

IBM StoredIQ

Deployment and Configuration Guide



Note

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

This edition applies to Version 7.6.0.17 of product number 5724M86 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation 2001, 2018.**

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

Contents

About this publication.....	V
IBM StoredIQ product library.....	V
Contacting IBM StoredIQ customer support.....	V
IBM StoredIQ components.....	1
Solution components.....	1
IBM StoredIQ products.....	1
Planning for deployment.....	3
Open Virtual Appliance (OVA) configuration requirements.....	3
Network and port requirements.....	5
Environment sizing guidelines.....	5
Stack-provisioning prerequisites.....	7
Security.....	7
Deployment of IBM StoredIQ.....	9
Open Virtualization Format (OVF).....	9
Deploying the virtual appliances.....	9
Installation of IBM StoredIQ on Microsoft Hyper-V.....	11
Prerequisites.....	11
Extracting the OVA.....	11
Converting VMWare .vmdk to Microsoft Hyper-V .vhdx.....	11
Building the Virtual Machine.....	11
Adding more disks.....	12
Installation of IBM StoredIQ.....	13
Deploying an Elasticsearch cluster.....	13
Installing the gateway.....	18
Installing the data server.....	20
Installing and configuring the application stack.....	21
Logging in to the application stack user interfaces.....	29
Configuration of IBM StoredIQ.....	30
Default open ports.....	30
Key and certificate management.....	32
Enabling encryption of IBM StoredIQ AppStack application data.....	37
Enabling or disabling FIPS.....	39
Securing Elasticsearch cluster communication.....	40
Configuring the gateway.....	41
Managing the status of secure gateway communication.....	41
Configuring the IBM StoredIQ data server.....	42
Configuring the application stack to receive reports and notifications.....	43
Configuring the application stack to synchronize data with the governance catalog.....	43
Installing IBM StoredIQ Desktop Data Collector.....	46
Backing up the IBM StoredIQ image.....	47
Upgrades for IBM StoredIQ.....	48
Considerations when not upgrading from the preceding version.....	48

Upgrading the Elasticsearch cluster.....	51
Updating the gateway and data servers.....	51
Updating the application stack.....	52
Installing the Livelink APIs.....	54
Configuring authenticated users for SMTP notifications.....	55
Creating an SMTP Receive Connector.....	55
Authenticating users.....	55
Configuring email notification with an authenticated user.....	56
Notices.....	57
Privacy policy considerations	58
Terms and conditions for product documentation.....	59
Trademarks.....	59
Index.....	61

About this publication

IBM StoredIQ Deployment and Configuration Guide provides information about how to plan, deploy, and configure the IBM StoredIQ product.

IBM StoredIQ product library

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

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

Contacting IBM StoredIQ customer support

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

- 1-866-227-2068

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

IBM Knowledge Center

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

Contacting IBM

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

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

IBM StoredIQ components

There are three components of the IBM StoredIQ solution: the application stack, the gateway, and the data server. These components work together as the IBM StoredIQ products.

Solution components

IBM StoredIQ provides three solution components: the gateway, data servers, and application stack (AppStack).

Gateway

The gateway communicates between the data servers and the application stack. The application stack polls the gateway for information about the data on the data servers. The data servers push the information to the gateway.

Data servers

IBM StoredIQ Data Server helps understand the data landscape of the enterprise. It obtains the data from supported data sources and indexes it. By indexing this data, you gain information about unstructured data such as file size, file data types, file owners.

Data servers can be categorized in two types: DataServer - Classic and DataServer - Distributed. A data server of the type DataServer - Classic uses the embedded PostgreSQL database for storing the index. With a data server of the type DataServer - Distributed, the index is stored in a distributed Elasticsearch cluster. Data servers of this type also provide better performance in search queries and can manage much larger amounts of data than those of the type DataServer - Classic.

The data servers push the information about volumes and indexes to the gateway so it can be communicated to the application stack. Multiple data servers feed into a single gateway.

In addition to completing standard administrative tasks, administrators can deploy the IBM StoredIQ Desktop Data Collector and index desktops from the data server.

Application stack

The application stack provides the user interface for the IBM StoredIQ Administrator, IBM StoredIQ Data Workbench, IBM StoredIQ Insights, and the IBM StoredIQ Policy Manager products.

The synchronization feature for integration with a governance catalog is also part of the application stack.

IBM StoredIQ products

The IBM StoredIQ products include IBM StoredIQ Administrator, IBM StoredIQ Data Workbench, IBM StoredIQ Policy Manager, and IBM StoredIQ Desktop Data Collector.

IBM StoredIQ Administrator

IBM StoredIQ Administrator monitors and manages the distributed infrastructure at a client site. IBM StoredIQ Administrator sits between the IBM StoredIQ interface and the applications and facilitates the transfer and communication of information. IBM StoredIQ Administrator understands and manages IBM StoredIQ concepts such as volumes, indexes, harvests, and configurations. At the same time, it manages the application concerns such as info set lifecycle and creation, data server and volume configuration, and action and target set management. To this end, it is divided into two sections, that is, platform and application so that the administrators know where to accomplish a task.

IBM StoredIQ Data Workbench

IBM StoredIQ Data Workbench enables visualizing the indexed data. You can identify potential red-flag issues to know how much and what types of data you have on different types of servers. It alerts people about potentially interesting or useful data.

If IBM StoredIQ is configured for integration with Information Governance Catalog, you can manage here which objects are published to the governance catalog. In addition, you can apply the classification rules that are synchronized from the governance catalog.

IBM StoredIQ Policy Manager

IBM StoredIQ Policy Manager acts on data in an automatic fashion at scale.

IBM StoredIQ Desktop Data Collector

The IBM StoredIQ Desktop Data Collector deploys from the IBM StoredIQ Data Server Administrator interface. It indexes desktops as volumes. The volumes appear in the Data Server Administrator interface, the IBM StoredIQ Data Workbench, where the data can be analyzed and acted upon, and the IBM StoredIQ Administrator application.

IBM StoredIQ uses a non-invasive deployment, specifically:

- Lightweight client deployment with no browser plug-ins and no client-side UI installations. All major browsers, such as Internet Explorer and Firefox, are supported natively.
- One (common) index that is shared across all use cases. This index supports a wide variety of data sources and is used by application dashboards for many use cases.
- No agents are placed on data sources. The native API/protocol is used whenever it is possible. Third-party or custom connectors are used only when necessary.

Note: Agents are used for desktop collection.

Planning for deployment

When you plan a deployment of IBM StoredIQ, evaluate several infrastructure considerations.

In addition to the information in this section, review the requirements detailed in the IBM Software Product Compatibility Reports (SPCR) tool at: [Software Product Compatibility Reports: StoredIQ 7.6](#)

If you plan to use IBM StoredIQ for Legal Identification and Collection to create and manage data boxes and data requests that are to be fulfilled by IBM StoredIQ, also check the system requirements for StoredIQ for Legal at: [Software Product Compatibility Reports: StoredIQ for Legal 2.0.3](#)

Generate customized reports with the SPCR tool

Go to the page at [Software Product Compatibility Reports](#) to create a high-level report for supported operating systems, related software, hypervisors, and supported translations for any product. You can also create an in-depth report to get detailed system requirements, hardware requirements, and end of service information for each product. You can search for a product in all of the report types and reports are generated based on your query values.

The following report types are the most commonly generated reports from software product compatibility reports:

Detailed system requirements

When you select your product version for the detailed system requirements report, you can set a report filter for **Operating system platforms**, **Product components**, and **Capabilities**, including prerequisites and support software. After you view the report, you can save it as a URL to generate anytime or download it as a PDF.

Hardware requirements

When you select your product version for the hardware requirements report, you can set a report filter by the **Operating system families** option. Set the operating system filter by selecting some or all of the operating systems that are supported by your product. After you view the report, you can save it as a URL to generate anytime or download it as a PDF.

End of service

The end of service report shows the service window of the products that you specify over an eight-year span. For example, you can find out when your product is scheduled to go out of service.

Open Virtual Appliance (OVA) configuration requirements

IBM StoredIQ is deployed as virtual appliances and currently only supported in VMware ESXi 5.0 (all fix pack levels) or VMware ESXi 6.0 (all fix pack levels) environments. You must have a virtual infrastructure that meets the IBM StoredIQ hardware requirements.

Application stack

- vCPU: 2
- Memory: 4 GB
- Storage:
 - Primary disk (vmdisk1): 21 GB
 - Data disk (vmdisk2): 10 GB

Gateway server

- vCPU: 4
- Memory: 8 GB
- Storage:

- Primary disk (vmdisk1): 100 GB
- Data disk (vmdisk2): 75 GB
- Swap disk (vmdisk3): 40 - 100 GB

DataServer - Classic

- vCPU: 4

Even though increasing the number of vCPUs increases performance, the actual benefits depend on whether the specific host is oversubscribed or not.

- Memory: 16 GB

While the minimum value works under light-load condition, as the load increases, the data server quickly starts using swap space. For high load situations, increasing RAM beyond 16 GB can benefit performance. Monitoring swap usage can provide insight.

- Storage:

- Primary disk (vmdisk1, SCSI 0:0): Default is 150 GB

This virtual disk has an associated VMDK that contains the IBM StoredIQ operating code. Do not change its size.



Attention: If you delete the primary disk, you delete the operating system, and the IBM StoredIQ software; the virtual machine might need to be scrapped.

- Data disk (vmdisk2, SCSI 0:1): Default is 2 TB

This virtual disk can be resized according to expectations on the amount of harvest data to be stored. For purposes of estimation, the index storage requirement for metadata is about 30 GB per TB of managed source data. Full-text indexing requires an extra 170 GB per TB. Therefore, the default data disk size is targeted for managing 10 TB of source information.

- Swap disk (vmdisk3, SCSI 0:2): Default is 40 GB

When under load, the data server can use much RAM; therefore, having ample swap space is prudent. The minimum swap size is equal to the amount of RAM configured for the virtual machine. For best performance under load, place this disk on the highest speed data store available to the host.

The general size limits for a data server are 150 million objects or 500 defined volumes, whichever limit is reached first. Assuming an average object size of 200 KB equals about 30 TB of managed storage across 30 volumes of 5 million objects each, the index storage requirement for metadata on ~30 TB of storage that contains uncompressed general office documents is ~330 GB (11 GB per TB). Add 100 GB per TB of managed storage for full-text or snippet index. For example, to support 30 TB of storage that is indexed for metadata, you need 8 TB indexed for full-text search and extracted text (snippet cache) of 8 TB for auto-classification. A total of 1.9 TB of storage is required (metadata: 330 GB, full-text: 800 GB, snippet cache: 800 GB).

Data-server performance is impacted by the IOPS available from the storage subsystem. For each data server under maximum workload, at least 650 IOPS generally delivers acceptable performance. In the situations where there is a high load on the system, the IOPS that is used can reach up to 7000 with main write operations.

DataServer - Distributed

- vCPU: 4
- Memory: 16 GB
- Storage:
 - Primary disk (vmdisk1, SCSI 0:0): Default is 150 GB
 - Data disk (vmdisk2, SCSI 0:1): Default is 2 TB
 - Swap disk (vmdisk3, SCSI 0:2): Default is 40 GB

If you deploy this type of data server, you must also deploy an Elasticsearch cluster with at least one node. If you deploy a cluster with more nodes, each of the Elasticsearch nodes must meet the listed requirements.

Each Elasticsearch node

- vCPU: 1
- Memory: 16 GB
- Storage:
 - Primary disk (vmdisk1): 100 GB
 - Data disk (vmdisk2): 1 TB

The required memory depends on the index size that is handled by the data node. For a better performance, consider increasing the memory. The total memory available to the node must be distributed between the host operating system, the Elasticsearch java heap space, and the kernel file system caches. For example, if the system has 32 GB memory, 2 GB must be allocated for the host operating system, 15 GB for the java heap space, and the remaining 15 GB for the file system caches.

The data disk (vmdisk2) can be resized according to expectations on the amount of harvest data to be stored.

Network and port requirements

Data sources and types can drastically impact the scale and scope of what needs to be deployed. The complexity of the source directly affects the number of data servers to be deployed, for example, exchange versus simple text documents in a CIFS location.

You must enable network connectivity from the following locations:

- The data server IP address to the gateway IP address on port 11103.
- The gateway IP address to and from the application stack IP address on ports 8765 and 5432.
- Ports 80, 443, and 22 from the administrative user's workstation (place from which the administrator is completing work with IBM StoredIQ Administrator) to the application stack and data server IP addresses.
- Port 22 from the administrative workstation to the gateway IP address.

TCP: port ranges for the firewall

To ensure the network access for desktop volumes, the following port ranges must be open through a firewall.

- 21000-21004
- 21100-21101
- 21110-21130
- 21200-21204

Environment sizing guidelines

To size an environment precisely, you must understand the factors such as harvest frequency, complexity of the source, and use case scenarios that drive application use and action execution.

The general design guidelines for IBM StoredIQ are as follows:

- For data servers of the type DataServer - Classic:
 - One data server for up to 30 TBs of data (which can vary based on the number of volumes, objects per volume, and object types).
 - Up to 500 volumes per data server.

