IBM Storage Sentinel anomaly scan software 1.1.5

Installation and User's Guide

Note:

Before you use this information and the product it supports, read the information in “Notices” on page 47.

This edition applies to version 1, release 1, modification 5 of IBM® Storage Sentinel (product number 5900-APZ) and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

About this publication

This publication provides overview, planning, installation, and user instructions for IBM Storage Sentinel anomaly scan software.

Audience and purpose

This publication is intended for administrators and users who are responsible for implementing a backup and recovery solution with IBM Storage Sentinel anomaly scan software in one of the supported environments.

System administrators can use this guide to help install, maintain, and start the application, manage users, and catalog resource information. Users can find procedures on how to search and browse for objects, generate and interpret reports, schedule jobs, and orchestrate backup and restore jobs.

What's new in IBM Storage Sentinel anomaly scan software 1.1.5

New features and enhancements are available in IBM Storage Sentinel anomaly scan software 1.1.5.

IBM Storage Enhancements

Feature	Description	Benefit
Support scanning volume group snapshots.	IBM Storage Sentinel supports scanning volume group snapshots. Any application data that is stored in a volume that is part of a volume group that is using the latest volume group snapshot technology can be scanned for ransomware.	Expands the IBM Storage Sentinel user base to customers running the latest volume group snapshot technology.

Chapter 1. IBM Storage Sentinel anomaly scan software overview

IBM Storage Sentinel anomaly scan software solution facilitates an end-to-end automated cyber resilience workflow that is designed to help protect copies of data, detect malicious code attacks, and enable accelerated and automated recovery of data from clean copies with IBM FlashSystem® family and SAN Volume Controller (SVC) storage.

IBM Storage Sentinel combines an IBM Storage Copy Data Management with an anomaly scan software to coordinate file & database corruption scanning with snapshot management and recovery orchestration.

IBM Storage Copy Data Management can take application-aware, immutable snapshots, commonly known as Safeguarded Copy, in IBM FlashSystems and SAN Volume Controller (SVC) storage starting with version 2.2.16.

The anomaly scan software with IBM Storage Copy Data Management and an anomaly scan engine provides scanning for corruption due to malicious code and cataloging for immutable snapshots in primary storage, enabling clients to automate recovery after an event.

Real-time cyber protection solutions are designed to protect from an attack. However, these solutions are not 100% effective and corporate data is corrupted daily. Anomaly scan software adds a layer of protection to these real time solutions and finds corruption that occurs when an attack has successfully penetrated the data center. Anomaly scan software enables early detection of issues so that IBM Storage Copy Data Management can coordinate fast application recovery, minimizing business interruption.

Anomaly scan software identifies files corrupted by malicious code using a set of statistics about files on the host being analyzed with a Machine Learning Model (MLM) trained using real world malicious codes to identify if a host was attacked by malicious code. In addition to identifying malicious code attacks, anomaly scan software checks the integrity of databases to detect corruption of the internal database data. The databases could be corrupted by an attacker, logical or physical data corruption or damage at the disk/volume level, or as a flaw in the process in the creation of a snapshot or backup of the database.

The anomaly scan software examines existing database pages and allocation tables if they exist to ensure that all the allocated database pages are present and located in their correct position. In cases where some type of page data signature is available and/or enabled by the database administrator, such as a checksum or CRC, anomaly scan software recalculates the signature based on the current page contents and verifies it against the value found in the page header. Other ancillary fields are also verified within each page depending upon the database application. The anomaly scan software Machine Learning Model (MLM) has been designed to tolerate a small amount of database corruption that is commonly observed in production database systems to avoid excessive false-positive alerts.

In addition to the anomaly scan software information that is available in IBM Documentation, other information that you might find helpful can be obtained through IBM Storage Sentinel Support.

Chapter 2. Preparation

This guide provides instructions for performing an initial installation of the IBM Storage Sentinel anomaly scan software.

This document covers installations where the server that will run the anomaly scan software does not have Internet access. Anomaly scan software now provides Linux repositories containing the packages required by the application.

Pre-installation

To ensure a smooth and successful installation, IBM Storage Sentinel anomaly scan software recommends performing a site survey prior to performing the installation. Consider the following:

- **Resources**
  Review the anomaly scan software requirements. To review, see *Server requirements and recommendations* in IBM Documentation that follow and confirm that the environment can adequately accommodate the system. If the environment cannot meet our minimum requirements, unpredictable application behavior may result.
- **License agreement**
  - The authorized user or license owner must sign into the anomaly scan software and sign the End User License Agreement (EULA). Accepting the EULA is required, otherwise, you will not be able to register the new system and activate/install the license.
- **TIP:**
  - In addition to reviewing the details that follow, see the Chapter 8, “Installation checklist,” on page 23 located at the very end of this guide. You may wish to print those two pages so you can mark off each item as you complete the steps.
  - The `CheckEngine.sh` script is available on the Passport advantage online. It should be installed in `/usr/local/bin` and made executable. The script should be run to assure the system is ready for a software installation and then to confirm the final configuration.

    Note: For steps on how to use Passport advantage online website, refer to Download Information.

Special considerations

Procedure

To avoid common mistakes during installing or upgrading procedures:

1. Ensure that you have an internet access enabled.
2. Ensure that your firewall is disabled or configured to open the ports specified in the port configuration table.
3. If you are using SLES 15.4 as your operating system, ensure that the following product modules are enabled and registered:
   - SUSE Linux Enterprise Server 15 SP4 (SLES/15.4/x86_64)
   - Base system Module (sle-module-basesystem/15.4/x86_64)
4. Ensure that you have a Mail Transfer Agent (MTA) running to enable email notifications. Set up the MTA to accept the email that the IBM Storage Sentinel anomaly scan software generates on port 25.
5. If you have multiple Network interface Cards (NIC) in your server, make sure the operating system is properly configured to use them.

6. Consistently use the hostname of your engine. For example, for host “engine1.example.com”, consistently use either the Fully Qualified Domain Name (FQDN), which is the preferred usage, "engine1.example.com" or "engine1" but not both.
7. Verify that your local host maps to the loopback address. If you need help, see “Map your local host to the loopback address” on page 4.
8. Ensure that the anomaly scan software server has Tomcat 7 or higher installed or enabled prior to installing the anomaly scan software. The GUI application depends on this service to run the application pages.

```
yum install tomcat
systemctl enable tomcat
```

Tasks prior to installation

There are several tasks that must be performed on your IBM Storage Sentinel anomaly scan software server to prepare the engine for software installation.

Hostname setting

Set your hostname on the engine. Once it has been set, you must reboot the system for the change to take effect. After the system reboot, make sure that the hostname has changed by running the `hostname` command.

Note: Consistently use the hostname of your engine (and that of other engines in a federation). For example, for host “engine1.example.com”, consistently use either the Fully Qualified Domain Name (FQDN), which is the preferred usage, "engine1.example.com" or "engine1" but not both.

Map your local host to the loopback address

Your local host must map to the loopback address on your engine. If not, edit the `/etc/hosts` file as in the example that follows.

To specify the hostname, modify the first line by adding the FQDN hostname ("myhost.domain" in this example), followed by the shortened name if you wish to use it (“myhost” in this example):

```
127.0.0.1   myhost myhost.domain localhost localhost.localdomain localhost4
localhost4.localdomain4
::1         localhost localhost.localdomain localhost6 localhost6.localdomain6
```

Important:

Do not remove the exiting localhost entries.

Do not add a second entry for 127.0.0.1

Chapter 3. IBM Storage Sentinel anomaly scan software installation

A new installation of the IBM Storage Sentinel anomaly scan software comprises of:

- installing the anomaly scan software
- registering your engine
- installing the license

After the software is installed, you will configure the system for initial use, Chapter 5, “Initial system configuration,” on page 11.

