creation of policies, and the timestamp when the events occurred can provide useful information to system administrators when needed.

The **Audit Logs** screen is available on Information Studio for all the admin roles, except for the Customer Super Admin and IT Admin roles. In the navigation pane on the left of the Information Studio administration console, click **Audit Logs** to view audit logs specific to your user roles. See “[About Information Studio organizations](#)” on page 17.

Note: You can give permissions to a user with a customized role to view audit logs.

Table 12-1

Veritas™ Information Studio role	Category of audit logs accessible to this role
Customer Admin	All the events that are generated by Customer Admin users
Tenant Admin	All the events that are generated by users created by this Tenant Admin and IT Admin for a customer account

Veritas™ Information Studio audit logging: key capabilities

- Robust framework to facilitate system-wide audit logging
- Easy availability of logs to authorized users
- Meets compliance requirements

Viewing audit logs metrics

You can view the **Audit Logs Metrics** on the **Audit Logs** screen using a toggle option. The metrics include:

- Number of **Successful Events**
- Number of **Failed Events**
- **Event Category**
- **Event Types**

Towards the right of the **Search Results** pane, there is a refresh option. You can click the refresh option to refresh or reload the metrics data.

To view the Audit Logs Metrics

- ◆ Use the toggle option on the top-right corner of the **Audit Logs** screen to show or hide the audit logs summary bar. The **Audit Logs** metrics are enabled by default.

Filters

The **Search Results** table in the **Audit Logs** screen provides various filters. Use predefined filters to filter the audit logs by various categories such as by **Event type**, **Start Date**, **End Date**, and others. If you select more than one category, the filter conditions are applied using the logical operator AND. For example, if you select the **Event type** as **Customer Management > Create**, a list of all events where a customer has been added to the Information Studio configuration are displayed.

The following table lists the supported filters:

Table 12-2 Filters supported on Information Studio

Filter	Purpose
Event Category	<p>The service that triggers an event; the event categories in the Audit Logs screen include:</p> <ul style="list-style-type: none"> ■ User Authentication ■ Tenant Management ■ Customer Management ■ Data Engine Management ■ Persistence POD Management ■ Policy Admin ■ User Management, Information Studio Hub
Event Object	<p>The object or asset on which an action takes place. For example, user or customer</p>
Event Type	<p>The type of event that the Event Category triggers; the event types in the Audit Logs screen include:</p> <ul style="list-style-type: none"> ■ User Authentication: Login, Logout ■ Tenant Management: Create, Delete, Add, Assign ■ Customer Management: Create, Delete ■ Data Engine Management: Create, Delete, Register, Unregister ■ Persistence POD Management: Register, Update, Delete ■ Policy Admin: Create, Update, Delete, Publish, Unpublish ■ User Management, Information Studio Hub ■ SMTP Configuration: Save ■ Data Disposition Request: Create, Status, Delete ■ Reporting API: Create, Delete ■ Role Management: Create, Update, Delete ■ Actions Configuration: Update <p>Note: You can view the description of the event type on hovering over the event itself in the Search Results table.</p>
Owner	<p>The user who initiates an action that results in an event</p>
Start Date	<p>Filter logs from date in "MM-dd-yyyy" format</p>
End Date	<p>Filter logs till date in "MM-dd-yyyy" format</p>
Date	<p>Timestamp of an event in "MM-dd-yyyy HH:mm:ssZ" format</p>

Table 12-2 Filters supported on Information Studio (*continued*)

Filter	Purpose
Event Result	Status of an event; for example, success, failure

Note: By default, the system displays logs from the past 7 days.

Exporting audit logs

On the **Audit Logs** screen, you can view limited number of logs (as per the **Page Size** you have set in the **User Preferences** screen).

You can set the user preferences using instructions in the About user preferences section. See [“About user preferences”](#) on page 141.

For viewing all the logs within a certain date range on a single page, you can export the logs either in a .csv or a .pdf format.

To export audit logs

- 1 In the **Audit Logs** screen, click the icon in the top-right corner of the **Search Results** table.
- 2 Select either of the following:
 - Export CSV
 - Export PDF

Monitoring Health

This chapter includes the following topics:

- [About monitoring health in Veritas™ Information Studio](#)

About monitoring health in Veritas™ Information Studio

Users with the Customer Super Admin role can monitor the health of the **Information Studio Hub** and the default Data Engine, while Tenant Admins can monitor health of the remote **Data Engines** in the Tenant. The Health screen lets you look at the overall activity and performance of your **Nodes** and **Services**. The metrics on the dashboard indicate key performance indicators of the resources such as **Memory**, **CPU**, and **Network I/O** by an individual node or service. You can easily detect a node or service that is down with the help of color-coded flags and notify the Customer Super Admin to identify and fix a critical situation.

Users accessing the **Health** screen are required to have domain knowledge to understand the data and make assumptions.

Note: Users accessing the **Health** dashboard are required to have domain knowledge to understand the data and make assumptions.

You can choose between **Information Studio Hub** (default selection) and **Veritas™ Information Studio Data Engines** as the origin with either an **Infrastructure** or **Services** (default) view. The **Services** view of **Information Studio Hub** also provides PPOD support. In the **Infrastructure** view of the **Information Studio Hub**, you can select from the **ppoddevx** and **platform** options to view services under it.