- Up to 150 million objects per data server.
- One gateway per 50 data servers.
- One application server.
- NFS is slightly faster than CIFS for metadata only, but assume CIFS/NFS even for this exercise.
- Full-content processing of file (for example, .ZIP, .RAR, .GZ) and email archive (.PST, .NSF, .EMX) processing are slower as items must be extracted from the archives. If there is a significant number of these files in the file system and they are not excluded from content processing, the full-content processing rate can be too high. Until you have an initial index of the file system, you do not know how to weigh full-content processing of archives.
- An object/time metric is appropriate for metadata only NOT computing a hash, membership in the National Institute of Standards and Technology (NIST) or enumerating objects that are contained in archives. Converting it to a bytes/time metric is a function of the average object size and might vary tremendously. An average object size of 250 KB was used for the metric that is provided earlier.
- A bytes/time metric is appropriate for metadata-only computing a hash and full-content processing. The object per second rate can vary tremendously depending on the object type and sizes encountered. For example, processing an email or file archive is much more expensive than a PDF document.
- Metadata-only not computing a hash, membership in the NIST list, or enumerating objects that are contained in archives is requesting only the file-attribute information from the NAS. Individual files are not opened and read. The processing rate is high, but that does not translate into a large amount of data that traverses a network between the NAS and data server. The bytes/time rate does not translate into bytes served by the NAS and sent over the network.
- Metadata-only computing a hash, membership in the NIST list, or enumerating objects that are contained in archives opens and reads the contents of each file. The content of all requested files traverses the network between the NAS and data server. The maximum load that the data server can place on a NAS is metadata-only processing. It requires all file content to be read to compute a hash or enumerate objects that are contained in archives. The bytes/time rate translates into bytes served up by the NAS and network traffic that must be considered.
- Full-content processing opens and reads the contents of each file to extract all text. The content of all requested files traverses the network between the NAS and data server. The processing time to enumerate archives, extract text, index words, and extract entities on the data server reduces the rate that data is requested from a NAS compared to metadata-only with full hash. The bytes/time rate translates into bytes served up by the NAS and network traffic that must be considered.
- The interrogator process count on the data server for "metadata only not reading all content indexing" is set to eight for optimal performance.
- The interrogator process count for all other processing that involves reading all content default setting is four per data server.
- The interrogator count can be viewed as the number of client connections that are made to a data source actively requesting data. It is important for capacity planning for the data source.
- The data servers are assumed to be "network close" to the NAS data sources. Network latency under 10 ms with at least 1000 Mbps bandwidth is assumed (connected through local area network). The data servers need a low latency high-bandwidth connection to a NAS data source for acceptable indexing performance.
- The gateway and application stack can be located remotely from the data servers. Network connections with latency greater than 10 ms and bandwidth of at least 2+ Mbps are acceptable.

VMware vSphere requirements

- VMware vSphere v5.0 and fix packs or v6.0 and fix packs.
- VMware virtual machine hardware version 8.0 or later. For more information, see [the VMware product documentation](#).
- The appropriate VMware license to enable the required processor cores and memory for the virtual machine.

Stack-provisioning prerequisites

Before a deployment, verify that you meet these prerequisites.

- At least one physical server with sufficient processor, RAM, and hard disk configuration for the planned management project.
- VMware ESX or ESXi on CD/DVD or USB drive.
- IP addresses, cables, and physical switch ports for at least the ESXi/ESX interface, one data server, one gateway server, and one application stack.
- Network connectivity that is enabled from the following locations:
 - The data server IP address to the gateway IP address on port 11103
 - The gateway IP address to and from the application stack IP address on port 8765
 - Ports 80, 443, and 22 from the administrative workstation to the application stack and data server IP addresses
 - Port 22 from the administrative workstation to the gateway IP address.
- Network connectivity that is enabled from the data server IP address to any data sources to be harvested and managed.
- A management station computer or notebook from where the load-management work is done.

Security

Plan and implement specific security measures to protect the application and the data it manages, especially when you deploy IBM StoredIQ into sensitive environments.

IBM StoredIQ keeps your data secure through encryption, security hardening, and auditing.

Federal Information Processing Standard (FIPS)

FIPS is a standard recommended by the National Institute of Standards and Technology (NIST) and the US Federal Government. It ensures certain security standards are met for software or hardware components deployed at US government sites. Enabling FIPS ensures that the SSL/TLS engine that is compliant with the US Government recommendation is used. IBM StoredIQ supports FIPS Level 1.

Secure gateway communication can be enabled without FIPS. If FIPS is enabled, IBM StoredIQ uses FIPS compliant versions of OpenSSL.

Secure communication and encryption of data in motion

In a production environment, you should configure or install certificates on the AppStack to enable HTTPS communication and to enable encryption of data in motion between the browser and the AppStack. You can do this during installation and initial configuration or at any time afterward. For details, see [the instructions for configuring certificates](#).

The gateway handles the communication between the data servers and the application stack. By default, the communication between the gateway, any data servers, and the AppStack is in plain text and is not encrypted. If your enterprise security policy mandates encryption of data in motion, enable secure gateway communication. In this case, secure gateway communication must be configured on all three IBM StoredIQ components. You can enable secure gateway communication during installation and initial configuration or at any time afterward. For details, see [“Managing the status of secure gateway communication” on page 41](#).

IBM StoredIQ then uses stunnel to ensure secure communication between the components. If your environment includes data servers of the type DataServer - Distributed, stunnel can also be used to encrypt the communication between the nodes within the Elasticsearch cluster but not for encrypting the communication between the data server and the Elasticsearch cluster.

To secure the communication between the data server and the Elasticsearch cluster and the communication within the Elasticsearch cluster likewise, you can enable Search Guard. For more information, see [“Securing Elasticsearch cluster communication”](#) on page 40.

If FIPS is not enabled, the following cipher suites and encryption algorithm are used for data at rest:

```
TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256
```

You can configure these cipher suites in the configuration files listed in the list of key and certificate files. However, if you run the utilities for enabling stunnel, you might need to make the respective configuration changes again.

Encryption of data at rest

Starting with IBM StoredIQ version 7.6.0.15, the disk volume on which the Elasticsearch indexes are stored is encrypted by default. IBM StoredIQ uses Linux Unified Key Setup (LUKS) for disk encryption. For details about key management, see [“Key and certificate management”](#) on page 32.

Optionally, you can encrypt the application data on the IBM StoredIQ application stack. For more information, see [“Enabling encryption of IBM StoredIQ AppStack application data”](#) on page 37.

Network isolation

If full-text harvesting and Step-up Analytics actions (cartridges) are applied, Elasticsearch indexes can contain potentially sensitive content. Therefore, you should deploy the Elasticsearch nodes in an isolated location on the network (for example, as an enclave or behind a firewall) that is properly secured according to the sensitivity of the data being harvested. Only the IBM StoredIQ application stack and data servers should be allowed to communicate with the Elasticsearch nodes.

Also, any data servers and the gateway should be deployed in an isolated network location to allow for communication with authorized clients only.

Access control

root and builder accounts on the Elasticsearch cluster nodes

Remote login for `root` can be disabled. However, local root login is required, either log in as `root` or use the `su` command to obtain root permissions temporarily.

You set the passwords for the `root` and `builder` accounts during the configuration process when you start the VM for the first time. You can change these passwords anytime.

siqadmin account on the AppStack

Administration of the AppStack usually does not require direct root access. For day-to-day administration, the `siqadmin` account can be used

You set the passwords for the `siqadmin` account during the configuration process when you start the VM for the first time. You can change this password anytime.

Deployment of IBM StoredIQ

IBM StoredIQ deployment utilizes the OVF file format and vSphere Client.

Open Virtualization Format (OVF)

An OVF is a file format that allows for the exchange of virtual appliances across products and platforms.

The OVF format offers these advantages:

- OVF files are compressed, allowing for faster downloads.
- The vSphere Client validates an OVF file before it imports it and ensures that it is compatible with the intended destination server. If the virtual appliance is incompatible with the selected host, it cannot be imported, and an error message appears.
- An OVF can encapsulate multitiered applications and more than one virtual machine.

The basic differences between OVFs and OVAs are as follows:

- An OVF is a collection of items in a single folder. Most commonly, it is a description file (.ovf), a manifest file (.mf), and virtual machine state files (*.vhd or *.vmdk).
- An OVA is a single file. The OVA is the OVF folder contents that are all compressed into a single file. The OVA is used when you want to take an OVF and share it, or give it as a download. The OVA needs to be opened into the OVF before it can be used, although it is not necessary when you use VMware. This step is automatically conducted during import. Both OVF and OVA files can be directly imported into vCenter and vSphere.

Deploying the virtual appliances

Use vSphere Client to deploy the virtual appliances to an ESX server. Three OVAs need to be deployed, preferably in this order: the gateway, the data server, and the application stack. However, if an Elasticsearch cluster is also being deployed, those OVAs must be deployed first.

Ensure that all prerequisites for the deployment are met before you start this task.

The number of data server OVAs deployed depends upon the number of data servers needed. Each data server requires a separate OVA deployment.

1. Connect to ESX server or vCenter server.
2. From the **File** menu, select **Deploy OVF Template**.
The **Deploy OVF Template** wizard appears.
3. Within the **Deploy OVF Template** wizard, complete these steps.
 - a) Within the **Select source** page, click **Local file**, and then browse to and select the appropriate OVF. Click **Next**.
 - b) Within the **Review details** page, review the OVF template details. These storage requirements are critical and will be used to select a data store during deployment. Click **Next**.
 - c) Within the **Select name and folder** page, enter a name for the deployed template or use the default name. Click **Next**.
 - d) Within the **Select a resource** page, select the resource pool where the deployed OVF template runs. Click **Next**.
 - e) Within the **Select storage** page, select a data store on which to store the deployed OVF template files. Click **Next**.

- f) Within the **Disk Format** list, select the disk format to be deployed. Note that although **Thin Provision** saves disk space, it can negatively affect performance. If possible, select **Thick Provision Lazy Zeroed**. Click **Next**.
 - g) Within the **Network Mapping**, map the network to the appropriate network. Click **Next**.
 - h) Within the **Ready to Complete** page, review the deployment settings. Click **Finish** to complete the Deploy OVF Template.
At this point, you have the option of selecting the Power on check box to turn on the VM once deployment is complete.
4. Repeat this procedure for each OVA.

Installation of IBM StoredIQ on Microsoft Hyper-V

IBM StoredIQ can be installed on Microsoft Hyper-V.

Note: Currently, Elasticsearch is not supported.

Prerequisites

To install IBM StoredIQ on Microsoft Hyper-V, you must complete the following installations.

1. Microsoft Hyper-V Manager must be installed on Windows system.
2. 7-zip for Windows or tar for Linux operating system must be downloaded to extract contents of the OVA.
3. 7-zip for Windows or gunzip for Linux operating system must be installed to decompress the `vmrk . gz` file.
4. `qemu-img` (Linux) or Microsoft Virtual Machine Converter 3.0 (Windows) must be installed to convert the files from VMWare . `vmrk` files to Hyper-V . `vhrx` files.

Extracting the OVA

Microsoft Hyper-V does not support OVAs and OVF. The virtual machine needs to be built manually.

Follow these steps to extract the OVA.

1. Download the OVAs for the gateway, the data server, and AppStack.
2. Extract the contents of the OVAs by using 7-zip for Windows or tar for Linux operating system.
3. Decompress the `vmrk . gz` by using 7-zip for Windows or gunzip for Linux.
4. Delete everything except the `vmrk` file when the OVA extraction is complete.

Converting VMWare .vmrk to Microsoft Hyper-V .vhrx

For instructions about converting the .`vmrk` file to .`vhrx` file, see the following links.

[Use Microsoft Virtual Machine Converter](#)

[Using qemu-img](#)

Building the Virtual Machine

Follow these steps to build the virtual machine.

1. Select **New Virtual Machine** > **Next** from Hyper-V Manager.
2. Enter the name of the virtual machine in the **Specify Name and Location** window and click **Next**.
3. Select **Generation of the machine** in the **Specify Generation** window.
4. Enter the startup memory in the **Assign Memory** window.
5. Select the correct network in the **Configure Network** window and click **Next**.
6. Select **Use an existing virtual hard disk** in the **Connect Virtual Hard Disk** window. Search the disk location and select one of the primary disks for the gateway, data server, and AppStack and then click **Next**.
7. Click **Finish**.
8. Repeat these previous steps for all IBM StoredIQ disks.

Adding more disks

Follow these steps to add more disks.

1. Select a virtual machine in Microsoft Hyper-V Manager.
2. Select the settings for the virtual machine under **Actions**.
3. Select the IDE controller and highlight Hard Drive in the **Settings** window.
4. Select Hard Drive and virtual hard disk and click **New**.
5. Click **Next** in the New Virtual Hard Disk wizard.
6. Select **VHDX** in Choose Disk Format and then click **Next**.
7. Select **Dynamically expanding** in Choose Disk Type and then click **Next**.
8. Enter a name for the disk in the **Specify Name and Location** window and then click **Next**.
9. Select **Create a new blank virtual hard disk** and enter a size in the **Configure Disk** window and then click **Finish**.
10. Repeat this procedure for all additional disks.