Downloading the software

Procedure

The next step is to download the engine software onto the system that it will be installed.

1. Use the Passport advantage online to download the IBM Storage Sentinel anomaly scan software.

   Note: For steps on how to use Passport advantage online website, refer to Download Information.
2. When prompted, save the anomaly scan software installation file, e.g., `indexengines-ibm-cyberv-8.1.0-1.3-suse15.x86_64.tar`, to a newly created, empty directory on the new anomaly scan software system (e.g., `/tmp/ie`). This directory should be accessible to the anomaly scan software system. Alternately, use winscp or a USB device to copy the downloaded file to the engine.

   Important: The temporary directory must be empty so that you install ONLY the .rpms for this release with the command shown in the next section.
3. Use the Fix Central to download the anomaly scan software updates.

Installing the software

Procedure

The next step is to install the IBM Storage Sentinel anomaly scan software .tar file.

Important:

The `install_ie` script will configure the following additional repositories during execution:

- For all operating systems - IBM Storage Sentinel packages for Enterprise Linux Leap 15.4 - x86_64
- For SLES 15.4 and OpenSUSE Leap 15.4:

  - SELinux (15.4) (security_SELinux)

  - Crypto applications and utilities (15.4) (security_privacy)
- For SLES 15.4 only:

  - OpenSUSE-Leap-15.4-1

  - Online updates for OpenSUSE Leap 15.4 (standard)

1. On the anomaly scan software system, login using root credentials.
2. Go to the directory containing the downloaded .tar file and unpack the .tar file.

```
indexengines-ibm-cyberv-8.1.0-1.3-suse15.x86_64.tar
```

Note: The installation will fail if the minimum disk requirements are not met. If this occurs, remediate the situation before attempting the installation again.

3. Run the command that follow to enable the permissions to execute the installation of `install_ie` script.

```
chmod 755 install_ie.sh
```

4. Next, install the anomaly scan software package.

```
./install_ie *.rpm
```

Important: The temporary directory must contain ONLY the downloaded .tar file so that you install ONLY the .rpms that were delivered with this release with the `zypper install *.rpm` command.

5. After you completed the installation on the command line, run the command that follows to enable the `iepasswd` command. User can use the `iepasswd` command to set the admin the password, for more information refer to “Setting the admin password” on page 7.

```
source /etc/profile.d/indexengines.sh
```

6. For the next steps in the process, refer to Chapter 4, “Final engine setup,” on page 7.

Chapter 4. Final engine setup

Before you can use the software, you must set the password for the admin user, accept the EULA, share your server MAC address or anomaly scan software ID to your IBM account or Business Partner representative, receive your anomaly scan software license file, and upload your license file to the anomaly scan software manager interface. For initial system configuration, see Chapter 5, “Initial system configuration,” on page 11.

Important:

1. Make sure that the contact email address is provided to your IBM account or Business Partner representative who is placing the order for anomaly scan software.
2. The MAC address of the server (also considered as the engine ID) where you intend to install the software is required to generate and provide you with the software license file. Make a note of the MAC address of the server, refer to “Finding your engine ID” on page 8.
3. Upon order, you will receive an order confirmation or proof of entitlement email from IBM. This proof of entitlement email contains your order number and can be used to claim your software license file for the anomaly scan software.
4. To claim your license file, reply to the order confirmation email received from IBM with your order number and the MAC address of the server where you intend to install the anomaly scan software.
5. After you complete the steps above, you will receive an email message from IBM with the software license file and corresponding instructions for license registration and activation.

Setting the admin password

About this task

The admin user does not have a password set at the end of the installation process. If you do not set a password for the user, then you will not be able to log into the IBM Storage Sentinel anomaly scan software GUI to proceed with the initial setup and configuration.

Procedure

1. Log in as root to the IBM Storage Sentinel anomaly scan software server using ssh.
2. On the command line, enter

   ```
   iepasswd admin
   ```

   The new password must meet the following requirements:

   - must be at least 15 characters in length
   - must contain at least one uppercase letter
   - must contain at least one lowercase letter
   - must contain at least one number
   - must contain at least one special character
   - must contain at least three characters that are not in the previous password.

   Note: These settings are also available in the anomaly scan software GUI. After setting the password so you can log into the GUI, access them by going to **Setup** > **Configuration** > **Login Settings** > **.**

What to do next

You can set the password to expire immediately and force a user to enter a new password when trying to log into the GUI by entering:

```
iepasswd -e <user>
```

If you have attempted to enter an incorrect password too many times and are now blocked from additional attempts, enter the following to unblock the user:

```
iepasswd -u <user>
```

Finding your engine ID

Procedure

The IBM Storage Sentinel anomaly scan software GUI is accessed via a web browser from the engine server.

1. Open a web browser and sign in to the anomaly scan software. In the URL bar, use https and type the hostname (preferred) or IP address of the host where the anomaly scan software system is installed. The **Sign In** page opens.

   Note: If you are unable to access this page, verify that ports 80 and 443 are open on the firewall.
2. In **Username**, type the default administrator login of `admin`. In the **Password** field, type the default password of `admin`. Select **Sign In** button. Since this is the first time you have logged in to the engine, the **Upload License** page opens.
3. The license upload screen should appear once you login to the GUI for the first time. If not, you can find the license upload screen at **Administration** > **System** > **Licenses** > **Upload License** page. On the **Upload License** screen, you will see the engine ID listed which appears in the title bar above the **Browse** button.

Completing the application licensing

Procedure

Once you have received the software license file, follow the steps below to install the license on your system.

Attention: If you are in a federation, upload a license only to the manager machines. Member machines share the license from the manager.

1. Save the email attachment that contains the license file to the system that will be used to access the IBM Storage Sentinel anomaly scan software web browser interface.
2. Open a web browser and sign in to the anomaly scan software. In the **URL** bar, use `https` and type the hostname (preferred) or IP address of the host where the anomaly scan software system is installed. The **Sign In** page opens.

   Note: If you are unable to access this page, verify that ports 80 and 443 are open on the firewall.
3. In **Username**, type the default administrator login of `admin`. In **Password**, type the default password of `admin`. Select **Sign In**.
4. After you sign in, the **Upload License** page opens. If not go to **Administration** > **System** > **Licenses** > **Upload License** page.
5. Use the **Browse** button to locate the license file you saved on the system. Then select **Upload**.

   If you have any questions or need support, contact through IBM Storage Sentinel Support.

Results

License installation in the manager engine is now complete.

Chapter 5. Initial system configuration

Now that the IBM Storage Sentinel anomaly scan software has been installed, the manager engine registered, and the license uploaded to it, the final step is to perform some initial system configuration. These initial configuration steps include:

- add member engine(s) to a federation
- add an index
- select an index
- setting the indexing service defaults
- creating an automatic merge job

Federation setup

A federation consists of two or more IBM Storage Sentinel anomaly scan software systems joined together to share work. Associating multiple engines in a federation has many advantages. For example, federated engines can perform parallel processing to accomplish tasks more quickly, such as indexing huge data stores. Federation Members share access to the data in the currently selected index. Segments of the shared index can reside on any of the federated engines, and searches, extractions, and archiving can be performed globally to all engines sharing the index.

Note: When you install the anomaly scan software for the first time its default role is manager and if it is used alone, it will continue to be a manager. In order to create a federation, you must choose a server to act as the federation manager, install the anomaly scan software on it, and then add a license. Now you can install new member anomaly scan software servers and connect them to the appropriate manager machines. Member machines share the license from the manager, as they do not receive a standalone license for installation. With the proper license any anomaly scan software server can take on the manager or member role as required. You can join, remove, and re-target the anomaly scan software servers as the requirements of the environment change. The main focus of this topic is on providing step-by-step instructions for creating federations.