Infrastructure view

The **Infrastructure** view provides a pictorial representation of the Information Studio Infrastructure. In the upper pane, the default node is selected. The left-side panel in the upper pane displays the following details:

- The node name
- IP of the Information Studio deployment
- The status of Kubernetes cluster
- The operating system used
- The OS Version

In the lower pane, the metrics display the percentage of CPU and memory used, and the number of disk I/O operations per millisecond for the selected node.

Services view

The **Services** view, in the upper pane, provides a honeycomb representation of various services, such as **Connector Framework Services**, **Visibility Services**, **Audit Log Services**, and so on. By default, PPOD namespace is selected. The left-side panel in the upper pane displays the following details of the selected service:

- The name of the service selected
- Number of pods that are allocated and pods down
- Health status of selected service. For example:
 - CRITICAL
 - NEEDS ATTENTION
 - HEALTHY
- Legend for service categories

In the lower pane, the metrics display the **CPU** and **Memory** (in GB) usage, and the **Network I/O** (in KBps) for the selected service. For Data Bus services (Kafka), the metrics display the number of **Messages**, **Bytes in** (KBps) and **Bytes out** (KBps).

Interacting with the Health screen

On selecting a view and a node (**Infrastructure** view) or service (**Services** view), you can view the corresponding metrics in the respective lower pane. You can change the filters in each of the metrics charts. By default the filter that is applied is 30 minutes, which means the data that is displayed is from the last 30 minutes. You can change the filter to select one hour, one day, or go back as far as one week.

New data is displayed when you click on any of the time filters.

In the metrics showing the **Bytes in** and **Bytes out**, click the topics at the top of each chart. This filters the results to show the metrics of the selected topics. **Bytes in** and **Bytes out** are only applicable to the data Streaming services like KAFKA.

Interpreting the data

The data in the metrics represents the performance of the node or service over a period of time (from 30 minutes to one week). The spikes or drops are relative to the previous data point since some of the metrics do not have 0 as the base value. For example, you can see a chart with a big drop, but on inspecting the data points it shows that the drop is from 90.7 % to 90.3%. The graph shows a big drop, but the performance is barely affected.

In terms of health, absence of an icon in a node or service is interpreted as healthy (as expected). A red or amber-colored icon indicates a probable cause of the health issue. For example, an amber icon indicates that the service is not working to its full capacity. Further deterioration can lead to the issue becoming critical.

Note: The metrics are subject to user interpretation. A big drop or spike may not exactly indicate the issue with the node or service. As a user of the **Health** dashboard, you may therefore be required to have the domain knowledge.

Errors you may encounter

You can encounter one or more of the following errors while trying to view the data on the **Health** screen.

Table 13-1 Health monitoring errors

View	Error message	Cause	Solution
Infrastructure	"Unable to fetch infrastructure nodes, please try again later"	There is no data coming back from the server; this can be because: <ul style="list-style-type: none"> ■ The monitoring service is not collecting data. ■ The service is not reachable. 	Contact the system administrator.
Services	"Unable to fetch list of services, please try again later" Note: This message appears only when there is a server error.		
Infrastructure or Services	"No results found for this filter." Note: This message appears only when there is no data available.		

Legend for health icons on the services and nodes

- RED indicates critical health issues on a node or service.
- YELLOW indicates whether a service needs attention (but is not critical).

Note: The Amber icon is not displayed for nodes.

- NO ICON indicates that the node or service is healthy.

Monitoring Jobs

This chapter includes the following topics:

- [About Jobs](#)
- [Viewing jobs](#)

About Jobs

Veritas™ Information Studio allows users to create various types of policies that help enterprises manage their data analytics as per business priorities and needs. See “[Overview of Information Studio policies](#)” on page 114..

When a policy is published or bound to specific enterprise information assets and a schedule is assigned to a policy, it results in the creation of a job definition. The Information Studio job scheduling framework uses these job definitions to create specific job instances that are based on the job definition schedule.

The Information Studio administration console has a **Monitoring** menu on the navigation pane on the left. Click **Monitoring > Jobs** to access the **Jobs** screen.

Note: Only users with the IT Admin role can view and interact with the **Jobs** screen.

Viewing jobs

The **Jobs** page lists all the jobs in a descending order based on the job creation date. As soon as a policy is bound with assets and a schedule, the job definition is created. Based on the job definition, a job instance is created and queues into the system for execution.

You can view the **Job Metrics** on the **Jobs** screen using a toggle option. The default toggle option setting can be configured in the **User Preferences** screen. The metrics include:

- Total number of jobs that are created for a particular Information Studio deployment
- Number of completed jobs
- Number of queued jobs
- Number of failed jobs

In the extreme-right corner of the **Job Metrics**, there is a refresh option. You can click the refresh option to refresh or reload the metrics data.

The **Jobs** screen displays the following details:

Table 14-1 Job details

Job status	Description
Job ID	The ID of a job.
Status	Whether a job is queued, running, failed, pulled, paused, has had partial success, blocked, or notification failure. For the description of each job status, see the section on job status descriptions.
Policy Name	The name of the policy associated with the job.
Job Definition Name	The name that is generated once a policy is bound to a schedule and asset.
Priority	Priority of the job. For example, on-demand jobs have higher priority than scheduled jobs, and are executed before the latter.
Type	Type of job. For example, Discovery, Scan, Classification, or Upgrade.
Start Time	The timestamp when the job started.
End Time	The timestamp when the job ended.
Asset(s)	List of names of the assets for which the job is being monitored.

You can filter the display of jobs by **Status**, **Type**, **Priority**, and **Start Time**.