Installation of IBM StoredIQ

After the OVF deployment, proceed to install IBM StoredIQ.

Note: You can install IBM StoredIQ Desktop Data Collector only after you finish configuring the data servers. For information about its installation, see [“Installing IBM StoredIQ Desktop Data Collector” on page 46](#)

Deploying an Elasticsearch cluster

If you are deploying a DataServer - Distributed type of the data server, you must first deploy the Elasticsearch cluster. If you are deploying a DataServer - Classic type of the data server, proceed to the topic [“Installing the gateway” on page 18](#).

The Elasticsearch cluster setup must be complete before you run the data server first-boot process.

The default setup is a three-node Elasticsearch cluster. However, in smaller enterprises or in non-production environments a single node Elasticsearch cluster can be set up.

Deploying a three node Elasticsearch cluster

For a default setup, deploy a three node Elasticsearch cluster.

Before you run the data server first-boot process, complete the setup of the Elasticsearch cluster.

1. Download the Elasticsearch OVA from [IBM Fix Central](#) and deploy the most recent Elasticsearch OVA to vcenter.
2. Edit the VM settings that were created from the OVA.
3. Optional: Reduce the memory.
4. Optional: Reduce the data disk.
 - a) Edit the VM.
 - b) Edit the second disk by expanding the configuration.
 - c) Change the SCSI ID assignment to 0:2.
 - d) Mark the disk for deletion and remove the files from the data store.
 - e) Add another second disk and choose its size.
 - f) Save the VM configuration.
 - g) Apply appropriate settings for disk when it is prompted.
 - h) Start the VM.
5. Run the first-boot procedure by using the VM console.
 - a) Select **I agree** on the IBM Eula screen.
 - b) Set up the passwords for `root` and `builder` accounts.
 - c) Configure the network.
 - d) Restart the node.

With the `builder` account, you can log in to the virtual machine via SSH.

6. Repeat Step 1-5 for each OVA for a minimum of three nodes.
7. Log in to the first node in the cluster that is used to set up or upgrade the cluster across the other nodes. Use the `builder` account and the password that are configured in step [“5” on page 13](#).
8. Copy `/etc/siq/cluster-setup.properties.sample` to `~/cluster-setup.properties`

At the prompt `[builder@localhost ~]$`, enter:

```
cp /etc/siq/cluster-setup.properties.sample cluster-setup.properties
```

9. Edit the file.

At the prompt [builder@localhost ~]\$, enter:

```
vi cluster-setup.properties
```

Configure these properties:

~/cluster-setup.properties	Note
CLUSTER_NAME=Elasticsearch	Cluster name string.
ES_HOSTS=192.0.2.0,192.0.2.10,192.0.2.24	Required. First host becomes node1.
STUNNEL_ENABLED=false	Set to <code>true</code> to have the inter-node communication within the Elasticsearch cluster encrypted using stunnel. If you want to use Search Guard to secure the Elasticsearch cluster, this property must be set to <code>false</code> .
COUNTRY_ABBR=US	Used for x509 certificate generation.
STATE_PROVINCE=Texas	Used for x509 certificate generation.
ES_USER=elasticsearch	Do not modify.
CITY=Austin	Used for x509 certificate generation.
COMPANY_NAME=IBM	Used for x509 certificate generation.
DEPT_NAME=StoredIQ	Used for x509 certificate generation.
CONTACT_EMAIL=storediqsupport@us.ibm.com	Used for x509 certificate generation.
EXPIRATION_DAY=3650	Used for x509 certificate generation.
LUKS=true	Do not modify.
SEARCHGUARD_ENABLED=false	Set to <code>true</code> to secure all communication with and within the Elasticsearch cluster by using Search Guard. Use of Search Guard requires additional configuration steps. For details, see “Securing Elasticsearch cluster communication” on page 40.

10. Run the cluster setup with the properties file that you created.

At the prompt [builder@localhost ~]\$, enter:

```
/siq/bin/cluster-setup.sh cluster-setup.properties
```

The setup script generates `properties/configs` for each node and runs the setup against each node by using SSH.

11. Accept the server certificate and enter the password for each node when prompted.
After the script is executed successfully, the setup of the Elasticsearch cluster is complete.

Data is stored in the `/siq/var/data/elasticsearch` directory. Log files are written to the `/siq/var/log/elasticsearch` directory.

12. Test the setup by using these commands:

```
curl -X GET 'http://primary_node_ip:9200'
curl -X GET 'http://primary_node_ip:9200/_cluster/health?pretty=true'
curl -X GET 'http://primary_node_ip:9200/_cluster/state?pretty'
```

Install the other IBM StoredIQ components.

Deploying a single node Elasticsearch cluster

Deploy a single node Elasticsearch cluster, for example, in a non-production environment.

Before you run the data server first-boot process, complete the setup of the Elasticsearch cluster.

1. Download the Elasticsearch OVA from [IBM Fix Central](#) and deploy the most recent Elasticsearch OVA to vcenter.
2. Edit the VM settings that were created from the OVA.
3. Optional: Reduce the memory.
4. Optional: Reduce the data disk.
 - a) Edit the VM.
 - b) Edit the second disk by expanding the configuration.
 - c) Change the SCSI ID assignment to 0:2.
 - d) Mark the disk for deletion and remove the files from the data store.
 - e) Add another second disk and choose its size.
 - f) Save the VM configuration.
 - g) Apply appropriate settings for disk when it is prompted.
 - h) Start the VM.
5. Run the first-boot procedure by using the VM console.
 - a) Select **I agree** on the IBM Eula screen.
 - b) Set up the passwords for `root` and `builder` accounts.
 - c) Configure the network.
 - d) Restart the node.

With the `builder` account, you can log in to the virtual machine via SSH.

6. Log in to the Elasticsearch node. Use the `builder` account and the password that are configured in step “5” on page 15.
7. Copy `/etc/siq/cluster-setup.properties.sample` to `~/cluster-setup.properties`
At the prompt `[builder@localhost ~]$`, enter:

```
cp /etc/siq/cluster-setup.properties.sample cluster-setup.properties
```

8. Edit the properties file:

```
vi cluster-setup.properties
```

Configure these properties:

~/cluster-setup.properties	Note
CLUSTER_NAME=Elasticsearch	Cluster name string.
ES_HOSTS=192.0.2.0	Required. Enter a single IP address.
STUNNEL_ENABLED=false	If you want to use Search Guard to encrypt the communication between the data server and the Elasticsearch node, this property must be set to false.

~/cluster-setup.properties	Note
COUNTRY_ABBR=US	Used for x509 certificate generation.
STATE_PROVINCE=Texas	Used for x509 certificate generation.
ES_USER=elasticsearch	Do not modify.
CITY=Austin	Used for x509 certificate generation.
COMPANY_NAME=IBM	Used for x509 certificate generation.
DEPT_NAME=StoredIQ	Used for x509 certificate generation.
CONTACT_EMAIL=storediqsupport@us.ibm.com	Used for x509 certificate generation.
EXPIRATION_DAY=3650	Used for x509 certificate generation.
LUKS=true	Do not modify.
SEARCHGUARD_ENABLED=false	Set to <code>true</code> to secure all communication with the Elasticsearch node by using Search Guard. Use of Search Guard requires additional configuration steps. For details, see “Securing Elasticsearch cluster communication” on page 40.

9. Run the cluster setup with the properties file that you created and the `--single-node` option:

```
/siq/bin/cluster-setup.sh cluster-setup.properties --single-node
```

The setup script generates `properties/configs` for the node.

10. Accept the server certificate and enter the password for the node when prompted.
After the script is executed successfully, the setup of the Elasticsearch cluster is complete.

Data is stored in the `/siq/var/data/elasticsearch` directory. Log files are written to the `/siq/var/log/elasticsearch` directory.

11. Stop Elasticsearch by using this command:

```
sudo systemctl stop elasticsearch
```

12. Edit the `docker-compose.yml` file to set the minimum number of master nodes to 1:

```
vi /siq/env/docker/docker-compose.yml
```

Make sure the `discovery.zen.minimum_master_nodes` property is set to 1. Then, save the change and exit the file.

13. Start Elasticsearch by using this command:

```
sudo systemctl start elasticsearch
```

14. Test the setup by using these commands:

```
curl -X GET 'http://node_ip:9200'
curl -X GET 'http://node_ip:9200/_cluster/health?pretty=true'
curl -X GET 'http://node_ip:9200/_cluster/state?pretty'
```

Because the cluster contains only one node, replicas are not enabled. Therefore, the node status will be yellow. To update the node status to green, set the number of replicas to 0 for each index:

```
curl -X PUT "http://node_ip:9200/volume_*/_settings" -d '{"index":{"number_of_replicas" : 0 }},'
```

Install the other IBM StoredIQ components.

Adding or removing a node from an existing Elasticsearch cluster

You can add nodes to any Elasticsearch cluster. Removing a node is possible only if the cluster has at least four nodes.

To ensure that the cluster works correctly and all nodes are available, enter

```
curl -X GET 'http://primary_node_ip:9200/_cluster/state?pretty'
curl -X GET 'http://primary_node_ip:9200/_nodes?pretty'
```

Where *primary_node_ip* is the IP address of either the single Elasticsearch node or the primary Elasticsearch node in a cluster of at least three nodes.

Adding a node

You can add a single node to any Elasticsearch cluster with three or more nodes. With a single node deployment, you must add at least two nodes as to create a default three node Elasticsearch cluster.

1. Download the Elasticsearch OVA version that matches the current cluster version from Fix Central and deploy it to vcenter.
2. Edit the VM settings that were created from the OVA.
3. Optional: Reduce the memory.
4. Optional: Reduce the data disk.
 - a) Edit the VM.
 - b) Edit the second disk by expanding the configuration.
 - c) Change the SCSI ID assignment to 0:2.
 - d) Mark the disk for deletion and remove the files from the data store.
 - e) Add another second disk and choose its size.
 - f) Save the VM configuration.
 - g) Apply appropriate settings for disk when it is prompted.
 - h) Start the VM.
5. Run the first-boot procedure by using the VM console.
 - a) Select **I agree** on the IBM Eula screen.
 - b) Set up the passwords for **root** and **builder** accounts.
 - c) Configure the network.
 - d) Restart the node.

With the **builder** account, you can log in to the virtual machine via SSH.

6. Log in to the ster node, where the installation was initially started, by using the **builder** account and password that were configured during the installation of the cluster.
7. Edit the `cluster-setup.properties` file by appending the new node's IP address to the `ES_HOSTS` line. For example, to add a fourth node to a three node cluster, the `ES_HOSTS` line looks like:

```
ES_HOSTS=192.0.2.0,192.0.2.10,192.0.2.24,198.51.100.0
```

, where 198.51.100.0 is the new node's IP address.

8. Run the cluster setup script. At the prompt `[builder@localhost ~]$`, enter:

```
/siq/bin/cluster-setup.sh cluster-setup.properties
```

9. Accept the server certificate and enter the password for the new node when prompted.
After the script is executed successfully, the new node is part of the Elasticsearch cluster.
10. Add the new node's IP address and port information to `/usr/local/storediq/etc/siq-elasticsearch.yml` on the data servers that use this Elasticsearch cluster.

11. Enter the command to restart services on the data servers that are using this Elasticsearch cluster:

```
service deepfiler restart
```

12. To ensure that the cluster works correctly and all nodes are available, enter:

```
curl -X GET 'http://primary_node_ip:9200'
curl -X GET 'http://primary_node_ip:9200/_cluster/health?pretty=true'
curl -X GET 'http://primary_node_ip:9200/_cluster/state?pretty'
```

Removing a node

You can remove nodes only from clusters with at least four nodes.

1. Log in to the primary node, where the installation was initially started, by using the **builder** account and password that were configured during the installation of the cluster.
2. Decommission a node by entering:

```
curl -XPUT localhost:9200/_cluster/settings -H 'Content-Type: application/json' -d '{
  "transient" : {
    "cluster.routing.allocation.exclude._ip" : "node_ip"
  }
};'
```

where *node_ip* is the IP address of the node that is being decommissioned.

The response must look like this:

```
{
  "acknowledged": true,
  "persistent": {},
  "transient": {
    "cluster": {
      "routing": {
        "allocation": {
          "exclude": {
            "_ip": "node_ip"
          }
        }
      }
    }
  }
}
```

3. Edit the `cluster-setup.properties` file and remove the decommissioned node's IP address from the `ES_HOSTS` line.
4. Run the cluster setup script. At the prompt `[builder@localhost ~]$`, enter:

`/siq/bin/cluster-setup.sh cluster-setup.properties`
5. Accept the server certificate and enter the password for each node when prompted.
6. Enter the following commands to ensure that the cluster works correctly and all nodes are available (the decommissioned node is still listed, but no shards are allocated to it):

```
curl -X GET 'http://primary_node_ip:9200'
curl -X GET 'http://primary_node_ip:9200/_cluster/health?pretty=true'
curl -X GET 'http://primary_node_ip:9200/_cluster/state?pretty'
```

7. Shut down and delete the VM in vCenter.
8. Remove the decommissioned node's IP address and port information from `/usr/local/storediq/etc/siq-elasticsearch.yml` on the data servers that use this Elasticsearch cluster.
9. Enter the command to restart services on the data servers that are using this Elasticsearch cluster:

```
service deepfiler restart
```

10. Ensure that the cluster works correctly and all nodes are available (now, the decommissioned node is not listed).

```
curl -X GET 'http://primary_node_ip:9200'
curl -X GET 'http://primary_node_ip:9200/_cluster/health?pretty=true'
curl -X GET 'http://primary_node_ip:9200/_cluster/state?pretty'
```

Installing the gateway

To install the gateway, you must first turn on the virtual machine and then configure its settings.