Important: All engines in a federation must run the same release of the anomaly scan software.

Becoming federation manager

To become the manager of a federation, the other engine(s) must join yours. Confirm the federation status on the **Administration** > **Home** > **Index Manager** > **Joined Engines** page.

As the manager, your engine hosts the index environment for the members of the federation.

Note: The **Remove** button on the **Joined Engines** tab provides a way to terminate a member's participation in the federation if the engine is permanently disconnected ("retired") from the federation and cannot leave the federation in the usual way (by selecting **Leave Federation** on the **Setup** tab of the member engine).

Becoming a member of a federation

Joining a federation provides the ability to connect to the remote indexes on the manager system. You will not be able to access local indexes on your engine until you leave the federation. Before joining, deselect your index:

Warning: Before leaving the federation or moving a member to another federation you must clear all the infections on both the database server and IBM Storage Sentinel anomaly scan software to avoid the failure of jobs. For steps refer to “Stop Reporting an infection on the IBM Storage Sentinel anomaly scan software” on page 32.

Note: As a member of a federation, your engine will inherit any licenses from the manager engine.

1. Go to **Administration** > **Home** > **Index Manager** > **Select Index**.
2. Select **Deselect**.
3. Next, select **Setup**.
4. Select **Join Federation**.
5. Enter the fully qualified domain name of the host or the IP address of the Federation Manager system, i.e., the IBM Storage Sentinel anomaly scan software system that holds the index you wish to access. Select **Submit**.
6. The **Index Manager Setup** tab confirms your status within the federation.

   Important: If you encounter any issues when attempting to join a federation, be sure that the ports required for operating in a federated environment are open; see Server requirements and recommendations.

The indexes on the manager system are now available to you on the **Select Index** tab. Any local indexes that you may have created before joining the federation become available again if you leave the federation.

Note: The **Locally Managed Objects** tab becomes available on this page when your engine is a federation member. The information provided there is needed when it's time to renew your system license.

License operation between IBM Storage Sentinel anomaly scan software federation manager and member servers

To make sure that the proper license operation between federation manager and members do the steps that follow:

1. Update the manager to the latest version, refer to Chapter 6, “Upgrading the IBM Storage Sentinel anomaly scan software,” on page 19.
2. License the manager with an enabled inheritable license, refer to “Completing the application licensing” on page 8.
3. Update the member to the latest version, refer to Chapter 6, “Upgrading the IBM Storage Sentinel anomaly scan software,” on page 19.
4. Add member to the federation manager, refer to “Becoming a member of a federation” on page 11.
5. Do the steps “1” on page 12 thru “4” on page 12 again to add the members to the manager.

Leaving a federation

Procedure

To leave as a member of a federation:

1. Go to **Administration** > **Home** > **Index Manager** > **Select Index** and deselect the current index on the **Select Index** tab.
2. Next, go to **Administration** > **Home** > **Index Manager** > **Setup**.
3. Finally, select **Leave Federation**, and select **OK** to confirm.

Results

This reverts the system back to its Manager role (manager of its own engine) and provides access to any local indexes created on your engine prior to joining the federation.

Adding an index

Procedure

1. On the manager engine, go to **Administration** > **Home** > **Index Manager**.

2. Select **Add Index**. In this example, we will add a new index called “Index-Test01”.
3. Type a name for the new index and select **Submit**. A confirmation message appears, indicating that your new index was added.

Selecting an index

Procedure

After adding the new index, you must select it for use on the member engine. Use the **Select Index** tab to select the index.

1. On the member engine, select **Select Index** (**Administration** > **Home** > **Index Manager** > **Select Index**).
2. Choose the index you wish to access and click **Select**. Once complete, a message appears, confirming your selection.

Indexing service defaults

This tab lets you change the Indexing Service defaults (previously called "Indexing Service Options"). Your changes affect subsequent indexing operations; existing index segments are not affected.

Access the indexing service default options by navigating to **Administration** > **Home** > **Indexing Services** > **Defaults**.

Important: Make sure to set the indexing services defaults as per the table below before doing the first IBM Storage Sentinel anomaly scan software analysis.

Setting	Selected Option	Description
Indexing Mode	Full Content	**Full Content**, the default mode, indexes all metadata plus the content of supported file types, provided the associated file filters are not disabled. On the **Search** > **Results** page, you can highlight a desired file in the search results, and then drag to resize the **Review** panel (previously named the **Reconstruction** panel) to view details about the selected file such as user and security information and metadata. You can "Click-through" the link (by clicking the name of the file) on the **Review** panel's **Metadata** tab to download the file for viewing. You can also get a plain text view of the file's text contents by clicking the **Review** panel's **Text** tab.
Case Handling	Map to Lowercase	Map all upper case letters to lower case. This feature specifies the case of keywords within documents and their metadata are handled in the selected indexing mode.

Setting	Selected Option	Description
Anomaly scan software Data Collection	Enabled	Gather the file information needed to calculate cyber analytics statistics. This applies the option to all file server indexing jobs. The anomaly scan software feature uses analytics, machine learning, and forensic tools to discover successful cyberattacks on your data. The feature requires anomaly scan software licensing.
Synthetic Incremental	Enabled	Ignore unchanged files in backup images if those files were previously indexed. This feature can dramatically improve the performance of indexing operations.
		When an engine performs indexing in a federation with **Synthetic Incremental indexing** enabled, the engine queries existing anomaly scan software segments, regardless of where they are located-even if the files were indexed on another engine in the federation-to determine if files in backup images should be indexed.
		The decision of whether or not to index a file is based on the filename and path, its modification time, change time, and size. If any of this information has changed, then the file is considered modified and it is indexed as normal. If, on the other hand, all of this information matches the prior version, then the file is marked as being present in the backup image, the processing of this file is skipped, and processing moves on to the next file.
Use Change Time for Incremental Indexing	Disabled Enabled	Controls whether modifications to a file's Change Time results in the file being indexed in an incremental CIFS, NFS, or Local File Server indexing job. This option does not apply to indexing within backup files using the **Synthetic Incremental** option. Set this option to the appropriate value for your setup.

Setting	Selected Option	Description
Base incremental indexing decisions on	`Content Currently in this Index`	Index a file based on the contents of the current index. This selection provides these advantages: • The index can span across engines in a federation, so incremental jobs can now be run from any engine in the federation instead of being restricted to running on the same engine every time. • If an indexing job spans across segments, the job can now use the new segment immediately instead of waiting for the old segment to close. • If a new index is selected, the job will index all files. • If segments are deleted, the files that were in those segments will be indexed again the next time the job runs. • If files are deleted from the index, those files will be indexed again the next time the job runs.

Segment merge

The merge segments operation takes one or more index segments as input and replaces them with a single segment. Segment merge provides these primary benefits:

- It can combine two or more segments into one, reducing the amount of disk space required by combining common elements, and reducing the time it takes to search the index, which is sensitive to the number of segments.
- By default, the output segment is logically equivalent to the inputs. Options exist to omit documents that have been marked for deletion when copying data from the input segments to the output segment. This saves space and completes the disposition process, where unwanted data is actually removed from the index.
- While the segment merge process is often done using multiple segments, merging a single index segment is fully supported and is practical when the segment contains documents to be deleted.
- Purge any empty backup sets that were a result of segment merge processes, which frees up available space on the system.

The process of merging segments is covered in the following sections:

1. Set up Segment merge options, refer to “Segment merge options” on page 15
2. Perform automatic merging, refer to “Automatic merging” on page 16

Segment merge options

To view the segment merge options, go to **Administration** > **Home** > **Index Segments** > **Merge Options**.

Choose segment merge options as described in the following table.