Job status descriptions

This section lists the various job statuses and their descriptions.

Table 14-2 Job status descriptions

Job status	Description
QUEUED	A job is scheduled and is in queue.
RUNNING	A job is running or is in process.
PULLED	The job state between, QUEUED and RUNNING .
FAILED	A job has failed.
COMPLETED	A job is successful and complete.
ABORTED	A job aborted by a user.
TO_BE_ABORTED	Initiation of aborting a job but not marked as aborted yet.
PAUSED	A job is currently in a paused state and can automatically resume at an appropriate time.
PARTIAL_SUCCESS	At least one of the sub-jobs is successfully complete; the rest have failed or aborted.
BLOCKED	The system has blocked a job from execution.
NOTIFICATION_FAILED	The system is unable to notify that a job is created and is as good as the job is in a failed state.

To abort a job instance

- 1 In the **Jobs** screen, click the vertical ellipsis at the right-end of a **Job ID**, and select **Abort**.
- 2 Click **Yes** to abort.

Unbind assets

You can unbind assets from a policy to un-assign a schedule for a job instance.

To unbind an asset from a policy

- 1 In the **Jobs** screen, click the vertical ellipsis at the right-end of a **Job ID**, and select **Unbind**.
- 2 Click **Yes** to unbind an asset.

Job details panel

You can click on the vertical ellipsis at the right-end of a **Job ID** and select **Show details** to open the **Job Details** panel.

The **Job Details** panel is divided into two panes.

The upper pane displays the following details.

Table 14-3 Upper pane of the Jobs Details panel

Element	Description
Type	Job type.
Job ID	The ID of a job.
Policy Name	The name of the policy associated with the job.
Status	The state the job is in.
Overall Progress	The overall progress of a Scan job. Discovery jobs do not display overall progress.
Total number of managed containers	For Discovery jobs, the total number of content sources being scanned by a job. For Classification and Scan jobs, the total number of repositories.

The **Progress Details** pane displays the following details:

Depending on the policy type, you can view the progress details (such as the jobs are in progress, have failed, or are successful) either in a donut chart or a table. For Discovery and Data owner policy types, you can view the progress details in a table while for Scan and Classification, you can view the progress in a donut chart.

Note: For Scan and Discovery policy, some of the sub-jobs may be skipped while running. The **Progress Details** pane then displays a **Skipped** state for those jobs along with the other statuses.

Table 14-4 Lower pane of the Jobs Details panel

Element	Description
Total Jobs	The list of total number of sub jobs that are scheduled (in case of Scan and Discovery policies) or files that are processed under a job (in case of Classification policies).

Table 14-4 Lower pane of the Jobs Details panel (*continued*)

Element	Description
Download Status Details	<p>Feature for downloading a report with progress details for a job. For example, for Scan and Discovery policy, you can download a complete report.</p> <p>For Classification policy, click Download Status Details to download a .csv file with error details.</p> <p>There is, however, no Download Status Details report for Data Owner policy type.</p> <p>The CSV file includes columns described in the table below.</p>

Table 14-5 Columns appearing in the CSV file

Column in the CSV file	Description
jobId	The platform job ID.
assetName	The name of the asset for which the job is listed.
jobType	The type of job, SCAN, DISCOVERY, or CLASSIFICATION.
jobStatus	Progress details for a job.
jobStartTime	The start time for a job.
jobEndTime	The time of completion for a job, if it is completed.
jobErrorId	Error ID of the failed job.
jobStatusMessage	Job status message for completion/failure of jobs as well as messages for failure of previous upload.
dataPlaneId	The ID of the Data Engine on which a job is running.
assetDataLevelId	The ID of the asset data in case of a Scan job

User Preferences

This chapter includes the following topics:

- [About user preferences](#)

About user preferences

As an active Information Studio user, you have access to provide your preferences through the Information Studio administration console. You can set the following preferences in the **User Preferences** screen.

- **Page size** - Number of rows per list view (global) or by screen (a default of 20 rows are displayed at a time on the screen).
- **Default landing page** - Radio button options for storing the last visited page to enable landing on the same page on the next sign-in. The options include:
 - **Open Last Visited Page**
 - **Open Dashboard**
- **Auto-logout** after X minutes (interval). The default auto-logout period is 15 minutes.
- **Page metrics** - show/hide by default.

To configure user preferences:

- 1 In the navigation pane on the left of the Veritas™ Information Studio administration console, click **User Preferences**.
- 2 Select **Information Studio** or the application for which you want to configure the preferences.
- 3 Enter the number of rows you want to view per page.
- 4 Select the **Default landing page**.

- 5** Enter a time for **Auto-Logout**.
- 6** Select the check box to show or hide the **Page Metrics** on every screen.
- 7** Click **Submit**.

Troubleshooting Veritas™ Information Studio

This chapter includes the following topics:

- [About troubleshooting](#)

About troubleshooting

You may encounter one or more errors that are listed in the tables below and can troubleshoot with the workaround mentioned.

Note: For additional questions, contact [Support](#).