- Keep in mind that one IP address cannot be used simultaneously by different systems.

- It is helpful to know your application stack's IP address before you install the gateway. If you do not have this information, the gateway's configurations can be edited to add this information later.
1. Within the vCenter console, click **Launch Console** to start the virtual machine console.
 2. Within the **CentOS EULA** screen, select **I Agree** to accept the CentOS license agreement, and then press **Enter**.
 3. Within the **IBM EULA** screen, select **I Agree** to accept the IBM license agreement, and then press **Enter**.
 4. Within the **Corporate Network** screen, select either **Static** or **Obtain IP via DHCP**.
 - If you select the **Static** option, complete these fields to configure the static address:
 - a. Gather this information to complete the static IP configuration for this system: **Host name, Corporate IP, Netmask, Gateway, and Primary DNS Host**.
 - b. In the **Hostname** field, enter a fully qualified host name and then press **Enter**.
 - c. In the **Corporate IP** field, enter an IPv4 address and then press **Enter**.
 - d. In the **Netmask** field, enter the netmask for the assigned IP address and then press **Enter**.
 - e. In the **Gateway** field, enter the IP address of the default gateway for the IP subnet and then press **Enter**.
 - Note:** This gateway is the network gateway, not the IBM StoredIQ gateway.
 - f. In the **Primary DNS Host** field, enter the IP address for the domain name server and then press **Enter**.
 - If you select the **Obtain IP via DHCP** option, complete these fields:
 - a. Gather the host name information to complete the DHCP IP configuration for the system.
 - b. Select **Hostname** and enter a host name.
 5. Click **Next**.
The **Enter root password** prompt appears.
 6. Enter the new password twice, and then press **Enter** when finished.
The installation of the IBM StoredIQ application starts.
Note: Depending on the performance of your virtual infrastructure, this process can be lengthy, taking from just minutes to multiple hours.
The **Select FIPS mode** screen appears.
 7. Select this option to run your system in FIPS-compliant security mode.
By default, FIPS mode is not enabled.
 8. In the **Gateway communication security** screen, optionally select **Enable secure gateway communication** to encrypt the communication between the gateway, the data server, and the application stack.
By default, the communication is in plain text and is not encrypted.
Secure communication via stunnel can impact performance. Therefore, enable this setting only if your enterprise security policy mandates encryption of data in motion. If you do so, the IBM StoredIQ application stack and the data server must be configured accordingly. Also, check the gateway settings as described in the section about configuring DA Gateway settings in the data server administration documentation.
You can change the enablement status at any time after the installation.
For more information, see [“Security” on page 7](#). For additional guidance, contact IBM Support.
 9. Select **Done**.
 10. The **Appstack IP** prompt, enter the application stack's IP address.
After the IP address is entered, the installation is finalized and the server restarts.

Installing the data server

To install the data server, you must first turn on the virtual machine and then configure its settings.

IP addresses cannot be used for multiple installations on the same system.

1. Within the vCenter console, click **Launch Console** to start the virtual machine console.
2. Accept the CentOS and IBM license agreements when prompted.
3. Within the **Corporate Network** screen, select either **Static** or **Obtain IP via DHCP**.
 - If you select the **Static IP** option, complete these fields to configure the static address:

Parameter	Value
Hostname	The fully qualified host name
Corporate IP	The IPv4 address of the application stack
Netmask	The netmask for the assigned IP address
Gateway	The IP address of the default gateway for the IP subnet Note: This is the network gateway, not the IBM StoredIQ gateway.
Primary DNS Host	The IP address for the domain name server

- If you select the **Obtain IP via DHCP** option, complete these fields:

Parameter	Value
Hostname	The fully qualified host name

4. Click **Next**.
The **Enter root password** prompt appears.
5. Enter the new password twice.
The **Select FIPS mode** screen appears.
6. Optional: Select this option to run your system in FIPS-compliant security mode.
By default, FIPS mode is not enabled.
7. If the data server type that is deployed is DataServer - Distributed, select **Enable elasticsearch storage**.
For a default Elasticsearch deployment, enter the IP address of the first node of the Elasticsearch cluster where the cluster was set up. For a single node Elasticsearch deployment, enter the IP address of the Elasticsearch node.
8. Select **Next**.
9. Optional: Select **Enable secure gateway communication** to encrypt the communication between the data server and the gateway.
By default, the communication is in plain text and is not encrypted.

Secure communication via stunnel can impact performance. Therefore, enable this setting only if your enterprise security policy mandates encryption of data in motion. If you do so, the IBM StoredIQ gateway and the application stack must be configured accordingly.

You can change the enablement status at any time after the installation.

For more information, see [“Security” on page 7](#). For additional guidance, contact IBM Support.
10. Select **Done**.
11. Press **Enter** when it is finished.
The installation of the IBM StoredIQ application starts.

Note: Depending on the performance of your virtual infrastructure, this process can be lengthy, taking from just minutes to multiple hours.

- Optional: For a data server of the type DataServer - Distributed, verify that the `/usr/local/storediq/etc/siq-elasticsearch.yml` exists and contains the proper information.

For a default Elasticsearch deployment, the `-host:` entry in the file must show the IP addresses of all nodes in the Elasticsearch cluster. With a single node Elasticsearch cluster setup, the entry must contain the IP address of this Elasticsearch node.

If the file does not exist on the data server, copy the `/etc/siq/siq-elasticsearch.yml` from the Elasticsearch node to the data server. In a three node cluster, you can find the file on the primary node.

If **Secure gateway communication** is enabled, you must take the following steps after the data server virtual machine and the deepfiler services are up:

- Log in to the IBM StoredIQ Data Server user interface.
- Go to **Configuration > System > DA Gateway settings**.
- The host shows 127.0.0.1
- Set the node name and click **OK**.
- Restart services in either of the following ways:
 - Using an SSH tool, log in as root and then type `service deepfiler restart`.
 - Go to the Data Server dashboard, click **About appliance**, and then click **Restart services**.

Installing and configuring the application stack

Follow these steps to install and configure the application stack.

You should have both the application stack's static IP and gateway IP address at hand. While the application stack's IP address must be specified during the installation, you can add the gateway IP address to the gateway configuration later.

If you plan to enable the synchronization with a governance catalog, a working deployment of one of these products must be available:

- IBM Information Server. The minimum required version is Version 11.7.0.
- IBM Cloud Private for Data. The minimum required version is Version 1.1.0.2.

For details about this feature, see the information about integrating with IBM Information Governance Catalog.

- Power on the virtual machine.
- Within the vCenter console, click **Launch Console** to start the virtual-machine console.
- Accept the CentOS and IBM license agreements when prompted.
- In the **Password Utility** screen, enter the new passwords for the root user and the siqadmin user twice and press **Enter**.
- Within the **Corporate Network** screen, select either **Static IP** or **Obtain IP via DHCP**.
 - If you select the **Static IP** option, complete these fields to configure the static address:

Parameter	Value
Hostname	The fully qualified host name
IP Address	The IPv4 address of the application stack
Netmask	The netmask for the assigned IP address
Gateway	The IP address of the default gateway for the IP subnet

Parameter	Value
	Note: This is the network gateway, not the IBM StoredIQ gateway.
Primary DNS Host	The IP address for the domain name server

- If you select the **Obtain IP via DHCP** option, complete these fields:

Parameter	Value
Hostname	The fully qualified host name
Primary DNS Host	The IP address for the domain name server

6. Select the **Restart services** option to commit the IP and restart services.

7. Click **Next**.

Within the **Appstack configuration** screen, set these options.

8. In the **Domain name** field, enter the fully qualified domain name (FQDN) or IP address of the application stack.

This information is used in generated URLs, such as links to reports.

If you enable the synchronization with the governance catalog, the domain name is used for building the base URLs for REST access to the application stack and for links to IBM StoredIQ artifacts that make these artifacts accessible from Information Governance Catalog.

Important: If you ever need to change the host name or IP address of the application stack (using the **appstackcfg** utility), you must restart all application stack services afterward by running the command `service appstack restart` from the command line. Selecting the **Restart appstack services** option is not sufficient because this option triggers the restart of only the **uwsgi** and **tomcat** services.

9. In the **StoredIQ Gateway** field, enter the IP address of the StoredIQ Gateway server.
10. Set the following SMTP options to enable the application stack's capability to send and receive notification email.

Tip: All SMTP settings are optional and can be configured before or after your deployment. If you choose to set or change the SMTP settings at a later time, see [“Configuring the application stack to receive reports and notifications”](#) on page 43.

- a) Set these options:

Parameter	Value
Server	The mail server's fully qualified domain name
Port	The SMTP port
Username	The login user name. For the default configuration, leave this field empty. Otherwise, enter the credentials of an authenticated user.
Password	The login password for the specified user. For the default configuration, leave this field empty.

- b) Select **Enable TLS** to enable TLS encryption, if it is supported by the mail server.

For additional information about SMTP notification, see [“Configuring authenticated users for SMTP notifications”](#) on page 55.

11. Optional: Set these options to enable the synchronization of specific objects between IBM StoredIQ and Information Governance Catalog.

Tip: These settings can be configured during or after the deployment. If you choose to set or change the synchronization settings at a later time, see [“Configuring the application stack to synchronize data with the governance catalog”](#) on page 43.

If the synchronization is not enabled, the values entered here are not validated. However, as soon as you enable synchronization, all entries must be valid. Otherwise, a warning is displayed and synchronization is implicitly disabled.

- a) Select **Enable synchronization with the governance catalog** and provide the following settings.

If the data catalog to which you want to publish the IBM StoredIQ object resides in an IBM Cloud Private for Data environment, select the **Server runs in IBM Cloud Private for Data** check box.

Provide or accept the values for these fields:

Parameter	Value
Host	The host name or IP address of the Information Server or IBM Cloud Private for Data installation. The specified host is part of the base URLs for REST access to the Information Governance Catalog or IBM Cloud Private for Data instance and for links to catalog artifacts that make these artifacts accessible from IBM StoredIQ. Therefore, you should provide the fully qualified domain name of the Information Server or IBM Cloud Private for Data host. If you specify a server port, this port also becomes part of such base URLs. In addition, the host name is also used to address the Information Server Apache Kafka server, which provides all Information Server events as Kafka messages. Specific Kafka messages are consumed by IBM StoredIQ and trigger the synchronization of objects from the governance catalog to IBM StoredIQ.
Port	The port of the Information Governance Catalog server. This setting is optional. For connections to an Information Server 11.7 environment: To ensure proper communication, you should set the port to the HTTPS port that is defined in Information Server. The default port is 9443. For connections to an Information Server 11.7 FP1 environment without Information Server Enterprise Search installed: To ensure proper communication, you should set the port to the HTTPS port that is defined in Information Server. The default port is 9446.. For connections to an Information Server 11.7 FP1 environment with Information Server Enterprise Search installed: Do not specify a port.

Parameter	Value
	For connections to IBM Cloud Private for Data, you can specify the IBM Cloud Private for Data port.
Kafka port	The port of the Information Server Kafka server. The port setting can be overridden. For connections to an Information Server 11.7 environment: the Kafka port defined in Information Server. The default port is 59092. For connections to an Information Server 11.7 FP1 environment without Information Server Enterprise Search installed: the Kafka port defined in Information Server. The default port is 59092. For connections to an Information Server 11.7 FP1 environment with Information Server Enterprise Search installed: the Kafka port defined in Information Server. The default port is 9092. For publishing to IBM Cloud Private for Data, this setting is not used. However, an error will be logged, which you can ignore.
Username	The user name for authenticating to Information Server when publishing IBM StoredIQ objects to the governance catalog. This user must be defined in Information Server with the following security roles: - Suite User - Common Metadata Administrator - Information Governance Catalog - Information Asset Administrator In IBM Cloud Private for Data, this user must be defined with the Data Stewart role.
Password	The password of the user set with Username .
Sync frequency (minutes)	Data is periodically propagated to the governance catalog at the specified interval. The value must be a positive number of minutes. The default value is 15 minutes.
StoredIQ instance name	The name identifying the IBM StoredIQ instance for which data is synchronized. This name can be freely chosen, but must be unique within the Information Governance Catalog instance.