Policy	Selected Option	Description
Job Retention	**Purge After 30 Days**	Retains the history and status of completed jobs before the jobs are automatically purged after 30 days.
Automatically Pause	**Yes**	Automatically pauses the merge when indexing or post-processing operations are running, reducing contention for engine resources.
Policy	**Purge Deleted**	Removes objects marked for deletion. When objects are deleted via the anomaly scan software search interface, they immediately become hidden
Days to Keep Deactivated Content	**30**	The number of days to retain deactivated content in the newly merged segment. This option automatically purges deactivated files after the specified period of time, thereby reducing the size of the index. A file in the index is deactivated either when the file is backed up subsequently or if the file is found to have been deleted when backing up that location again. Use this option to specify the number of days that the history of a particular file is kept. For example, IBM Storage Sentinel anomaly scan software customers typically don't care about retaining a long history of a file, while e-discovery customers typically do not want to purge any historical versions of a file.

Automatic merging

Procedure

To perform an automatic segment merge:

1. Go to **Administration** > **Home** > **Index Segments** > **Merge Segments**.
2. To automatically merge segments, click **HERE**. The **Add Job Definition** form opens.
3. Complete the **Add Job Definition** form as described in the following table.

   You can choose to run the schedule daily or weekly. In this example, we chose weekly: a Monday when no backups are created on Sundays.

   Set the Beginning date to a date 30 days after the first full analysis job is started to avoid segment merges from occurring during what might be a long first pass.

Field	Selected Option	Function
Description	Example: auto-mergejob	Provide a description for the auto-merge job in this optional field.
Remove Source Segments Upon Successful Merge	**Yes**	Removes the original segments when the merge successfully completes.
Apply Merge Policy Option	**Yes**	Applies the policy option you selected from the drop-down on the **Merge Options** page.
Purge Empty Backups	**Yes**	Will purge backups that would be kept even if they have no remaining files in them after applying policy.
Maximum Merge Time	**Do Not Pause**	Specifies that the system should not pause the merge job if it does not complete in the selected amount of time (hours or days).

Field	Selected Option	Function
Minimum Segments	**2**	The merge process is run only if there are at least two segments.
Minimum Segment Age	**720**	A segment will only be merged if it is the minimum segment age or older, i.e., if it was completed at least the **Minimum Segment Age** hours ago.
Schedule Type	**Run Weekly**	Specify that this merge job will be performed on a weekly basis. Additional fields may be required, depending on the selected schedule type. You will be prompted for these as necessary.

4. When completed, select **Submit** to define the merge job.

Ready for operation

The system is now ready for operation. Each time you log in to the IBM Storage Sentinel anomaly scan software GUI, check the status of your system. The **Engine Status** page provides this information.

Once you begin running your indexing jobs, it is important to set merge options and to create a merge job definition. See “Segment merge” on page 15 for more information.

Chapter 6. Upgrading the IBM Storage Sentinel anomaly scan software

This section provides instructions on upgrading the IBM Storage Sentinel anomaly scan software through the GUI. For new installations, see Chapter 3, “IBM Storage Sentinel anomaly scan software installation,” on page 5.

Note: Read the following sections in Chapter 2, “Preparation,” on page 3:

- *Server requirements and recommendations* in IBM Documentation
- “Special considerations” on page 3
- “Tasks prior to installation” on page 4

They contain additional information and steps that you will need to perform to ensure a smooth upgrade. This includes tips on server requirements, what to do if your host system does not have Internet access, and how to install the package dependencies from the repository bundle downloaded from the Passport advantage online.

Procedure to upgrade the software

Procedure

It is recommended that you do what is necessary to allow Analyze and Merge jobs to complete before performing an upgrade. Please follow these steps to update the IBM Storage Sentinel anomaly scan software:

Important: It is recommended that you check the Fix Central for repository bundle updates before updating. This is especially the case when upgrading between major releases (e.g., 7.10.0 to 7.12.0) as there may be new dependencies. If you choose not to do this step, you may encounter dependency errors while updating and you will need to download and install the new repository bundle.

1. Download the desired .tar file from the Passport advantage online.

   Note: For steps on how to use Passport advantage online website, refer to Download Information.
2. Log in to the host system via the command line interface using the root credentials.
3. Change directories to the directory containing the new .tar file.
4. Untar the file:

   ```
   indexengines-ibm-cyberv-7.11.0-1.25.patch_ibm.3-el7.x86_64.tar
   ```
5. Execute the following command to install the software:

   ```
   yum install *.rpm
   ```

   Important: If members show a loss of communication with the manager or if members report a license issue, you might need to unjoin and rejoin with the manager, restart the member, and reselect the original index to resume work.

Chapter 7. Commands to manage the IBM Storage Sentinel anomaly scan software services

The IBM Storage Sentinel anomaly scan software “dservice” command is found in the `/opt/ie/bin` directory. The anomaly scan software installation adds this directory to your path so that this command can be easily executed, however, it does not take effect after the initial installation until you log out and back in again.

Action	**Command**
Get status of anomaly scan software Services	`dservice status [all] <service name>`
Restart anomaly scan software Services	`dservice restart [all] <service name>`
Stop anomaly scan software Services	`dservice stop [all] <service name>`
Start anomaly scan software Services	`dservice start [all] <service name>`

Chapter 8. Installation checklist

Follow the required steps below to complete the IBM Storage Sentinel anomaly scan software installation.

- **Prerequisites - Portal access**
  - The License Owner or other authorized user has signed the End User License Agreement (EULA).
- **Prerequisites - Host server**
  - Server is built that meets the requirements. For more information about requirements, see *Server requirements and recommendations* in IBM Documentation and ensure that the OS is updated.
  - Drives are partitioned to meet requirements.
  - Firewalls are either stopped/disabled OR configured.
  - SELinux is disabled or permissive
  - `/etc/hostname` configured with hostname.
  - `/etc/hosts` is configured with long and short hostname added to local host `127.0.0.1`.
  - Optional – postfix or sendmail is active and enabled. Successful email verified from anomaly scan software server to intended domain.
  - `atop` is installed, enabled, active and rotating properly.
  - DNS or `/etc/hosts` is configured as needed.
  - CLI access verified from your administration server (PC).
  - Run `CheckEngine.sh` to confirm that above before proceeding.
- **Prerequisites – Software**
  - Download/Install the IBM Storage Sentinel anomaly scan software provided repositories.
  - Download the IBM Storage Sentinel anomaly scan software .tar file.
- **Installation – Software**
  - Install the anomaly scan software – About 20 minutes.
  - Verify all services are running.
- **Configuration**
  - Log in to GUI - Record the manager Engine ID.
  - Receive email with the license (.txt file).
  - GUI – Upload the license.
  - GUI - **Administration** > **Home** > **Index Manager** > **Add Index** to add an index; then **Select Index** to select the index.
  - GUI - **Administration** > **Home** > **Index Manager** > **Defaults** to make these selections:

    Full Content, Map to Lowercase, CS Data Collection: Enabled, Synthetic Incremental: Enabled, Use Change Time for Incremental Indexing: Disabled
  - GUI - **Administration** > **Home** > **Index Manager** > **Merge Options** to set options.

    Purge after 30 Days, Pause Yes, Policy Purge Deleted, Days to Keep 30
  - GUI - **Administration** > **Home** > **Index Manager** > **Merge Segments**to add a merge job.

    Yes, Yes, No, Do Not Pause, 2, 720 – Schedule as required after first full scan
- **Validation**
  - Reboot the manager engine.
  - Run `CheckEngine.sh` to confirm the configuration and that all anomaly scan software processes are running.

- Confirm that you can log into the anomaly scan software GUI.
- If email is supported, confirm that email was received by the expected recipients.

Chapter 9. Navigating the IBM Storage Sentinel anomaly scan software

This chapter will guide you through the IBM Storage Sentinel anomaly scan software and show you how to navigate the software application.