Policies

Table 16-1 Troubleshooting errors with policies

Event	Probable cause	Action
User is not able to see policy types on the Create Policy screen	Issue with database connectivity	Contact Veritas Support
User is not able to see default or the user-created policies on Policy List screen	Issue with database connectivity	Contact Veritas Support
Connectors are not visible on the Create Policy or Edit Policy screens	Services may be down	Contact Veritas Support

Table 16-1 Troubleshooting errors with policies (*continued*)

Event	Probable cause	Action
User is unable to save a policy	Some services may be down	Contact Veritas Support
Assets are not visible while publishing a policy	Respective jobs may not have been processed as expected	Check if the connections are added properly for the connector type selected on the said policy
User is unable to publish a policy after selecting assets and adding a schedule	Some services may be down	Contact Veritas Support
Job remains in queued or running state for a long time		Contact Veritas Support with Job IDs and policy names
Job state is failed		<ul style="list-style-type: none"> ■ Check the job status and download the error report ■ Contact Veritas Support
Job Details not shown for a job or Download Errors/Full Report buttons are disabled	Services specific to that job type may be down or causing errors	Contact Veritas Support

Data engines

Table 16-2 Troubleshooting errors with data engines

Event	Probable cause	Action
User is not able to see the Data Engines link in the navigation pane on the left	User does not have sufficient permissions as only users with the IT Admin and Tenant Admin role have access to the Data Engines screen	Check that the logged-in user is a Tenant or an IT Admin
User is not able to see the New Data Engine option navigation pane on the left	User does not have sufficient permissions as only users with the IT Admin role can add a new data engine	Check the role and permissions
A newly registered data engine is not visible on the Data Engines list page	Some services may be down	Contact Veritas Support

Table 16-2 Troubleshooting errors with data engines (*continued*)

Event	Probable cause	Action
On the Setup Data Engine screen, the Country and City drop-downs are empty	Some services may be down	Contact Veritas Support
On the Setup Data Engine screen, selecting an option in the Country drop-down does not auto-populate the City drop-down	Some services may be down	Contact Veritas Support
Clicking Register Data Engine , does not download the starter.zip	Some services may be down	Contact Veritas Support with attempted on-prem data engine

Information Studio CLI

This chapter includes the following topics:

- [Using the Information Studio CLI](#)

Using the Information Studio CLI

Information Studio CLI has a restrictive reach through which users can execute a very limited set of commands meant only for a specific purpose. Some of the commands that you can execute include those for:

Command	Description
<code>system security change-password</code>	To change the password for the virtual machine
<code>manage upgrade-ui</code> <code>application-start</code>	To enable and disable the upgrade User Interface
<code>manage upgrade-ui application-stop</code>	
■ For Information Studio Hub <code>manage infostudio</code> <code>infostudio-stop</code> <code>manage infostudio</code> <code>infostudio-start</code>	To shut down and start Information Studio services
■ For Remote Data Engine <code>manage rde rde-stop</code> <code>manage rde rde-start</code>	

Changing the password for the virtual machine

After you have deployed the Information Studio Hub and the Remote Data Engine(s) (optional) using the respective wizards, follow the instructions below to change the password for the virtual machine using the Information Studio CLI.

To change the password for the virtual machine

- 1 Log into the virtual machine as the host admin using the command:

```
ssh hostadmin@<IP address of the starter virtual machine>
```
- 2 Enter the default password, Admin123.
- 3 Type the command `system security change-password` to change the password.
- 4 At the prompt to confirm if you want to change your password, enter **Yes**.
- 5 You are first prompted for your old password if one is present. Type the default password, and hit Enter.
- 6 On prompt for `newpassword`, enter a new password and hit enter.

Note: Make sure the new password has at least 8 characters and includes 1 capital letter, 1 special character, and 1 number.

- 7 On prompt for confirmation of the new password, re-enter the new password and hit enter.
- 8 On successful validation of the data entered, the new password is set and you are logged out of the virtual machine.
- 9 Log into the virtual machine with the new password.

Enabling and disabling the upgrade UI

Use the commands given below to enable and disable the upgrade UI.

To enable the upgrade UI

- 1 Log into the virtual machine using the host admin credentials.
- 2 Execute the command `manage upgrade-ui application-start`.
- 3 To access the Upgrade UI, go to <https://<IP address of the Information Studio Hub or Remote Data Engine>:8080>.

To disable the upgrade UI

- 1 Log into the virtual machine using the host admin credentials.
- 2 Execute the command `manage upgrade-ui application-stop`.

Starting and shutting down Information Studio services

Use the commands given below to start and shut down the Information Studio services.

To shut down the Information Studio Hub services

- 1 Log into the Information Studio Hub virtual machine using the host admin credentials.
- 2 Run the command `manage infostudio infostudio-stop` to stop all Information Studio Hub services.

To start the Information Studio Hub services

- 1 Log into the Information Studio Hub virtual machine using the host admin credentials.
- 2 Run the command `manage infostudio infostudio-start` to start all Information Studio Hub services.

To shut down the remote Data Engine services

- 1 Log into the remote Data Engine virtual machine using the host admin credentials.
- 2 Run the command `manage rde rde-stop` to stop all remote Data Engine services.

To start the remote Data Engine services

- 1 Log into the remote Data Engine virtual machine using the host admin credentials.
- 2 Run the command `manage rde rde-start` to start all remote Data Engine services.

Back-up and restore Information Studio

This chapter includes the following topics:

- [Backing up and restoring Information Studio](#)

Backing up and restoring Information Studio

About backing up Information Studio

This section describes the recommended process to take a back-up of the Information Studio Virtual Machine(s). You can take a snapshot of the Information Studio Virtual Machine(s) to be used by backup solutions to protect the data.