12. Optional: Select **Enable FIPS mode at boot time** to enable running your system in FIPS-compliant security mode. By default, FIPS mode is not enabled.
13. Optional: Select **Enable secure gateway communication** to encrypt the communication between the application stack and the gateway.
By default, the communication is in plain text and is not encrypted.

Secure communication via stunnel can impact performance. Therefore, enable this setting only if your enterprise security policy mandates encryption of data in motion. If you do so, the IBM StoredIQ gateway and the data server must be configured accordingly. For more information, see [“Security” on page 7](#). For additional guidance, contact IBM Support.

You can change the enablement status at any time after the installation.

14. Click **Next**.

Important: For synchronization with the governance catalog to work, HTTPS must be enabled on the AppStack. Therefore, generate and install at least a self-signed certificate.

15. Optional: Within the **Certificate configuration** screen, perform the procedure in its entirety to generate a self-signed SSL or TLS certificate.

SSL or TLS certificates are used to establish secure communications. You can generate self-signed certificates, which should be used in test and development environments only, or certificates that are signed by an internal or a third-party certificate authority (CA). To avoid certificate trust issues, you should obtain and install a certificate signed by a third-party CA.

To skip certificate configuration, tab to **Exit** and click **Enter**. The certificate can be configured at a later time by logging in as siqadmin user and running this command: certcfg

Important: Synchronization with the governance catalog requires a certificate to be installed. Therefore, do not skip certificate generation now if you enabled the synchronization.

If you choose to generate a certificate, complete the steps of this procedure in the described sequence. In the wizard, use the Up and Down Arrow keys to navigate between options and the space bar to select an option.

a) Generate a self-signed root certificate. Make sure that option 1 is selected and press Enter.

The resulting certificate can be used as a certificate authority (CA).

If you want to use a root certificate from a third party CA to sign your certificates, you can skip this step.

The following table lists the configuration settings for a self-signed root certificate. Required settings are denoted by an asterisk. Edit the settings as required.

Table 1: Creating self-signed root certificate	
Setting	Notes
Common Name *	The name of the certificate. You can use the prefilled value or choose a different name. However, make sure not to use the AppStack host name.
Email Address *	The email address that is used in the certificate.
Country (two-letter) *	An acceptable entry is an ISO-3166-1 alpha-2 code. A listing is available here .
State/Province *	
City *	
Organization *	
Department *	
Key length	The length of the key to be created. The default value is 2048.
Days to expiry *	The number of days before the certificate expires. The default value is 3650.

<i>Table 1: Creating self-signed root certificate (continued)</i>	
Setting	Notes
Root key location *	The fully qualified file name of the root key file. This name can be freely chosen. If the file does not exist, it is created when the certificate is created. However, it is recommended to use the prefilled default file name.
Root certificate location	The fully qualified file name of the root certificate file. This name can be freely chosen. If the file does not exist, it is created when the certificate is created. However, it is recommended to use the prefilled default file name.

Click **Next** to proceed to creating a certificate-signing request.

- b) Create a certificate-signing request. Make sure that option 2 is selected and press Enter.

Generate a certificate signing request (CSR) to be signed by a certificate authority. The process creates a key or uses a provided key and generates the CSR from it.

The following table lists the settings for a certificate signing request. Required settings are denoted by an asterisk. Edit the settings as appropriate.

<i>Table 2: Creating certificate signing request</i>	
Setting	Notes
Common Name *	The host name of the AppStack. It must match the domain of the URL that you use. Important: This value is prefilled. If you create a self-signed root certificate, make sure to change this value so that it is different from the common name of Step “15.a” on page 25.
Email Address *	
Country (two-letter) *	An acceptable entry is an ISO-3166-1 alpha-2 code. A listing is available here .
State/Province *	
City *	
Organization *	
Department *	
Key length	The length of the key to be created. The default value is 2048.
Key location *	The fully qualified file name of the key file. This name can be freely chosen. If the file does not exist, it is created when the certificate is created. However, it is recommended to use the prefilled default file name.

<i>Table 2: Creating certificate signing request (continued)</i>	
Setting	Notes
Certificate request location *	The fully qualified file name of the CSR file. This name can be freely chosen. If the file does not exist, it is created when the request is created. However, it is recommended to use the prefilled default file name.

Click **Next** to proceed to signing the certificate.

- c) Generate the signed certificate. Make sure that option 3 is selected and press Enter.

Sign a certificate with a certificate-authority-eligible root certificate based on the certificate-signing request.

The following table lists the settings for signing the certificate. All settings are required. Edit them as appropriate.

<i>Table 3: Signing request with root certificate</i>	
Setting	Notes
Days to expire	The number of days before the certificate expires. The default value is 3650.
Certificate request location	The fully qualified file name of the CSR to sign as specified in step “15.b” on page 26.
Certificate location	The fully qualified file name of the signed certificate. This name can be freely chosen. If the file does not exist, it is created when the certificate is created. However, it is recommended to use the prefilled default file name.
Root key location	The fully qualified file name of the root key file as specified in step “15.a” on page 25. Or, if you chose to use a root certificate signed by a third-party CA, the fully qualified file name of the respective root key file. However, it is recommended to create a copy of the third-party root key file with the default name assigned by IBM StoredIQ.
Root certificate location	The fully qualified file name of the root certificate as specified in step “15.a” on page 25. Or, if you chose to use a root certificate signed by a third-party CA, the fully qualified file name of the respective root certificate. However, it is recommended to create a copy of the third-party root certificate with the default name assigned by IBM StoredIQ.

Click **Next** to proceed to updating the AppStack HTTPS certificate.

- d) Update the AppStack HTTPS certificate. Make sure that option 4 is selected and press Enter.

Update the application stack to use the provided certificate and key for HTTPS access.

The following table lists the settings for updating the certificate. These settings are required and are prefilled with the information from the previous steps. Do not change these settings.

<i>Table 4: Updating the AppStack HTTPS certificate</i>	
Setting	Notes
Key location	The fully qualified file name of the key file.
Certificate location	The fully qualified file name of the certificate.

Click **Finish** to complete the certificate configuration.

The application stack installation begins. When the installation is done, it restarts to open a console login prompt. The installation and configuration of the application stack is complete.

You can open a browser to log in to the IBM StoredIQ applications.

Logging in to the application stack user interfaces

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

The system administrator must ensure that these tasks are completed:

- Install and configure IBM StoredIQ.
- Ensure that the data servers are up and running.

The default administrator account is `superadmin`. The default password for this account is `admin`. For security purposes, change the password as soon as possible. Also, create extra administrators for routine administration so that their actions can be audited.

1. In a browser, enter the IP address or host name of the application stack server.
2. Log in with your IBM StoredIQ credentials.

The `superadmin` account has access to all IBM StoredIQ applications on the application stack. To switch applications, click the triangle arrow icon and select the application that you want to access from the list of the available applications.

Configuration of IBM StoredIQ

After you deploy virtual appliances, you must configure IBM StoredIQ before you can fulfill the product installation process.

Default open ports

The following ports are open by default on the IBM StoredIQ.

Default open ports on the AppStack

Port number	Protocol
22	tcp
80	tcp
443	tcp

Default open ports on the IBM StoredIQ data server

Port number	Protocol	Service
22	tcp	PROD-ssh
80	tcp	PROD-web
443	tcp	PROD-https (UI and Web Services APIs)
11103	tcp	PROD-transport (IBM StoredIQ transport services; communication between the gateway and the data server)
11104		

Enable or disable ports or services on the IBM StoredIQ data server

To manage ports, you can use the `/usr/local/storediq/bin/util/port_handler.pyc` script with the appropriate parameter:

```
python /usr/local/storediq/bin/util/port_handler.pyc -parameter
```

-s

To list the current rules in iptables

-l

To list the supported services

-d port_number|port_range'

To delete a port or a range of port numbers from iptables, for example:

```
python /usr/local/storediq/bin/util/port_handler.pyc -d '21200:21299'
```

-e 'service_name'

To enable a specific service, for example, to enable HTTPS services:

```
python /usr/local/storediq/bin/util/port_handler.pyc -e 'PROD-https'
```

-d 'service_name'

To disable a specific service, for example, to disable HTTPS services:

```
python /usr/local/storediq/bin/util/port_handler.pyc -d 'PROD-https'
```

Default open ports on the nodes in the Elasticsearch cluster

Port number	Protocol
22	tcp
80	tcp
443	tcp

Default open ports on the IBM StoredIQ gateway

Port number	Protocol	Service
22	tcp	PROD-ssh
80	tcp	PROD-web
443	tcp	PROD-https (UI and Web Services APIs)
5432	tcp	PROD-postgres
5434	tcp	PROD-transport (IBM StoredIQ transport services; communication between the gateway and the data server)
8765		
7766		
11102		
11103		
11104		

Supported chain and rules on the IBM StoredIQ gateway

In iptables, the following firewall and chain rules are defined:

```
'PROD-transport': ['5434', '8765', '7766', '11102', '11103', '11104'],
'PROD-https': ['443'],
'PROD-ssh': ['22'],
'PROD-web': ['80'],
'PROD-postgres': ['5432']
```

```
'desktop' service:
'PROD-broker': ['21000'],
'PROD-collectionsvc': ['21300:21399'],
'PROD-desktopupgrade': ['21004'],
'PROD-objlistmgr': ['21100:21199'],
'PROD-objlistsvc': ['21200:21299'],
'PROD-registration': ['21001'],
```

```
'PROD-session': ['21002'],  
'PROD-task': ['21003'],
```

Open ports for desktop client access to the data server

To open ports for desktop client access to the data server on OVA deployed systems, follow these steps:

1. Log in to the data server as root and run this command:

```
python /usr/local/storediq/bin/util/port_handler.pyc -e desktop
```

2. Run this command: `iptables -L INPUT`

In the output of the command, check the list position of the rule that is named `PROD-reject`, for example, the 6th position on the list.

3. Run this command: `iptables -A INPUT -j PROD-reject`

4. Run this command: `iptables -D INPUT list_position`

list_position is the position number of the `PROD-reject` rule that you determined in step 2.

5. Run the following command:

```
python /usr/local/storediq/bin/util/port_handler.pyc -e desktop
```

Tip: These steps are required only on an IBM StoredIQ OVA deployed system. The correct ports are open on an upgraded system.

Key and certificate management

IBM StoredIQ uses keys and certificates to ensure secure communication and encrypt data.

The following table shows the key files, certificates, and configuration files used by the AppStack, the data servers, and the gateway for secure communication. You can replace all keys and self-signed certificates that IBM StoredIQ uses.

Table 5: List of keys and certificates.			
Component	Key or certificate location	Configuration file	Comment
AppStack Application	/etc/siq/ssl/client.crt	/siq/svc/nginx/ssl.conf	Run the certcfg utility as siqadmin user to manage certificates.
Data Server Web application	/usr/local/apache/conf/ssl.crt/server.crt	usr/local/apache/conf/httpd.conf	Hashing algorithm: PKCS #1 SHA-256 with RSA

Table 5: List of keys and certificates. (continued)

Component	Key or certificate location	Configuration file	Comment
AppStack stunnel	/etc/siq/ssl/ stunnel.pem	/etc/siq/ stunnel.conf	If you want to use your own certificates, see “Installing your own certificates” on page 34 for details. For details about stunnel key management, see “Managing disk encryption keys for Elasticsearch index volumes” on page 33.
Gateway stunnel	/etc/deepfile/ gateway/ stunnel.pem	/etc/deepfile/ gateway/ stunnel.conf	
Data Server stunnel	/etc/deepfile/ dataserver/ stunnel.pem	/etc/deepfile/ dataserver/ stunnel.conf	
ElasticSearch stunnel	/siq/env/docker/ node1-stunnel.pem	/siq/env/docker/ node1-stunnel.conf	

Managing disk encryption keys for Elasticsearch index volumes

IBM StoredIQ uses LUKS to encrypt the disk volume on which the Elasticsearch indexes are stored. You can add your own private keys or passphrases to the LUKS encryption system, or even remove the encryption key that is generated when the cluster is set up.

To protect the private key from unauthorized access, you must be an administrator with root access to be able to complete this procedure.

LUKS provides eight slots to specify the encryption keys. You can manage keys and passphrases by using the **cryptsetup** command.

1. Check which device is encrypted by running the following command:

```
[builder@hostname ~]$ su
Password:
[root@hostname builder]$ blkid -t TYPE=crypto_LUKS
```

The output looks similar to what is shown in this example:

```
/dev/sdb1: UUID="3755df51-cf96-46d9-b5b1-b301a0284bc8" TYPE="crypto_LUKS"
PARTLABEL="primary" PARTUUID="66715c0d-0e22-4e44-a5a3-cd5cd9bbdb5f"
```

By default, the encrypted device should be `/dev/sdb1`. If you converted the OVA to a `.vhdx` file for deployment to a Microsoft Hyper-V server, the encrypted device might be `/dev/sdc1`.

2. Check the LUKS key slots to detect which slots are enabled and which slots are disabled.

```
[root@hostname ~]# cryptsetup luksDump /dev/sdb1 | grep "Key Slot"
```

The output looks similar to what is shown in the example:

```
Key Slot 0: ENABLED
Key Slot 1: DISABLED
Key Slot 2: DISABLED
Key Slot 3: ...
```

3. Extract the private keys used for encrypting the device and convert them into a binary format.

Important: You must keep this key file protected from unauthorized access. Otherwise, security is compromised.