Navigating the IBM Storage Sentinel anomaly scan software

After you received the license, then logged in to the IBM Storage Sentinel anomaly scan software to activate your license, it's a good idea to check the status of your system. The **Engine Status** page provides this information, and opens by default from now on, whenever you log in to the anomaly scan software.

Navigation is easy, for example to navigate to the **Engine Status** page at any time:

1. Start by selecting the **Administration** link, located along the top right portion of the screen.
2. Next, point to the **Administration** menu located on the far left. The **Administration** sub-menus open.
3. Point to the **Home** sub-menu. The **Home** sub-menus open.
4. Finally, select the **Engine Status** menu.

In the following sections, the navigation instructions have been shortened as follows: **Administration** > **Home** > **Engine Status**, which assumes that the **Administration** link on the navigation bar was selected.

The **Engine Status** page is a good place to begin exploring the anomaly scan software. The key components of this page are covered next.

The numbered markers on the **Engine Status** page indicate the five key areas of the screen that follows:

1. Navigation Links
2. System Identification
3. Administration Menu
4. Engine Status
5. Service Status

An overview of each key area appears next, with additional details following.

Navigation links

The navigation links provide quick access to the following pages in the GUI:

Link	Action
Message Center	Opens the **Message Center** page that displays alerts and messages, providing a centralized location for information about IBM Storage Sentinel anomaly scan software services and jobs.
Setup	Opens the **Setup** page that provides system configuration and network settings as well as a tab for uploading anomaly scan software updates.
Manage	Opens the **Manage** page that lets you add, edit, and delete user accounts and roles, check engine and service status, manage Productions, Collections, archive jobs, and action queries.
Index	Opens the **Index** page where you can configure and launch file server indexing and/or catalog ingestion jobs. You can also view all related reports, such as backup hosts, segment and tape/extraction reports.

Link	Action
Search	Opens the **Search** page, providing access to the powerful features of the anomaly scan software.
Administration	Opens the (default) Engine Status page and access to the **Administration** drop-down menus and sub-menus such as **Tape Manager** and **Index Manager**, among many others.
Help	Provides access to anomaly scan software documentation and a link to the IBM Storage Sentinel Support.
Sign Out	Logs you out of the anomaly scan software.

For detailed information about each of these pages and options, please see Chapter 13, "IBM Storage Sentinel anomaly scan software site map," on page 43.

System Identification

Your login and engine names always appear at the top right of the screen, along with the name of the currently selected index. This information is very important if you have:

- more than one index
- a platform consisting of several IBM Storage Sentinel anomaly scan software systems
- different users accessing the system.

Administration menu

The Administration menu becomes accessible after you select the **Administration** located along the top right of the screen or in the horizontal menu bar. As discussed in the previous example, use your mouse to hover over this menu option and then point to an option in the sub-menu that appears. This in turn sometimes opens another sub-menu of additional choices. The Administration menu options are shown and explained next.

The Administration menu provides access to the following sub-menus:

- **Home**
  Provides access to the Engine Status page and the following pages and their sub-menus: Extraction Service, Index Manager, Index Segments, Indexing Service, File Server Indexing, Action Query, and Tape Manager.
- **System**
  Provides access to various pages so you can accept the latest EULA, adjust settings, upload licenses, access log files, set up users and email alerts, select a collection, and recover a system backup.
- **Network**
  Provides direct access to the Location Manager page where you can access location data, path replacements, and define skip directories. You can also access NIS settings from this menu, as well as the User Identity pages for downloading and uploading user identity files. In addition, this menu provides an alternate way to access the **Setup** > **Network** > **Configuration** and **Software pages** for configuring Active Directory, Security settings, iSCSI, Login settings, Email alerts; and for installing software updates.
- **Reports**
  Provides an alternate way to access the **Index** > **Reports** > **Backup Hosts**, **Backupsets**, **Segments**, and **Tape and Extraction** pages.

For detailed information about each menu, please see Chapter 13, "IBM Storage Sentinel anomaly scan software site map," on page 43.

Engine status

The **Engine Status** panel provides a high level view of the current state of your IBM Storage Sentinel anomaly scan software system. Several small buttons associated with the Engine Status (and Service Status) listings also appear on this page. In the preceding screen image, for example, the square icon located to the right of the "Index Segments" status is a button. Clicking that button opens the **Index** > **Reports** > **Segments page**, providing detailed information about the segments that make up your index. Similarly, you can access additional information regarding other listings by clicking their associated square icon buttons.

The two most important pieces of information on the **Engine Status** panel are **Disk Status** and **Disk Space**:

- **Disk Status**
  Informs you of the current state of the disk array on a physical system. This information does not pertain to VMs.
- **Disk Space**
  Reports on the currently available space in the /opt/ie file system. The anomaly scan software will not properly function without adequate free space in this file system. Please open a new case in the IBM Storage Sentinel Support as soon as this starts to become an issue.

Service Status

The **Service Status** panel, located on the lower portion of the **Home** page, provides the current status of each service and indicates if any jobs are currently running. If any services are not running, or if any issues arise with the system license, appropriate messages appear on or below this section of the screen.

Chapter 10. Post Attack workflow

This section explains what to do if you suspect a ransomware attack and also the steps you should take to qualify the alerts and identify known good backups.

Post Attack workflow

Once ransomware is detected, it is critical to take immediate action to resolve any potential issues. You need to immediately contact your internal Security team and IBM Storage Sentinel Support to investigate the extent of the potential threat and fix it if necessary.

In some cases, you may need to recover your server's operating system and application installation before IBM Storage Sentinel anomaly scan software can recover the last known good application data. It's a good practice to have cold/offline servers standby that can be used as backups in an emergency.

If you have any questions about the detection status, contact IBM Storage Sentinel Support. IBM may assist with questions about the solution, troubleshoot recovery needs, and may refer you to IBM Services for more extensive threat mitigation assistance.

IBM Storage Sentinel anomaly scan software analyze dashboard

The analyze dashboard displays information related to alerts resulting from a analysis job. The dashboard is partitioned into several sections; each section displays information about the possible alerts, infected files, the state of the infection, and where the suspect files are located.

The alerts are reported in real-time to the dashboard and can be used to find which files were possibly infected by a ransomware attack. With this information, you can then perform company-instituted policies to clean up the suspect files and clear the infection from your system. The possible suspect files are identified during the analysis phase in the post-processing of an indexing job run on a particular policy. In the analyze dashboard, you can clear the alert, which stops the reporting of the infection alerts.

Logging into the analyze dashboard

Procedure

To log into the analyze dashboard:

1. In a web browser, type:

   `https://<hostname>/sentinel`

   Note: The URL is case sensitive and should be entered as shown.

   Where <hostname> is the hostname or IP address of the server running the IBM Storage Sentinel anomaly scan software.
2. In **Username**, enter your login, or the default administrator login of "admin". In **Password**, enter your password, or the default password of "admin". Select **Log In**.

Results

The analyze dashboard is displayed.

The dashboard

The analyze dashboard is made up of three main sections, as shown in the below table, and displays a summary of the state of a system as alerts occur and then breaks down the information into more detail as you scan down the page.

Below table contains the section and the information that they provide is described below:

Section	Display information on	Description
A	Alerts summary and table	Indicates new alerts and also lists all new and cleared alerts in table form. Any alerts that you clear are moved from the **New Alerts** tab to the **Cleared Alerts** tab.
B	Selected alert details	Displays detailed information on an alert that was selected in the **Alerts** table. You can also clear the alert to stop the software from reporting it.
C	Alerts graphic and files list	Displays an overall summary line of the details in the alert and graphical representations of what hosts were infected, the type of file extensions of the affected files, and the modification time of the files. The table in this section lists the suspect files, which can be grouped by column headings. Change the view of this section with the buttons to display only the list or graphs or use the **Search** field to find a specific set of files in the list.