With the back-up procedure described below, you can capture all of the Information Studio application data such as:

- Configuration data such as Connections, Users, Roles, Preferences
- User-generated reports
- Item metadata and classification tags collected from content sources
- Item metadata on remote Data Engines which has not yet been sent to the Information Studio Hub

To back-up Information Studio Hub virtual machine

Note: In this case the Information Studio Hub virtual machine runs on Red Hat Enterprise Linux 7.6.

- 1 Stop all Information Studio Hub services.

- Log into the Information Studio Hub virtual machine using the host admin credentials.
 - Run the command `manage infostudio infostudio-stop` to stop all Information Studio Hub services.
- 2 Using vSphere, take a snapshot of the Information Studio Hub virtual machine.
 - 3 Start all Information Studio Hub services.
 - Log into the Information Studio Hub virtual machine using the host admin credentials.
 - Run the command `manage infostudio infostudio-start` to start all Information Studio Hub services.

To back-up one or more remote Data Engine virtual machines

Note: In this case the remote Data Engine virtual machine runs on Red Hat Enterprise Linux 7.6.

- 1 Stop all remote Data Engine services.
 - Log into the remote Data Engine virtual machine using the host admin credentials.
 - Run the command `manage rde rde-stop` to stop all remote Data Engine services.
- 2 Using vSphere, take a snapshot of the remote Data Engine virtual machine.
- 3 Start all remote Data Engine services.
 - Log into the remote Data Engine virtual machine using the host admin credentials.
 - Run the command `manage rde rde-start` to start all remote Data Engine services.

Note: If you have deployed multiple remote Data Engine virtual machines, stop remote Data Engine services from all remote Data engines, before stopping Information Studio services on the Information Studio Hub to avoid errors reported due to a loss of connectivity between the Data Engine and its assigned Information Studio Hub.

About Windows virtual machines

These virtual machines host a subset of Information Studio Connectors (CIFS and SharePoint). These do not need to be protected as they are stateless. You can,

however, back up the virtual machines using the same mechanism as above to reduce the time to rebuild the node if it were to be lost due to, say, disk failure causing loss of the VMDK.

The services running on the Window Server collect metadata from Content Sources like CIFS and SharePoint and send that metadata to their associated remote Data Engine for further processing. If one of these virtual machines were to be lost, it could be rebuilt with a snapshot, the Windows Services re-installed and any metadata that was in the process of being transferred to a Linux Data Engine would then be recollected.

Restoring Virtual Machine Snapshots

To restore just one Information Studio Hub virtual machine

- 1** Follow the VMware guidelines on how to restore the desired snapshot for the Information Studio Hub Virtual Machine.
- 2** Redo any configuration changes performed since the last snapshot was taken.
For example, any new content source Connections added since the last snapshot was taken, need to be added again

About deployments of Information Studio Hub virtual machine and one or more remote Data Engine virtual machines

- Recovering from an issue with the Information Studio Hub virtual machine
If you need to restore the Information Studio Hub virtual machine to an earlier snapshot, all the remote Data Engines must also be restored to the snapshot taken at the same time (or earlier) as the Information Studio Hub snapshot. This is required since restoring only the Information Studio Hub virtual machine will very likely lead to inconsistencies in item metadata. For example, if you restored the Information Studio Hub to a snapshot taken one week ago and did not restore any of the remote Data Engines, any item updates on content sources such as new items and classification tags that the remote Data Engines sent to the Information Studio Hub within the last week would be lost. That metadata would not be sent to the Information Studio Hub again unless the items are modified and re-scanned by the remote Data Engine, which may never happen.
- Recovering from an issue with one or more remote Data Engine virtual machines
If you need to restore one or more remote Data Engines to an earlier snapshot, you can do it without having to restore any other remote Data Engine or Information Studio Hub virtual machine snapshots.
For example, if you have two remote Data Engines, A and B, and wish to restore A, that is all that is required. You do not need to restore the Information Studio Hub virtual machine to the same snapshot as Data Engine A; the Data Engine

catches up with any configuration changes made on the Information Studio Hub. In the worst case, Data Engine A can send the same item metadata a second time to the Information Studio Hub but this is de-duplicated as it is an expected scenario.

To restore Information Studio Hub virtual machine and one or more remote Data Engine virtual machines

- 1 Follow the VMware guidelines on how to restore the desired snapshot for the Information Studio Hub Virtual Machine.

Note: The Information Studio Hub should be started before the remote Data Engines.

- 2 Power on the Virtual Machines that have been restored.

To restore Information Studio Hub virtual machine and one or more remote Data Engine virtual machines from an abrupt shutdown

If the Information Studio Hub or remote Data Engine virtual machines shut down abruptly, you can start them using the following steps.

- 1 To start Information Studio Hub, log into the Information Studio Hub virtual machine using the host admin credentials and run `manage infostudio infostudio-stop` followed by rebooting of the Information Studio Hub virtual machine.

After the machine is up, log in with the host admin credentials again and run `manage infostudio infostudio-start`.

- 2 To start the remote Data Engine, log into the remote Data Engine virtual machine using the host admin credentials and run `manage rde rde-stop` followed by rebooting of the remote Data Engine virtual machine.

After the machine is up, log in with the host admin credentials again and run `manage rde rde-start`.

Reporting API

This appendix includes the following topics:

- [Overview of reporting API](#)
- [Elasticsearch schema](#)

Overview of reporting API

The reporting API is a generic reporting framework that allows users with reporting permissions to create custom reports using REST API. For example, you can use reporting API to run complex queries such as getting distribution of file extensions (that is, number of files for each unique extension) in all folders at a given depth.