The **xxd** utility is used to convert the key file from text to binary. This utility is not available on the Elasticsearch nodes but only on the data server. Therefore, the conversion must happen on the data server.

a) Extract the private keys:

```
[root@hostname]# dmsetup table --showkeys|grep siq-elasticsearch-luks | gawk '{ print $6}' > key.txt
```

b) Copy the key .txt file to the data server by using a tool such as **scp** or **sftp**.

c) Convert the file from text to binary.

```
[root@hostname]# xxd -r -p ./key.txt key.bin
```

d) Copy the resulting key .bin file to that Elasticsearch node from where you copied the key .txt file.

e) Check the content of the key .bin file.

Use the following command to make sure that the binary file has the same information as the text file.

```
[root@hostname]# hexdump -C key.bin
```

4. Add a new passphrase to the next available key slot.

```
[root@hostname]# cryptsetup luksAddKey /dev/sdb1 --master-key-file <(cat key.bin)
Enter new passphrase for key slot:
Verify passphrase:
```

Check again which key slots are enabled and which ones are disabled:

```
[root@hostname]# cryptsetup luksDump /dev/sdb1 | grep "Key Slot"
Key Slot 0: ENABLED
Key Slot 1: ENABLED
Key Slot 2: DISABLED
```

5. Optional: Add a new key to a selected LUKS key slot.

```
[root@hostname ~]# cryptsetup luksAddKey /dev/sdb1 secret.Key -S 3 --master-key-file <(cat key.bin)
[root@hostname ~]# cryptsetup luksDump /dev/sdb1 | grep "Key Slot"
Key Slot 0: ENABLED
Key Slot 1: ENABLED
Key Slot 2: DISABLED
Key Slot 3: ENABLED
Key Slot 4: DISABLED
```

6. Optional: Remove the key that is provided by LUKS for immediate use.

```
[root@hostname ~]# cryptsetup luksKillSlot /dev/sdb1 0
```

Important: If you remove a key without adding a new key or a passphrase, you will lock yourself out of the encrypted device. In this case, you will no longer be able to access the device and the data that is stored on this encrypted device is permanently lost.

Installing your own certificates

You can install and use your own certificates when secure gateway communication using stunnel is enabled.

Back up your virtual machines before replacing the certificates. The following procedure requires root authority, even on the AppStack. To perform these steps, siqadmin authority is not sufficient.

If you enable secure gateway communication using stunnel, IBM StoredIQ generates its own self-signed root certificate and uses that one. If you want to use your own root certificates along with the IBM StoredIQ intermediate certificate, follow these instructions.

The following certificates and keys must be available:

- Root certificate
- Intermediate certificate (gateway certificate and key)

- AppStack certificate and key
- Data server certificate and key (one for each data server)
- Certificates for the Elasticsearch nodes (one for each Elasticsearch node)

For some steps of this procedure, you must copy files between servers. Use a tool such as **scp** or **sftp** to do so.

For the gateway, the AppStack, and each data server, complete these steps:

1. Copy the root certificate to the gateway, the AppStack, and each data server.

The `ca.pem` root certificate must be copied to the following locations:

- On the gateway: `/etc/deepfile/gateway/certs`
- On the data server: `/etc/deepfile/dataserver/certs`
- On the AppStack: `/etc/siq/ssl/certs/`

2. Create key pairs on the gateway, the AppStack, and each data server.

A key pair is a combination of a key and a certificate. For example, you can use the **cat** command to create a combined PEM file as shown here:

- On the gateway:

```
cat gateway.key gateway.pem > gateway.keypair.pem
```

- On each data server:

```
cat dataserver.key dataserver.pem > dataserver.keypair.pem
```

- On the AppStack:

```
cat appstack.key appstack.pem > appstack.keypair.pem
```

3. Copy the AppStack and data server certificates to the gateway.

Store all certificates (PEM files) in the certificate directory. Copy the `gateway.pem` file, the `dataserver.pem` file, and the `appstack.pem` file to the `/etc/deepfile/gateway/certs` directory.

4. Copy the gateway certificate to the AppStack and to each data server.

Store the PEM file in the certificate directory.

- On the data server, copy the `gateway.pem` file to the `/etc/deepfile/dataserver/certs` directory.
- On the AppStack, copy the `gateway.pem` file to the `/etc/siq/ssl/certs` directory.

5. Rehash the certificate directories.

Use the following commands:

- On the gateway:

```
/usr/sbin/cacertdir_rehash /etc/deepfile/gateway/certs
```

- On each data server:

```
/usr/sbin/cacertdir_rehash /etc/deepfile/dataserver/certs
```

- On the AppStack:

```
/usr/sbin/cacertdir_rehash /etc/siq/ssl/certs
```

6. Copy the key pairs as `stunnel.pem` file to the appropriate location.

Use the following commands:

- On the gateway:

```
cp gateway.keypair.pem /etc/deepfile/gateway/stunnel.pem
```

- On each data server:

```
cp dataserver.keypair.pem /etc/deepfile/dataserver/stunnel.pem
```

- On the AppStack:

```
cp appstack.keypair.pem /etc/siq/ssl/stunnel.pem
```

7. Restart service as follows:

- a) On the gateway, run this command:

```
service deepfiler restart
```

- b) On each data server, run this command:

```
service deepfiler restart
```

- c) On the AppStack, run these commands:

```
/siq/bin/monit restart uwsgi
/siq/svc/stunnel/stunnel-ctl.sh stop
/siq/svc/stunnel/stunnel-ctl.sh start
```

8. Test the connectivity as follows:

- Connectivity between the gateway and the data server

On the gateway, issue the following API call:

```
curl http://localhost:7766/backchannel/1.0/server
```

- Connectivity between the gateway and the AppStack

On the AppStack, issue the following API call:

```
curl -i http://localhost:8765/administrative/1.0/server
```

9. To install certificates on the Elasticsearch nodes, complete the following steps on each node:

- a) Create a key pair, combining key and certificate:

```
cat node.key node.pem > node.keypair.pem
```

- b) Copy the key pair to the appropriate location.

```
cp node.keypair.pem /siq/env/docker/node1-stunnel.pem
```

- c) Restart the Docker service by using the following commands:

```
cd /siq/env/docker
docker-compose restart
```

- d) To verify that the certificates were accepted, check the `/var/log/stunnel_service.log` log file.

The communication between the nodes within the Elasticsearch cluster as well as the communication between the data server and the Elasticsearch nodes can also be encrypted by using other methods than stunnel. For more information about those other methods, contact IBM Support.

Enabling encryption of IBM StoredIQ AppStack application data

IBM StoredIQ uses Linux Unified Key Setup (LUKS) for encrypting data at rest at the file system level. You can enable disk encryption on the AppStack for security reasons.

You must be logged in to the AppStack as an administrator with root access to be able to complete this procedure.

Complete this procedure for new installations or for update installations where LUKS encryption was not yet enabled. After you complete the procedure, all application data on the AppStack that is stored in the `/var/siq` directory structure is encrypted.

To enable disk encryption:

1. Stop all IBM StoredIQ services on the AppStack:

```
services appstack stop
```

2. Check whether the `cryptsetup-luks` package on the AppStack.

This package contains the **cryptsetup** utility that you can use for setting up and managing encrypted file systems. Use the following command:

```
repoquery -i cryptsetup-luks
```

If the package is installed, the output of this command looks similar to this example:

```
Name : cryptsetup-luks
Version : 1.2.0
Release : 7.el6
Architecture: x86_64
Size : 302744
Packager : CentOS BuildSystem <http://bugs.centos.org>
Group : Applications/System
URL : http://cryptsetup.googlecode.com/
Repository : centos-65-x86_64
Summary : A utility for setting up encrypted filesystems
Source : cryptsetup-luks-1.2.0-7.el6.src.rpm
Description :
This package contains cryptsetup, a utility for setting up encrypted filesystems using
Device Mapper and the dm-crypt target.
```

If the package is not yet installed, install it by running this command:

```
# yum install cryptsetup-luks
```

3. Configure the LUKS partition.

For the LUKS partition, select an unused disk with at least 10 GB of space. You might need to add a new disk.

- a) List all available partitions by using the following commands:

```
fdisk -l
blkid
```

- b) Set up the partition by running the **cryptsetup luksFormat** command.

The sample command shows how to encrypt the `/dev/sdc` partition.

```
cryptsetup -y -v luksFormat /dev/sdc
```

Important: The **cryptsetup luksFormat** command removes all data on the partition that you are encrypting.

- c) Create a logical device-mapper device, mounted to the LUKS-encrypted partition.

Choose a meaningful mapping name. In the following example, `siq` is used as the mapping name for the opened LUKS partition.

```
cryptsetup luksOpen /dev/sdc siq
Enter passphrase for /dev/sdc:
```

This command initializes the volume and sets an initial key or passphrase.

Important: Make sure to remember the passphrase. It is not recoverable.

- d) To check the mapping details, use the following command:

```
ls -l /dev/mapper/siq
```

The output should be similar to the following example:

```
lrwxrwxrwx 1 root root 7 Jun 27 19:21 /dev/mapper/siq -> ../dm-0
```

- e) To view the status of the mapping, use the following command:

```
cryptsetup -v status siq
```

The output should be similar to the following example:

```
/dev/mapper/siq is active.
type: LUKS1
cipher: aes-cbc-essiv:sha256
keysize: 256 bits
device: /dev/sdc
offset: 4096 sectors
size: 52424704 sectors
mode: read/write
Command successful.
```

- f) To check that the device was configured successfully for encryption, run the following command:

```
cryptsetup luksDump /dev/sdc
```

4. Format the mapped device.

- a) Wipe the partition by writing zeros to it:

```
dd if=/dev/zero of=/dev/mapper/siq
```

Tip: Depending on the size of the partition, this command will take some time to complete.

- b) Create a file system on the mapped device, for example, an ext4 file system:

```
# mkfs.ext4 /dev/mapper/siq
```

5. Transfer the IBM StoredIQ application data to the LUKS partition.

- a) Mount the new file system to a temporary location.

In the following example, the mount point `/mnt/siq` is used.

```
mkdir /mnt/siq
mount /dev/mapper/siq /mnt/siq
df -H
cd /mnt/siq
ls -l
```

- b) Copy the content of the `/var/siq` directory to the new file system on the LUKS partition.

```
cp -pr /var/siq/. /mnt/siq/.
```

This command copies all files including any hidden files to the LUKS partition.

6. Mount the LUKS partition to the `/var/siq` directory and restart the appstack service.

- a) Unmount the LUKS partition from the temporary location.

In the following example, the file system is at `/mnt/siq`.

```
umount /mnt/siq
```

b) Rename the `/var/siq` directory, for example, to `/var/siq.orig`:

```
mv /var/siq /var/siq.orig
```

c) Mount the LUKS partition to `/var/siq`.

```
mount /dev/mapper/siq /var/siq
```

Check the content of the `/var` directory:

```
ls -l /var
```

The output should list the `/siq` and the `/siq.orig` directories:

```
drwxr-xr-x 9 root root 4096 Feb 7 18:28 siq
drwxr-xr-x 8 root root 4096 Feb 7 18:28 siq.orig
```

d) Restart the appstack service.

```
service appstack start
```

The application data on the IBM StoredIQ application stack is now encrypted.

At any time, you can change the passphrase for the encrypted disk. To change the passphrase for key slot 0, run the following commands:

```
cryptsetup luksDump /dev/sdc
cryptsetup luksAddKey /dev/sdc
```

Enter the passphrases when prompted.

Enabling or disabling FIPS

At any time, you can enable or disable FIPS as required on the IBM StoredIQ components.

On the gateway or any data server, you must be logged in as root to check or modify the FIPS enablement status. On the AppStack, you must be logged in as `siqadmin` user.

- To check the status of FIPS on the gateway, any data server, or the AppStack, check the `/proc/sys/crypto/fips_enabled` file on each component.

On the gateway or any data server, use `edit` to edit the file directly. On the AppStack, run the **`edit_proc_sys_crypto_fips_enabled`** command. The entry 0 indicates that FIPS is disabled and 1 means it is enabled.

- To enable FIPS:
 - On the gateway or any data server, run this command: `python /usr/local/storediq/bin/fipscfg.pyc --enable`
 - On the AppStack, run the `appstackcfg` utility and select **Enable FIPS mode at boot time**.
- To disable FIPS:
 - On the gateway or any data server, run this command: `python /usr/local/storediq/bin/fipscfg.pyc --disable`
 - On the AppStack, run the `appstackcfg` utility and deselect **Enable FIPS mode at boot time**.

Securing Elasticsearch cluster communication

To better secure the IBM StoredIQ Elasticsearch cluster, you can enable Search Guard on each node in the cluster.