A - Alerts summary and table

The alerts summary section and table listing the alerts is located at the top of the dashboard page.

The list of new alerts displays any recently occurring alerts and groups them into three severity levels:

- Critical - indicates a possible ransomware infection found during a policy job.
- Error - indicates that a anomaly scan software policy job was started but ended abruptly. For future use.
- Warning - indicates issues during a anomaly scan software policy job, but the job completely. For future use.

The table of alerts displays more detail about each alert and groups them by the status in two tabs of either **New Alerts**, which indicates they have not been dealt with yet, or **Cleared Alerts**, which indicates that the alert for the suspect files affected by the possible ransomware infection has been manually cleared and will no longer be reported.

B - Selected alert details

This second section of the dashboard displays the details of a selected alert in the top alert table.

The details include:

- the file location of the infection
- what type of encryption or infection method was used by the ransomware
- the policy name
- the MAC address of the anomaly scan software server.

You can scroll to another alert by selecting the arrow on the left or right of the text box. After verifying the alert per company protocol, select **Clear** to change the status of the alert from **Pending** to **Resolved OK** and the alert will appear under the **Cleared Alerts** tab in the upper section.

C - Alerts graphics and files list

The last section on the dashboard lists the statistics of the selected alert in various configurable ways.

This section can be displayed in three different ways. Select the desired icon to display:

Icon	Displays...
	Only the graphical representation of the suspect hosts, file extensions, and modification times of the suspect files.
	Only the list of suspect files.
	Both the graphs and the file list.

The top portion lists the number of suspect files, the number of suspect hosts, and the policy name that was run. These statistics change depending on which alert you selected in the top alert table.

The three graphs that are displayed show the distribution of:

- the suspect files per each host
- the suspect file extension per type
- the suspect file's modified time per date.

The list of suspect files is displayed at the bottom of the dashboard. The list can be grouped by the fields listed in the column headers. You can also choose which fields and columns to display by selecting the **Settings** icon.

Analysis and alerts

The analyze dashboard displays critical alerts that indicate a possible ransomware infection has corrupted files. A typical workflow to remediate the infected files are:

- view the details of the critical alert indicating the ransomware attack
- restore those suspect files from a clean backup copy
- clear the alert in the analyze dashboard.

Identifying the suspect files

New critical alerts may appear after a policy job runs, and, during the post-processing phase, analyzed by anomaly scan software. When a new alert is listed in the **New Alerts** table, it displays the number of infected hosts, the policy name used for the anomaly scan software analysis job, the status, the type of alert, and the severity.

Note: At this time, only **Critical** alerts with a **Pending** status are displayed in the **New Alerts** table.

To see more detail on the critical alert, select it in the **New Alerts** table. The details are displayed just below the table with the **Clear** button, which is used to stop the alert reporting once the files have been restored to a clean condition.

From the alert details, you learn the policy name, which is the identifier of the anomaly scan software analysis job, along with the type of encryption or infection method used by the ransomware, and the MAC address of the anomaly scan software server. The next step is remediation of the suspect files.

Remediate the suspect files

Once the suspect files are located, follow your company's policy to remediate the files, which could include instructions on:

- remove the corrupted files
- restore the files from a recent backup
- verify that the files are no longer infected.

Stop Reporting an infection on the IBM Storage Sentinel anomaly scan software

About this task

This procedure describes how to stop reporting if the infection on database server is cleaned up or removed and the IBM Storage Sentinel anomaly scan software continues to report an infection for each future scan.

Procedure

Once the files have been restored, the alert for that infection will continue to report previously cleared infection unless you manually stop the reporting. To stop reporting previously cleared infections do the steps that follow:

1. Open a web browser. In the browser's URL field, enter the IP address of the host where the anomaly scan software system is installed. The **Sign In** page opens.

   Note: If you are unable to access this page, verify that ports 80 and 443 are open on the firewall.
2. In the **Username** field, type the administrator username. In the **Password** field, type the administrator password. Select the **Sign In** button.
3. Select the **Alerts** tab to access the message center.
4. The infection found are listed in the message table as logs with a **jobname number** (four-digit number), this is where the **Stop Reporting** is performed. Select **Stop Reporting** to stop reporting the infection that is reported.

   Remember: For the Federation setup the procedure for stop reporting only needs to be performed on the member.
5. The alert will stop reporting the infection and move to the **Cleared Alerts** tab in the **Alerts** table.

Chapter 11. System configuration

This chapter covers system configuration, which involves confirming tmpfs partition size appropriate for your system, setting security options, enabling Audit Trail logging, integrating Active Directory on your engine, configuring NIS settings, managing users and login settings, setting up email alerts, and managing the Message Center.

Confirm tmpfs partition size

As noted in the IBM Storage Sentinel anomaly scan software Installation and Upgrade Guide, you may need to change the size of the tmpfs partition to ensure that the size is adequate. For example, if you expect to run queries that will match large numbers of documents (more than 100 million), you will need to increase the size of the tmpfs partition. This can only be done after the anomaly scan software is successfully installed. In a future release, you will be able to change the tmpfs partition size directly from the GUI; however, until then it must be done manually on the command line.

To change the tmpfs partition size using your preferred command line tool:

1. Edit the `/etc/fstab` file by changing the **`size=`** parameter for this partition.

   The default configuration is:

   ```
   tmpfs /opt/ie/var/cache/qserv tmpfs size=64G,mode=0700,uid=0,gid=0 0 0
   ```

   To increase the tmpfs partition size to 256GB, for example, change this line as follows:

   ```
   tmpfs /opt/ie/var/cache/qserv tmpfs size=256G,mode=0700,uid=0,gid=0 0 0
   ```

2. Reboot the system. Alternatively, you can manually remount the tmpfs partition with the new size:

   ```
   mount -o remount,size=256G /opt/ie/var/cache/qserv
   ```

Configuring security settings

The security settings are found at **Setup** > **Network** > **Security Settings** in the GUI.

1. Go to **Setup** > **Network** > **Security Settings**.

   Note: SSL/TLS certificates must be in PEM format.

   The **Security Settings** window provides the following controls:

   - **Upload SSL/TLS Files**
     Selecting **Upload SSL/TLS Files** opens the dialog box below that lets you upload the private key or certificate. Browse to the file that contains the private key or certificate; then select **Submit**.
   - **SSL/TLS Connections:**
     - **Required** limits SSL/TLS connection to only HTTPS. This is the default setting.
     - **Enabled** allows HTTPS access in addition to HTTP.
   - **Audit Trail Format**
     Sends the audit trail report in the format you select.
   - **Enable Audit Trail Heartbeat**
     Sends an event every hour to the Audit Trail log to indicate that the system is operational.
   - **Enable Status Report**
     Sends status reports to anomaly scan software.
   - **Enable Crash Report**
     Sends software crash reports to anomaly scan software.
   - **Enable anomaly scan software CSV report to anomaly scan software**
     Sends anomaly scan software reports to anomaly scan software.

- **Disable Concurrent Logins**

  When this option is unchecked (the default behavior), multiple concurrent logins using a single username are permitted. For example, different users may be logged in with the same credentials at the same time on different computers, or a single user may be logged in at the same time via different tabs or browsers.

  When this option is checked and another session is connected using the same credentials, the next attempt to log in with those credentials displays the **Concurrent Users** dialog box with two options. Select:

  – **OK** to terminates the other session(s) and permits the user to log in; or

  – **Cancel** to cancel the new log-in attempt and make no changes to any currently logged-in sessions for that username.

2. Once you make your selections in **Setup** > **Network** > **Security Setting**, confirm them by selecting **Submit**. Next, the web server is restarted and the **Sign In** page appears.