Information Studio stores scanned metadata in Elasticsearch indices. The reporting API provides you access to these indices by allowing you to run ad-hoc Elasticsearch queries on scanned metadata to derive better insights. See “[Elasticsearch schema](#)” on page 161. for details on the index mapping. Detailed description of Elasticsearch Query DSL can be found at [Query DSL](#).

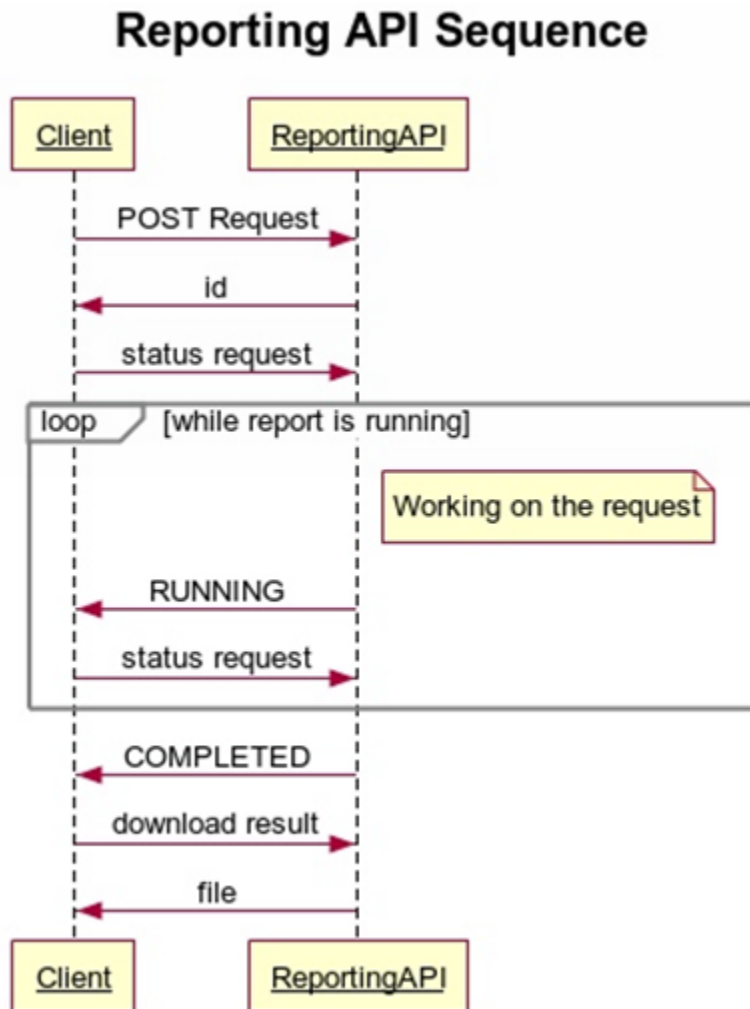
Before you use the reporting API, ensure the following:

- You have the permissions to use the reporting API
By default, only users with IT Admin role have permissions to use the reporting API. These permissions can also be assigned to any custom role. See “[Creating a custom role](#)” on page 21. for creating custom roles.
- You have the Tenant ID
Tenant ID is a UUID that you can find by logging into the Information Studio UI as a Customer Admin and navigating to **Organizations > Tenants**. You can see the Tenant ID among other tenant details on the right-hand side.

Note: Getting the access token from the IDM server is not a one-time pre-requisite, but a part of the workflow for creating a report.

- You have an access token
A valid access token authorizes the bearer to use reporting API. You can obtain the access token from the Identity Management service by providing the Tenant ID, username, password, and domain via REST call (see [Creating a new report](#)). Note that each access token has a validity of 30 minutes after which a new token has to be obtained.
Access tokens have to be provided in the 'Authorization' header of reporting APIs (see [Creating a new report](#)). If the access token is not present, the reporting API REST calls return the 401 (Unauthorized) status code. If you do not have permissions to use reporting API or if the token is not valid, the reporting API REST calls return the 403 (Forbidden) status code.
- You know the IP address of the Information Studio Hub.

Figure A-1 Reporting API sequence



The Reporting API provides the following capabilities:

- Authenticate the user creating the report
- Create a new report
- Get the status of a report
- Download a report
- Cancel or delete a report

- Get details of a given report

Getting the access token

The access token is required to authenticate the user who has permissions to run reports using the Reporting API. Each access token has a limited validity. You must get a new token after the validity expires.

Request Type: POST

URL: <https://hub-ip/idm/v1/login/token?contextId=<tenant-Id>>

Header

Content-Type: application/json

Body

```
{
  "username": <user name>,
  "password": <password>,
  "domain": <domain name>
}
```

The response of this request will contain the access token in the `accessToken` field that must be used with reporting APIs.