If you want to use Search Guard to secure the Elasticsearch cluster, the primary Elasticsearch node must be configured with specific property settings in the `cluster-setup.properties` file. To edit the file, you must be logged in with the `builder` account. Update the properties file as follows:

- The `STUNNEL_ENABLED` property must be set to `false`.
- For deployments upgraded to IBM StoredIQ 7.6.0.17, add the following entry to the end of the properties file:

```
SEARCHGUARD_ENABLED=true
```

For new deployments (starting with IBM StoredIQ 7.6.0.17), the `SEARCHGUARD_ENABLED` property must be set to `true`.

After editing the properties file, run the cluster setup with the updated properties file. At the prompt `[builder@localhost ~]$`, enter:

```
/siq/bin/cluster-setup.sh cluster-setup.properties
```

Search Guard enables HTTPS connections between the data server and the Elasticsearch cluster and between the nodes of the Elasticsearch cluster. It also ensures that all connections to the Elasticsearch other than those with authorized clients are blocked.

1. Initialize the Search Guard configuration.

- a) Log in to the primary Elasticsearch node as `builder`. Then, switch to the root user:

```
su root
```

- b) Log in to the Docker shell by using this command:

```
docker exec -it node-name bash
```

where `node-name` is the name of the primary Elasticsearch node.

- c) Navigate to the Search Guard configuration directory:

```
cd /usr/share/elasticsearch/config
```

- d) Run the Search Guard **sgadmin** tool with the following options:

```
/usr/share/elasticsearch/plugins/search-guard-5/tools/sgadmin.sh -cd /usr/share/elasticsearch/plugins/search-guard-5/sgconfig -cn cluster_name -cert sgcerts/es-admin.crt.pem -key sgcerts/es-admin.key.pem -cacert sgcerts/chain-ca.pem -nhnv
```

where `cluster_name` is the name of the Elasticsearch cluster.

2. Test the setup by using this command:

```
curl -vk https://localhost:9200/_cluster/health?pretty -E ./sgcerts/es-admin.crt.pem --key ./sgcerts/es-admin.key.pem --cacert ./sgcerts/chain-ca.pem
```

3. To enable secure communication between the data server and the Elasticsearch nodes, copy the data server certificates that are available on the Elasticsearch nodes to each data server.

- a) Log in to an Elasticsearch node as `root`.
- b) Navigate to the directory where the certificates are stored:

```
cd /siq/env/docker/sgcerts
```

- c) Copy the following files to the `/etc/deepfile/dataserver` directory on the data server.

Use a secure copy tool such as scp to do so.

- es-admin.crt.pem
- es-admin.key.pem
- chain-ca.pem

d) Log in to the data server as root and rename the certificates as follows:

es-admin.crt.pem to client.crt.pem
es-admin.key.pem to client.key.pem
chain-ca.pem to cert-chain.pem

e) Restart all services on the data server.

4. Test the connection from the data server to the Elasticsearch cluster:

```
curl -vk https://primary_node_ip:9200/_cluster/health?pretty -E ./client.crt.pem --key ./client.key.pem --cacert ./cert-chain.pem
```

where *primary_node_ip* is the IP address of the primary Elasticsearch node.

Configuring the gateway

Use the SSH tool to configure the gateway.

1. Using an SSH tool, log into the gateway server as root.
2. If you did not enter the appstack IP address when you deployed the OVA, run the `/usr/bin/set-appstack-ip.sh` script and then enter the appstack IP address when prompted.

Note: This script adds an entry into the `pg_hba.conf` file, allowing trusted access to the gateway database from the appstack IP. Without this access, reports cannot run from the appstack. This script must be run whenever the appstack IP address happens to change.

Managing the status of secure gateway communication

The gateway handles the communication between the data servers and the application stack. You can check whether secure gateway communication is used to encrypt the communication and change the enablement status as required.

You must be logged in to the component that you want to manage: on the gateway or any data server as root user, on the application stack as `siqadmin` user.

Secure gateway communication must be enabled or disabled on all components. Follow these steps to see the enablement status of these components, and to disable or enable secure gateway communication as required.

Secure communication via stunnel can impact performance. Therefore, enable this setting only if your enterprise security policy mandates encryption of data in motion. If you do so, the IBM StoredIQ gateway and the data server must be configured accordingly. For more information, see [“Security” on page 7](#). For additional guidance, contact IBM Support.

You can change the enablement status at any time after the installation.

- To check the status:

- On the application stack, enter `service appstack status` on the command prompt.

If **Process appstackStunnel** is listed and running, secure gateway communication is enabled on the application stack. Otherwise, it is disabled.

- On the gateway, enter `service deepfiler status` on the command prompt.

If **Process GatewayStunnel** is in the **running** state, secure gateway communication is enabled. If it is in the **not monitored** state, it is disabled.

- On a data server, enter `service deepfiler status` on the command prompt.
If **Process DataserverStunnel** is listed and running, secure gateway communication is enabled. Otherwise, it is disabled.
- To enable secure gateway communication:
 - a) On the gateway, enter `/usr/local/storediq/bin/util/stunnelcfg` on the command prompt. Then, enter `service deepfiler restart` to restart the services.
 - b) On a data server, enter `/usr/local/storediq/bin/util/stunnelcfg gateway_ip` on the command prompt.
Then, you have to log in to the IBM StoredIQ Data Server user interface to complete this configuration:
 - a. Go to **Administration > Configuration > DA Gateway settings**.
 - b. Enter `127.0.0.1` in the **Host** field.
 - c. Enter the name in the **Node name** field.
 - d. Click **OK**.
 - e. Log off.

Back on the data server, enter `service deepfiler restart` on the command line to restart the services.
 - c) On the application stack, complete these steps:
 - a. Run the `appstackcfg` utility.
 - b. Select **Enable secure gateway communication**.
 - c. Select **Restart appstack services** to restart the services on exiting the configuration utility and then select **Save and exit**.
Alternatively, you can just select **Save and exit** and then restart the services from the command line by using this command: `service appstack restart`
- To disable secure gateway communication:
 - a) On the gateway and on the data server, enter `/usr/local/storediq/bin/util/stunnelcfg disable` on the command prompt. Then, enter `service deepfiler restart` to restart the services.
 - b) On the application stack, complete these steps:
 - a. Run the `appstackcfg` utility.
 - b. Deselect **Enable secure gateway communication**.
 - c. Select **Restart appstack services** to restart the services on exiting the configuration utility and then select **Save and exit**.
Alternatively, you can just select **Save and exit** and then restart the services from the command line by using this command: `service appstack restart`

Configuring the IBM StoredIQ data server

When configuring the IBM StoredIQ data server, you must configure the gateway settings and then restart services.

You must complete this task for every data server.

1. Go to **Administration > Configuration > DA Gateway settings**.
2. In the **Host** field, enter the IP of the gateway.
3. In the **Node name** text box, enter the name of the data server.
4. Restart services in either of the following ways:

- Using an SSH tool, log in as root and then type `service deepfiler restart`.
- Go to the data server dashboard, click the **About Appliance** link, and then click **Restart Services**.

Configuring the application stack to receive reports and notifications

You must configure the application stack accordingly if you want to receive reports and notifications through email.

If the email server was not configured during deployment, follow these steps.

1. Using an SSH tool, log into the application stack as `siqadmin`.
2. Launch the Appstack Configuration utility using this command: `appstackcfg`
3. In the **Domain name** field, enter the fully qualified domain name (FQDN) or IP address of the application stack.
4. In the **StoredIQ Gateway** field, enter the IP address for the IBM StoredIQ gateway server.
5. Set the following SMTP options:

Parameter	Value
Server	The mail server's fully qualified domain name
Port	The SMTP port
Username	The login user name. For the default configuration, leave this field empty. Otherwise, enter the credentials of an authenticated user.
Password	The login password for the specified user. For the default configuration, leave this field empty.

6. Select **Enable TLS** to enable TLS encryption, if it is supported by the mail server.
7. Select **Restart appstack services**.
Alternatively, you can restart the application stack services from the command line after you save the configuration by running the following command: `monit restart uwsgi`
8. Select **Save and exit** and wait for all IBM StoredIQ services to restart.
You can now send email notifications from the IBM StoredIQ application stack or when generating reports.

For more information about SMTP notification, see [“Configuring authenticated users for SMTP notifications”](#) on page 55.

Configuring the application stack to synchronize data with the governance catalog

Configure the application stack to make data from IBM StoredIQ data sources discoverable in IBM Information Server or IBM Cloud Private for Data, to make IBM Information Governance Catalog data classes available for use in IBM StoredIQ, and to keep the information in sync.

Before you can configure and use the synchronization feature, a working deployment of one of these products must be available:

- IBM Information Server. The minimum required version is Version 11.7.0.
- IBM Cloud Private for Data. The minimum required version is Version 1.1.0.2.

If the synchronization was not enabled during deployment, complete these steps to configure the application stack accordingly.

Important: Synchronization requires HTTPS to be enabled on the AppStack. Therefore, it is mandatory that at least a self-signed certificate is installed.

To enable the synchronization with Information Governance Catalog or the IBM Cloud Private for Data catalog at any time after deployment of the AppStack:

1. Using an SSH tool, log in to the application stack as `siqadmin`.
2. Launch the Appstack Configuration utility by using this command: `appstackcfg`
3. Check the value in the **Domain name** field.

The information that you provide here is used to build the base URLs for REST access to the application stack and for links to IBM StoredIQ artifacts that make these artifacts accessible from Information Governance Catalog. Therefore, you should provide the fully qualified domain name or the IP address of the application stack.

4. If you want your data experts to receive email notifications for any changes to the synchronized Information Governance Catalog objects, ensure that the application stack configuration includes the appropriate SMTP settings.

For details, see [“Configuring the application stack to receive reports and notifications”](#) on page 43 or, for Microsoft Exchange specific setup, [“Configuring email notification with an authenticated user”](#) on page 56.

5. Select **Enable synchronization with the governance catalog** and provide the following settings.

If the data catalog to which you want to publish the IBM StoredIQ object resides in an IBM Cloud Private for Data environment, select the **Server runs in IBM Cloud Private for Data** check box.

Provide or accept the values for these fields:

Parameter	Value
Host	The host name or IP address of the Information Server or IBM Cloud Private for Data installation. The specified host is part of the base URLs for REST access to the Information Governance Catalog or IBM Cloud Private for Data instance and for links to catalog artifacts that make these artifacts accessible from IBM StoredIQ. Therefore, you should provide the fully qualified domain name of the Information Server or IBM Cloud Private for Data host. If you specify a server port, this port also becomes part of such base URLs. In addition, the host name is also used to address the Information Server Apache Kafka server, which provides all Information Server events as Kafka messages. Specific Kafka messages are consumed by IBM StoredIQ and trigger the synchronization of objects from the governance catalog to IBM StoredIQ.
Port	The port of the Information Governance Catalog server. This setting is optional. For connections to an Information Server 11.7 environment: To ensure proper communication, you should set the port to the HTTPS port that is defined in Information Server. The default port is 9443. For connections to an Information Server 11.7 FP1 environment without Information Server Enterprise Search installed: To ensure proper communication, you should set the port to the

Parameter	Value
	HTTPS port that is defined in Information Server. The default port is 9446.. For connections to an Information Server 11.7 FP1 environment with Information Server Enterprise Search installed: Do not specify a port. For connections to IBM Cloud Private for Data, you can specify the IBM Cloud Private for Data port.
Kafka port	The port of the Information Server Kafka server. The port setting can be overridden. For connections to an Information Server 11.7 environment: the Kafka port defined in Information Server. The default port is 59092. For connections to an Information Server 11.7 FP1 environment without Information Server Enterprise Search installed: the Kafka port defined in Information Server. The default port is 59092. For connections to an Information Server 11.7 FP1 environment with Information Server Enterprise Search installed: the Kafka port defined in Information Server. The default port is 9092. For publishing to IBM Cloud Private for Data, this setting is not used. However, an error will be logged, which you can ignore.
Username	The user name for authenticating to Information Server when publishing IBM StoredIQ objects to the governance catalog. This user must be defined in Information Server with the following security roles: - Suite User - Common Metadata Administrator - Information Governance Catalog Information Asset Administrator In IBM Cloud Private for Data, this user must be defined with the Data Stewart role.
Password	The password of the user set with Username .
Sync frequency (minutes)	Data is periodically propagated to the governance catalog at the specified interval. The value must be a positive number of minutes. The default value is 15 minutes.
StoredIQ instance name	The name identifying the IBM StoredIQ instance for which data is synchronized. This name can be freely chosen, but must be unique within the Information Governance Catalog instance.

6. Select **Restart appstack services**.

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

7. Select **Save and exit** and wait for all IBM StoredIQ services to restart.

On the first synchronization run, Information Governance Catalog is initially populated with the IBM StoredIQ artifacts that you selected for publishing to the catalog.

Installing IBM StoredIQ Desktop Data Collector

IBM StoredIQ Desktop Data Collector deploys from the IBM StoredIQ Data Server Administrator interface. You cannot install it until you finish configuring the data server and bring it up running.

For the protocol TCP, the required port ranges are as follows:

- 21000-21004
- 21100-21101
- 21110-21130
- 21200-21204

How to open ports for desktop client access to the data server on OVA deployed systems is described in [“Open ports for desktop client access to the data server” on page 32](#).