Login settings

The login settings are found at **Setup** > **Configuration** > **Login Settings**. This page provides the following options:

- **Enable/Disable Timeout**

  The IBM Storage Sentinel anomaly scan software lets you set a system inactivity timeout for the number of seconds you specify. If you enable the timer, you are forced to re-log in once the timeout setting is reached. Any indexing, action queries, extractions, etc., continue to run while the GUI is timed out. A query run from the **Search** page also continues to run during the timeout.

- **Enable/Disable Suspension**

  To turn on this setting, select **Enable** and specify the number of login attempts a user is permitted before being blocked from logging in.

- **Enable/Disable Splash Screen**

  If you wish to post a customized message on the **Sign In** page of the anomaly scan software, select **Enable** and type the message in the text field provided.

Email alerts

The email alert settings are found at **Setup** > **Configuration** > **Email Alerts**.

They are also accessible from **Administration** > **System** > **Email Alerts**.

IBM Storage Sentinel anomaly scan software can be integrated into your email infrastructure. Once configured, users will be able to email file restoration requests to your backup administrator.

1. Select **Add**. The **Add Email Alerts** panel opens..
2. In **Add Email Alerts**, enter the email address where you want the alerts to be sent.
3. Choose the subscription options to receive an email when a condition for a service has been reached. Select **Submit**.

A summary of your selections appears. Selecting the email address you entered enables the **Edit** and **Delete** buttons if you wish to make any changes.

Note: Systems licensed to use the anomaly scan software feature have a separate email addresses option to which anomaly scan software alerts are sent.

User account administration

Go to **Manage** > **Logins** to manage user accounts.

This page is also accessible from **Administration** > **System** > **Users**.

User profile information tells IBM Storage Sentinel anomaly scan software if your user account is active, which modules you can access, the type of access you have to each module, and your role in the system and/or the Collection if a Collection is created for the index. Each user profile includes a user account name and password.

User accounts

Select **Logins** tab to access the **Accounts** page.

Adding a user account

Procedure

To add a user account:

1. Select **Add**. The **Add Account** dialog box appears.
2. Enter the **Username**, **Password**, and **System Role** for the user you are adding. The system role determines the access privileges to the entire system for this user.

   Note: If you want to add an account for a user to whom you will grant access only to a collection (via **Add User** on the **Manage** > **Collections** page), assigning a system role is not required here; you can leave "None" as the system role.
3. If you select the check box for **AD Authentication**, you will need to provide the domain and AD username in addition to other fields in this dialog box.

   For further details on the available system roles, please see “User roles” on page 36.
4. Select **Submit** to create the user account. You can also edit and delete user accounts as needed.

Editing a user account

Procedure

To edit a user account:

1. Select **Logins** > **Accounts**.
2. In the display area, click the user account you want to edit and then select **Edit**. The **Edit Account** dialog opens.

   Note: There is one pre-defined user account: **admin**. For this account, you can edit only the password.
3. Enter a password for the selected account in the **Password** field and then confirm the password by entering it again in the **Confirm Password** field.

   Note: This strong password rule set enhancement is starting with the IBM Storage Sentinel anomaly scan software 7.11.0 build 1.20.

   Remember: All passwords should meet the following requirements:

   - must be at least 15 characters in length
   - must contain at least one uppercase letter
   - must contain at least one lowercase letter
   - must contain at least one number
   - must contain at lease one special character
   - must not contain three or more characters from the old password.
4. From **Active**, select **Yes** to make the account active or **No** to make the account inactive.
5. From **System Role**, select a role for this account. The system role determines the access privileges for this user's account.
6. Select **Submit** to submit the changes to the account.

Deleting a user account

Procedure

To delete a user account:

1. Select **Logins** > **Accounts**.
2. In the display area, select the user account you want to delete and then select **Delete**. A confirmation dialog box appears to confirm the deletion of the selected user account.

   Note: You cannot delete the admin pre-defined user account.
3. Verify that this is the user account you want to delete and select **OK**.

User roles

Use user roles to define a set of access privileges to the IBM Storage Sentinel anomaly scan software system and assign a role to each user. If you create a collection for this index, you can also assign a role to a user when you add a user to the collection.

Pre-existing system roles are **Admin** and **Search**; you can add new roles, assigning privileges as desired. Select the **Roles** tab to view the privileges associated with each.

When you create a new role, the available privileges that you can assign to the user role appear on the form shown below.

Note: The Search privileges in this group, "Results/Reports/Action Queries Actions", remain disabled unless you select either the **Search** or **Query privilege** check box. A user must be granted Search or Query privileges to enable those additional privileges.

The form above has been condensed and the privilege options for user roles are expanded in the following tables. If a privilege has a dependency on another privilege, it will be noted in the table.

Privilege	Access	Provides the Ability to...	Dependencies
Indexing. Extract, Archive, & Segments	Configure	Control the indexing, extraction, archiving, and segment configuration of an engine Change device settings within Tape Manager, but may not alter other configurations	None
	Run Jobs	Review and modify Indexing Service Defaults such as ingestion mode and the file filter enable/disable statusRun indexing, extraction, and archiving jobs	
Reports	All	Access the **Index** > **Reports** page	None
Status/Logs	All	View system logs	None

Privilege	Access	Provides the Ability to...	Dependencies
System Management	All	Modify the network and system configuration of the engine	None
Production Management	All	Create and remove productions	None
User Management	All	Create/modify/delete/ activate/deactivate users and user roles	None
Manage Entity Model	All	View the **Manage** > **Entities** page and build, download, etc., entity recognition models.	None
Collection	Add/Delete	Create/delete Collections	None
	Edit	Modify Collections, add/ edit collection users, and remove users from Collections	
Search	Edit Preferences	• Change index/project preferences within the `Search` application • Use the `Search` application	None

Privilege	Access	Provides the Ability to...	Dependencies
Query	Results Reports Edit	• Use the Search application • Query all documents • Query and reconstruct the content of documents in search results • See query results • See backup information • See query reports, and create CSV downloads • Modify the filters, and filter on Summary reports if applicable; may upload filters • See locations in the Search application • Modify the locations checkbox tree in locations panel of the search application • Make the query box editable and fill in a free-form query; can also upload a query	None
Review Panel	All	• Use the Search application • Query and reconstruct content of documents in the search results • See query result reconstruction and filter on details	None

Privilege	Access	Provides the Ability to...	Dependencies
Results/Reports/Action Queries/Actions Dependent on Search being enabled	Tag	Add or remove tags; can run action queries that add or remove tags	Search
	Delete	Activate the delete/ undelete button, may run action queries that affect delete/undelete	
	Delete from LAN	Delete files from file servers. **Note: Delete from Source** is not supported for SharePoint or OneDrive.	
	Archive	Archive and run action queries that mark results for archiving	
	Extract	Create extract jobs and see job status in the Search application	
	Download	Download CSV of query results	
	Rename Owner	Change the owner of objects	
Saved Queries	Load	View/load saved queries	Query
	Edit	View/add/save/delete saved queries	
Action Queries	Run	View/run action queries	Query
	Edit	View/create/modify/ delete action queries	

Adding a user role

Procedure

To add a user role:

1. Go to **Manage** > **Logins** > **Roles**.
2. Select **Add** to add a user role.
3. Give the new role a name and assign the desired privileges. Select **Submit** to create the role.

Results

The newly defined role now appears as a valid choice when adding new user accounts. You can also edit and delete user roles as needed.

Editing a user role

Procedure

To edit a user role:

1. Go to **Manage** > **Logins** > **Roles**.
2. Select **Logins** and then **Roles**. Select the role name that you wish to edit.