Sample response from the Identity Management service

```
{
    "redirectUri": "",
    "accessToken":
        "eyJhbGciOiJIUzI1NiJ9.eyJkZXBSb3ltZW50TW9kZSI6Im9uX3ByZW0iLCJyb2xlcyI6WyJpdCIhZG1pbjJdLdCjcpc3MiOiJJRE0iLCJjb250ZXh0SWQoiOiJNzU2ZGQ3MiOzZTc4LTQzMmYtYTJBMOSlhYmFlNmIxZTY4MjgiLCJwbGF0Zm9ybUlkiOiYyYTYwQlMjAtMmVhMy00NGM0LWI2N2YtMDAwMDIjYWYxZGUxIiwidXNlcikIjoizEYyZmQ4NTgtNGNkYy00ZWrmLTg2ZTMtYTU1ODkzNjRlMDNJLiwiY3VzdG9tZXJJZCI6IjBkMWUxZDQyLWQ3YzUtNDdkzi05MjZmLTFjYzgwN2EzN2RkZSIsInRlbmFudElkiOiYzclNmRkNzItM2U3OC00MzZmLWIwZjktYWJhZTZhZTIwZDI4IiwicGFydG5lcikIjoimTA0NTU3MDUtYmJlMS00ZTZQTg4ZmMtMzliMzc1NmRhZjIlIiwidXNlc1R5cGUoiOiJUZW5hbnQiLCJhcHBsaWNhdGlvbklkiOiMzZjkmZzkM2UtYjg4MS00YTdhLWF1MzItYmFlMGNI2jcyOTBiwiWiZxhwIjoxNTY2MzgwnZAzLCJhdXR0TGZ2ZWwiOiJ0ZW5hbnQiLCJpYXQiOiJlNjYzNzg5MDN9.FYchzerG8ckFlnyQxcjR2RoMkd0twOK1Mtl_TmPbdis",
    "expires_in": 1800000,
    "token_type": "bearer"
}
```

Creating a new report

Request Type: POST

URL: `https://hub-ip/v1/reports/`

Header

Content-Type: `application/json`

Authorization: `Bearer <access token>`

`tenant-id: <ID of the tenant>`

Note: The Tenant ID is visible in the Information Studio UI. Sign in to Information Studio administrator console as a Customer Admin and navigate to **Organizations > Tenants**, to view the Tenant details on the right side of the screen.

Body

```
{
  "queryObject": "{\"query\":{\"match-all\":{}}}", (Elasticsearch query as a json)
  "outputType": "csv", (csv/sqlite/json)
  "reportName": "testreport", (Using a report name is optional; Report ID is u
  "timeout": timeout-for-this-report, (Optional parameter. By default, the ti
}
```

Note the following:

- `'/'` at the end of URL is mandatory for this endpoint.
- For non-aggregation queries output type can be csv or sqlite (csv is default).
- For aggregation queries output type must be JSON.

If your report is successfully created, the response will contain the id of the report. For example:

```
{
  "id": 1234
}
```

Examples of creating new reports

- The following body creates a report that gets a file name and extension from `picoshare12` repository. The output format is `sqlite` and is sorted by file name.

```
{
  "outputType": "SQLITE",
```

```
"queryObject": {
  "sort": [
    {
      "name.keyword": "asc"
    }
  ],
  "_source": [
    "name",
    "extension"
  ],
  "query": {
    "match": {
      "repository": "picoshare12"
    }
  }
}
```

- The following body creates an aggregation report whose output format is JSON. This report lists the top 100 repositories by file count. The timeout for this report is 240 minutes.

```
{
  "timeout": "240",
  "outputType": "json",
  "reportName": "mkreport",
  "queryobject": {
    "aggs": {
      "container": {
        "terms": {
          "field": "repository.keyword",
          "size": 100,
          "order": {
            "_count": "desc"
          }
        }
      }
    }
  }
}
```

Getting the status of the report

Request Type: GET

URL: `https://hub-ip/v1/reports/{id}`

Header

Authorization: Bearer <access token>

tenant-id: <ID of the tenant>

The response body contains 2 keys, `status` and `response`. The `response` key as an additional message.

For example, the following response(s) can be expected when the report is successfully created:

```
{
  "status": "CREATED/RUNNING/COMPLETED/FAILED",
  "error": "Failed to execute query, ES not reachable" #optional
}
```

Some examples include:

Report completed successfully

```
{
  "status": "COMPLETED",
  "error": "Reporting API query successfully completed"
}
```

Report completed, but has been purged

```
{
  "status": "COMPLETED",
  "error": "Report has been purged"
}
```

Report is running

```
{
  "status": "RUNNING",
  "error": "Writing query output to file"
}
```

Report failed

```
{
  "status": "FAILED",
```

```
"error": "{ \"root_cause\": [{ \"type\": \"parsing_exception\", \"reason\": \"Unknown\", \"line\": 1, \"col\": 9 }], \"type\": \"parsing_exception\", \"reason\": \"Unknown\" }
```

Downloading results of a report

Request Type: GET

URL: `https://hub-ip/v1/reports/{id}/download`

Header

Authorization: Bearer <access token>

tenant-id: <ID of the tenant>

Returns the results of the report in response in a zip file. If the output is in .csv format, after every 50,000 records, a new file is created and all files are zipped together.

Canceling a report

Request Type: DELETE

URL: `https://hub-ip/v1/reports/{id}`

Header

Authorization: Bearer <access token>

tenant-id: <ID of the tenant>

Note that you can only cancel jobs that are in-progress.

Response 200 - Job successfully canceled is expected when the report is successfully canceled:

Validating a query

Use this endpoint to validate whether the Elasticsearch query is correct.

Request Type: POST

URL: `https://hub-ip/v1/reports/validate`

Header

Content-Type: application/json

Authorization: Bearer <access token>

tenant-id: <ID of the tenant>

Body

```
{
  "queryObject": "{ \"query\": { \"match-all\": {} } }", - Elasticsearch query as a
}
```


Returns validation response as **true** if the query is valid and response as **false** with an explanation for an invalid query.