   Note: You cannot edit the pre-defined **admin** user role.
3. Select **Edit** to open the **Edit Role** dialog.
4. Select the plus sign (+) next to a privilege to expand the listings under a header; select the minus sign (-) to collapse the list.
5. Select to place or remove a check mark associated with each privilege you wish to grant or deny. Select the top level of an expandable list to select all of the subordinates in that group.
6. If a set of privileges is grayed out, it may have dependencies. For example, **Results** > **Reports** > **Action Queries** Actions is only available if you have access to **Queries**.
7. Commit the changes by selecting **Submit**.

Deleting a user role

Procedure

To delete a user role:

1. Go to **Manage** > **Logins** > **Roles**.
2. Select the role name that you wish to delete.

   Note: If the user role is assigned to a user account, it cannot be deleted.
3. Select **Delete**. When prompted, select **Yes** to confirm the deletion.

Chapter 12. Additional Administrative Functions

This section of the guide covers additional administrative tasks that you may need to perform. These administrative tasks include managing collections, how to use the IBM Storage Sentinel anomaly scan software for disaster recovery, and shutting down the engine.

Application backup and recovery

The backup and recovery procedures provide disaster recovery options so you can restore an engine back to its state at the time of the backup. The procedures let you back up, restore, and recover your anomaly scan software configuration files, indexes, databases, log files, and license files. This is largely achieved via commonly used third party commercial backup and restore software, with the final recovery phase accomplished through the IBM Storage Sentinel anomaly scan software graphical user interface.

Once you restore a backup image to /opt/ie/backup using a backup tool, open the freshly installed anomaly scan software and use the graphical user interface to recover the anomaly scan software system configuration, database, and index data.

Note: Requirements for a Successful Recovery:

- Stop or cancel any indexing, querying, extraction, or archive jobs before attempting a recovery. Otherwise, the recovery procedure will stop all anomaly scan software services and any running jobs, which may have a negative impact.
- The hardware must have the same MAC address as the Backup.
- The installed anomaly scan software version must be the same or newer than the Backup version.
- The restore from the backup server must have successfully completed, with the backed-up files now residing on the client system with the anomaly scan software.
- Anomaly scan software systems depend on the ability to resolve host names to IP addresses. The recommended and normal way to accomplish this is to use the Domain Name System (DNS). If, however, you are not using DNS, then the `/etc/hosts` file on each appliance must have the hostnames/IP addresses of the manager.

System shutdown/reboot

If maintenance is needed on your IBM Storage Sentinel anomaly scan software system, you will first need to perform a shutdown. The **Administration** > **System** > **Shutdown** page lets you initiate a **System Reboot** or a **System Halt**. In each case, access becomes unavailable to the browser-based administrative and search interfaces, although you will regain access after the system restarts if you chose a System Reboot.

Note: All indexing is stopped while your anomaly scan software system reboots or halts. If an indexing job was running, it terminates and is marked incomplete. Whenever possible, it is best to reboot your engine when indexing will not be disrupted.

To shut down the engine:

1. Go to **Administration** > **System** > **Shutdown**.
2. Select a shutdown action from the **Action** list:

Option	Description
System Reboot	Reboots the system. Normal operation resumes when the reboot completes. Your anomaly scan software system does not normally need to be rebooted; instead, you may be asked to perform this action at the direction of IBM Storage Sentinel Support personnel.
System Halt	Halts the system in preparation for power-off. If you need to perform hardware maintenance, use this option before turning off the power.

Note: You may not see any feedback while the reboot progresses. If you resubmit a Reboot Now request, the system informs you that you do not have access to the requested page and directs you to Sign In again-you will be able to sign in once the reboot completes. Otherwise, if you click elsewhere on the page while a reboot is in progress, an "Unable to Connect" error appears. In this case as well, you will be able to sign in once the reboot completes.

3. Next, from the **When** drop-down list, select when you want the action to be performed:

Option	Description
Now	Perform the selected action immediately, without delay. You should use this option only if you are certain that no one is accessing the engine at the console or via an SSH session.
In 1 minute	Perform the selected action with a one-minute delay. This option allows any console or SSH user sessions a chance to react to the shutdown request, possibly by cancelling it at the command line. This is the recommended delay. If the shutdown request is canceled by a console or SSH user, you will receive a shutdown cancelled message on the Administration page from which you submitted the request. In the absence of such a cancellation message, you can assume that shutdown is proceeding according to the selection options.
In 5 minutes	Similar to the 1-minute delay, but provides a little more time for users to react.

4. Select **Submit**.

Chapter 13. IBM Storage Sentinel anomaly scan software site map

The information in this chapter provides details about the IBM Storage Sentinel anomaly scan software GUI, which includes details about the various pages, menus, sub-menus, and the functionality of these pages.

About

The About link displays the version and build numbers of your IBM Storage Sentinel anomaly scan software, as well as copyright information and trademark acknowledgments.

Help

The **Help** menu links to the following:

- The IBM Storage Sentinel anomaly scan software Installation and User Guide, i.e., the main User Guide, the Archive Guide, Search Guide, Catalog Ingestion Guide, and Query Operators Guide.
- The IBM Storage Sentinel Support.

Sign Out

Sign Out logs you out of the IBM Storage Sentinel anomaly scan software.

Chapter 14. Troubleshooting

This section covers log files and warning and error messages. When you need IBM Storage Sentinel Support to investigate any issues that arise, you may be asked to send log files. This section provides information about log files and specific warning and error messages that may appear as you use the application.

System and application log files

To view the system and application log files, go to **Administration** > **System** > **Logs**.

On this page, view and download various log files such as Tape Manager logs, query logs, and extraction logs. The logs are particularly useful when you need IBM Storage Sentinel Support to investigate any issues that arise. An excerpt from the long list of available log files follows.

Note the following:

- Each file has an associated **Download** button that provides convenient access to the related information.
- To view and/or save a log file, select **Download** next to the desired file version. When several versions of the same log file are available, note that each version covers a different date range.
- The **Modified** time stamp tells you when the file was last updated. Use this information to determine which version is important to you.

Some of the log files available for download are listed in the following table.

SYSTEM COMPONENT LOG	LOGS THE FOLLOWING INFORMATION
Application Server (catalina.out)	Graphical user interface operations
Catalog Ingestion Messages (ceng.log)	All catalog ingestion jobs
Catalog Ingestion Server (ceng_svr.log)	Catalog ingestion server process
Catalog Ingestion Feed (tdbfeed.log)	All ingested tapes
Extraction Service (extract.log)	All extraction jobs
Index Management (indexman.log)	Index management-related activities
Indexing Service (ie_run.log)	All catalog and indexing jobs
Location Manager (locman.log)	All location manager activities during system startup
Merge History (arcmerge.log)	Information regarding merged segments
Post Processor (postproc.log)	All post-processing activities
Query Server (qserv.log)	All query activities
Scheduler (sched.log)	Logs information pertaining to the scheduler process
Share Discovery (getshares.log)	Information pertaining to discovered domains. Also records all discovery activity (users, groups, computers) when an IBM Storage Sentinel anomaly scan software system is attached to Active Directory.
System Configuration (setconfig.log)	System configuration information

SYSTEM COMPONENT LOG	LOGS THE FOLLOWING INFORMATION
System Messages (messages)	Global system messages, i.e., the messages plus mail, cron, daemon, kern, auth, etc.
Tape Manager (tapeman.log)	Tape operations during catalog, index and extraction jobs
Upgrade Logs (rpmutil.log)	Information regarding your anomaly scan software upgrades
Web Server (error_log)	Errors related to the Apache web server

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

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

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linear Tape-Open, LTO, and Ultrium are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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

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

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

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

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

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

Rights

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

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

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

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

Privacy policy considerations

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

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

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

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

Glossary

A glossary is available with terms and definitions for the IBM Storage Sentinel family of products.

See the IBM Storage Sentinel glossary.

Product Number: 5900-APZ