Audit events are created whenever a report is created or canceled. See [“About audit logs”](#) on page 128. for information on and viewing audit logs.

Reports are purged if they are older than 90 days or if the disk fills up beyond a threshold (in that case, even reports that are not older than 90 days will be deleted until the disk usage comes within the threshold).

Elasticsearch schema

Information Studio stores scanned metadata in an Elasticsearch index. This index contains information about items/files.

Information Studio stores scanned metadata in Elasticsearch. Information Studio 1.1 uses Elasticsearch version 6.1.3. See [Elasticsearch query DSL](#) for more information. The following table describes fields in the Elasticsearch index.

Table A-1 Fields in the Elasticsearch index

Name	Type	Keyword field	Description
name	text	name.keyword	Name of the file
size	long (number)		Size of the file
owner	text	owner.keyword	Filesystem owner
aDate	date		Time of last access
mDate	date		Time of last modification
extension	text	extension.keyword	Extension of the file
depth	long (number)		Depth of the file relative to the share, for example, /a.txt has depth 0, whereas /folder1/folder2/a.txt has a depth of 2

Table A-1 Fields in the Elasticsearch index (*continued*)

Name	Type	Keyword field	Description
contentSource	text	contentSource.keyword	Name of the content source
repository	text	repository.keyword	Name of the repository
absName	text (keyword)		Absolute name of the file, for example, file /folder1/a.txt in content source filer1 and repository share1 will have absname \\filer1\share1\folder1\at
parentPath	path (split at /)	parentPath.keyword	Absolute path of the parent, for example, \\filer1\share1\at has parentPath /filer1/share1, whereas \\filer1\share1\folder1\at has parentPath /filer1/share1/folder1
classification	nested object		Classification information of the file (See Table A-2 , Table A-3 , and Table A-4 .)

Table A-2 Fields for classification

Name	Type	Keyword field	Description
source	text (keyword)		The classification source can be 'MICROSOFT' (for OneDrive tags) or 'VERITAS' (for files classified by the Veritas classification engine.)
classificationDate	date		Time of classification
peopleTags	text	peopleTags.keyword	Names of people in the file content Note: Ensure that Name Entity Recognition is enabled when you submit a classification request in Information Studio.
placesTags	text	placesTags.keyword	Names of places in file content Note: Ensure that Name Entity Recognition is enabled when you submit a classification request in Information Studio.
organisationTags	text	organisationTags.keyword	Names of organizations in file content Note: Ensure that Name Entity Recognition is enabled when you submit a classification request in Information Studio.

Table A-2 Fields for classification (*continued*)

Name	Type	Keyword field	Description
matchedContents	text	matchedContents.keyword	Actual content that matched some rule in the Classification Engine (Content Extraction has to be enabled during classification) Note: Ensure that content extraction is enabled when you submit a classification request in Information Studio.
rulesByTag	Nested object		Mapping of tags defined in the Classification Engine to matching rules.

Table A-3 Fields for classification tags and matching rules

Name	Type	Keyword field	Description
Name	text		Tag name
rules	Nested object		Contains Classification Engine rules that matched content and led to the corresponding tag being applied.

Table A-4 Fields for rules

Name	Type	Keyword field	Description
names	text		Rule names

Note: - Keyword fields are useful for aggregation. So if we were to compute a count of documents owned by each person, we would aggregate over `owner.keyword` instead of `owner`.

- Fields like **absname** and `[source]` are keywords themselves, so they can be used for aggregations.

For detailed information about writing Elasticsearch queries, see the [Elasticsearch reference](#).

Sample Elasticsearch queries

- Get all files with `txt` where owner is, John.

```
{
  "query": {
    "bool": {
      "must": [
        {
          "match": {
            "extension": "txt"
          }
        },
        {
          "match": {
            "owner": "john"
          }
        }
      ]
    }
  }
}
```

- Get file count per extension for repositories.

```
{
  "aggs": {
    "container": {
      "terms": {
        "field": "containerId",
        "size": 100,
        "order": {
          "_count": "desc"
        }
      }
    }
  }
}
```

```
    },  
    "aggs": {  
      "extension": {  
        "terms": {  
          "field": "extension.keyword",  
          "order": {  
            "_count": "desc"  
          }  
        }  
      }  
    }  
  }  
}
```

Getting help

This appendix includes the following topics:

- [Displaying the online help](#)
- [Using the Veritas™ Information Studio product documentation](#)

Displaying the online help

You can access the online help through the management console of Veritas™ Information Studio by clicking the question mark icon.

Using the Veritas™ Information Studio product documentation

The latest version of the Veritas™ Information Studio product documentation is available on the Veritas Services and Operations Readiness Tools (SORT) website.

<https://sort.veritas.com/documents>

You must specify the product and the platform and apply other filters for finding the appropriate document. Make sure that you are using the current version of documentation. The document version appears on page 2 of each guide. The publication date appears on the title page of each document. The documents are updated periodically for errors or corrections.

The following documents are available on the SORT site:

- *Veritas™ Information Studio Release Notes*
- *Veritas™ Information Studio Deployment Guide*
- *Administering Veritas™ Information Studio*
- *Veritas™ Information Studio User Guide*

- *Veritas™ Information Studio Software Compatibility List*
- *Veritas™ Information Studio Third-party Software Licence Agreements*