1. Enter the https:// IP address of the data server in the browser.
2. Enter admin in the email field and password admin on the login page.
3. Under the **Administration** tab, click the **Configuration** tab.
4. Click the **Desktop Settings** link under the **Application** column on the **System Configuration** page.
5. Click the **Download the desktop client installer** link.
This is being installed on a Windows workstation.
6. Click **Run** when the installer window prompts.
7. Complete the installation steps in the InstallShield Wizard.
8. After the installation is complete, IBM StoredIQ Desktop Data Collector appears as volumes on the primary data source page of the data server and in the IBM StoredIQ Administrator interface on the application stack.

Backing up the IBM StoredIQ image

Backing up the IBM StoredIQ image is a good method for disaster recovery. It is also a best practice before you start any upgrades on your image. If you need to back up the IBM StoredIQ image, you must complete the following steps.

An active IBM StoredIQ image must not be backed up by using VMWare VCenter or other product backup utilities. If you do so, the data servers might hang and become unresponsive. Running a backup snapshot on an active IBM StoredIQ image might result in transaction integrity issues.

If a backup snapshot of IBM StoredIQ image is needed, follow these steps:

1. Enter this command to stop all IBM StoredIQ services on all data servers and gateway: `service deepfiler stop`

Note: Wait 10 minutes after a harvest before you use this command to stop services.

2. Enter this command to stop the postgresql database service on all data servers and gateway: `service postgresql stop`
3. Enter this command to stop the IBM StoredIQ services on the application stack: `service appstack stop`
4. Contact the VMWare VCenter administrator to have the IBM StoredIQ image manually backed up. Confirm the work completion before you proceed to the next step.
5. Enter this command to restart the IBM StoredIQ services on all data servers and gateway: `service deepfiler restart`
6. Enter this command to start the IBM StoredIQ services on the application stack: `service appstack start`

To prepare for disaster recovery, another method is to back up the system configuration of the IBM StoredIQ data server to an IBM StoredIQ gateway server.

Upgrades for IBM StoredIQ

You can obtain the upgrade ISO from IBM Fix Central. Upgrades consist of upgrading the gateway and data server first and then the application stack. These instructions apply to upgrades from the preceding release version to the current version. Direct upgrades from earlier versions are not supported.

Important:

- Back up all of your IBM StoredIQ VMware images (gateway, data servers, application stack, and all nodes of the Elasticsearch cluster, if deployed) before you start the upgrade in case the upgrade fails. You can use one of the following methods:
 - Create a snapshot of the current state of the image.
 - Create a clone of the current state of the image.
- In deployments with a DataServer - Distributed, you must upgrade the Elasticsearch cluster before you upgrade any data server. For more information, see [“Upgrading the Elasticsearch cluster”](#) on page 51.

Go to IBM Fix Central. In the **Product selector** field, enter StoredIQ and then select the latest fix pack as **Installed Version**. The upgrade ISO is part of the fix pack.

Considerations when not upgrading from the preceding version

IBM StoredIQ update packages are not cumulative. Therefore, you must upgrade to the preceding release version before you can upgrade to the current version.

Direct upgrades are supported from IBM StoredIQ 7.5.1.0 to 7.6.0.1. The upgrade path for any release version starting with version 7.6.0.1 requires you to install all 7.6.0.x upgrade packages one by one in ascending order.

For new deployments, directly install the most current fix pack without installing the base version 7.6.0. Each fix pack does not only contain an upgrade package but also an OVA for a fresh installation.

Gateway upgrades from 7.6.0.2 to 7.6.0.4

When updating the gateway from 7.6.0.2 to 7.6.0.4, run this command: `sed -i '/^icc/d' /etc/deepfile/package_excludes` and then proceed with the normal upgrade process.

Application stack upgrades from 7.6.0.2 to 7.6.0.4

For application stack upgrades from 7.6.0.2 to 7.6.0.4, complete the following steps before the actual upgrade:

1. Update the gateway and open port 80.
2. Log in as root using PuTTY and run the following commands:
 - a. `wget http://<gatewayIP>/products/appstack/centos65/7.6.0.4+32/siq-appstack-1.1.15+81-1.x86_64.rpm`
 - b. `rpm -Uvh siq-appstack-1.1.15+81-1.x86_64.rpm`
3. Proceed with the normal application stack upgrade process.

Upgrading the gateway and data server from 7.6.0.3 to 7.6.0.4

For DVD and ISO upgrades from 7.6.0.3 to 7.6.0.4, the following steps must be done before the upgrade of the gateway and data servers.

1. Copy the 7.6.0.4 upgrade ISO into the tmp directory on the gateway and data servers.
2. Log in as root using PuTTY.

3. Create mount point by running this command:

```
mkdir /mnt/cdrom
```

4. On the data server and gateway that are being updated, mount the update ISO:

```
mount -o loop /tmp/7_6_0_4-STOREDIQ-38-upgrade.iso /mnt/cdrom
```

5. Upgrade the updater package:

```
rpm -Uvh /mnt/cdrom/storediq-appliance-updater*.rpm
```

6. Unmount the ISO:

```
umount /mnt/cdrom
```

7. When upgrading the gateway, run this command:

```
sed -i '/^icc/d' /etc/deepfile/package_excludes
```

Do **not** run this command on the data servers

8. Proceed with the gateway and data server normal upgrade process.

Upgrading to 7.6.0.15

The 7.6.0.14-IBMStoredIQ_IF001 was uploaded to Fix Central in January, 2018. Before you upgrade to IBM StoredIQ Version 7.6.0.15, deploy the Spectre (CVE-2017-5715, CVE-2017-5753) and Meltdown (CVE-2017-5754) fix pack to the 7.6.0.14 system and then upgrade.

Also check the information about enabling disk encryption on an Elasticsearch cluster.

Upgrading to 7.6.0.16

The following migration steps are required:

Change existing Linux account passwords

Starting with IBM StoredIQ Version 7.6.0.16, hashing algorithm for Linux account passwords was changed to SHA512 with 999,999 rounds, which is now the default for new installations. For upgraded deployments, the passwords for all Linux user accounts that can log in to the system remotely or that can execute tasks with elevated privileges must be updated manually on each OVA after the upgrade. Do this to have upgraded deployments match the security standards for new installations where the passwords are hashed using SHA512 with 999,999 rounds by default.

For each user ID, change the current password and specify the hashing algorithm that you want to use as shown in the examples:

```
/usr/sbin/authconfig --passalgo=sha512 --update passwd userID
```

This command changes the password for the given user ID using SHA512 as the hashing algorithm.

```
echo userID:mynewpassword | chpasswd --crypt-method SHA512 --sha-rounds 999999
```

This command changes the password for the given user ID using SHA512 as the hashing algorithm with 999,999 rounds of hashing.

Set the password for the synchronization anew

If synchronization with the governance catalog is enabled, you must update the configuration of the application stack after the product upgrade is complete by setting the password for the governance catalog user once again. If you do not update the password, synchronization with the governance catalog will no longer be possible. If the synchronization is not enabled, you can skip this update.

1. Launch the Appstack Configuration utility by using the **/siq/bin/appstackcfg** command.

2. In the section where the settings for the synchronization are configured, navigate to the **Password** field and specify the password for the user whose credentials are used for authenticating to Information Server.
3. Select **Restart appstack services**.
4. Select **Save and exit** and wait for all IBM StoredIQ services to restart.

Enabling disk encryption on an Elasticsearch cluster upgraded to 7.6.0.15

For an Elasticsearch cluster that you upgraded to IBM StoredIQ 7.6.0.15, follow these instructions to encrypt the volume on which the Elasticsearch indexes are stored. IBM StoredIQ uses Linux Unified Key Setup (LUKS) for disk encryption.

Make sure that the cluster works correctly and all nodes are available by running these commands:

```
curl -X GET 'http://your first node IP:9200/_cluster/state?pretty'
curl -X GET 'http://your first node IP:9200/_nodes?pretty'
```

New installations of IBM StoredIQ Version 7.6.0.15 and later by default encrypt the disk volume on which the Elasticsearch indexes are stored. However, if you installed a previous version and upgraded to Version 7.6.0.15, the existing volumes were not encrypted automatically during the upgrade process. For security reasons, you should encrypt the volume. However, you cannot decrypt the volume again later.

1. For each node, determine whether the Elasticsearch data on the disk is less or greater than the available storage on the root partition.
The root partition can be used for backing up the existing data.
 - If the data is less than the available storage, proceed with step “2” on page 50.
 - If the data on the disk is greater than the available storage, add a disk of the required size to the VM and mount it at /mnt/backup.
2. Restart the VM on each node to ensure that no outside processes or connections are using the data partitions. If you added a disk, ensure that it is still mounted at /mnt/backup.
3. Log in to the primary master node from PuTTY by using the **builder** account and password that were configured during the installation of the cluster. Edit the `cluster-setup.properties` file by appending `LUKS=true` on a new line at the end of the file.
4. On the primary master node, run the cluster setup script. At the prompt `[builder@localhost ~]$`, enter:

```
/siq/bin/cluster-setup.sh cluster-setup.properties
```

Note: If these steps were not followed correctly, the process might fail.

Enter the password to continue for each node.

5. Indexes are now encrypted but the key is accessible on non-encrypted disks.

If you want to secure the key with a LUKS password, complete these steps for each node in the Elasticsearch cluster:

- a) Log in from the vCenter Console by using the root account and password that were configured during the installation of the cluster.
- b) Add a key by using the **cryptsetup** tool. Issue the following command:

```
[root@localhost ~]$ sudo cryptsetup -d /root/siq-elasticsearch-luks.key luksAddKey /dev/sdb1
```

- c) Enter the new passphrase for the key slot and verify it.

Important: You must remember that passphrase or keep it in a safe place. If you store the passphrase, make sure to keep it in a place away from the information it protects.

- d) Remove the LUKS key after you verify the password by using `shred` utility:

```
[root@localhost ~]$ sudo shred -zvu -n 5 /root/siq-elasticsearch-luks.key
```

Upgrading the Elasticsearch cluster

If you have a DataServer - Distributed, you must upgrade the Elasticsearch cluster first before you upgrade any data server. If you have only the DataServer - Classic, skip this procedure and proceed to the next topic.

Follow these steps to upgrade the Elasticsearch cluster.

1. Download the release upgrade ISO to the primary master node of the Elasticsearch cluster where the installation was initially started.
Go to IBM Fix Central. In the **Product selector** field, enter StoredIQ and then select the latest fix pack as **Installed Version**. The upgrade ISO is part of the fix pack.
2. Log in to the primary master node from vCenter Console for that node by using the root account and password that were configured during the installation of the cluster.
3. In the Console, run these commands to prepare the upgrade:

```
[root@hostname ~]# cd /home/builder
[root@hostname builder]# mount 7.6.0.17-STOREDIO-Upgrade-23.iso /mnt
[root@hostname builder]# cd /mnt
[root@hostname mnt]# /usr/bin/python -m SimpleHTTPServer 80 >/dev/null 2>&1 &
[root@hostname mnt]# export killpid=$! && cd
```

4. Log on to the primary master node from PuTTY by using the **builder** account and password that were configured during the installation of the cluster. Run the following commands to upgrade the cluster.

```
[builder@hostname ~]$ bootstrap product.set elasticsearch repository http://IP of primary
master node/products/elasticsearch
```

Press **Enter** to continue for each node.

```
[builder@hostname ~]$ bootstrap product.upgrade elasticsearch 7.6.0.17+2
```

Press **Enter** to continue for each node.

5. Go back to the **Console** window from step “3” on page 51 and run these commands to finish the upgrade:

```
[root@hostname ~]$ kill $killpid
[root@hostname ~]$ umount /mnt
```

Updating the gateway and data servers

Download the upgrade ISO and then update the gateway and data servers.

You must upgrade the gateway before upgrading any data servers.

1. Download the release upgrade ISO 7.6.0.17-STOREDIO-upgrade-23.iso from IBM Fix Central and store it in the /tmp folders on the gateway and on each data server that you want to upgrade.
Download the ISO to your local machine by using one of the options provided on Fix Central. Then, copy it to the gateway and data servers by using a copy tool, for example, PuTTY Secure Copy client on Windows.

Complete the following steps for the gateway first and then for each data server that you want to upgrade.

2. Log in to the virtual machine as `util util`.
 - a) Use the arrow and Tab keys to select **Manage Software Repositories**, and then select **OK**.
 - b) Use the arrow and Tab keys to select **DVD Update Repository** and then select **Edit**.
 - c) Use the arrow keys, the space bar, and the Tab key within **Edit Update Repository** to select **ISO** within **Type**, and then select **Next**.
 - d) Enter the location of the upgrade ISO and the name of the ISO such as /tmp/7.6.0.17-STOREDIO-Upgrade-23.iso, and then select **Finish**.

- e) In the **Manage Update Repositories** screen, select **Done**.
3. When prompted, save the repository configuration changes by selecting **Yes**.
4. In the **Appliance Manager Screen**, select **Update This Node** and select **OK**.
5. In **Choose Update Repositories**, select **DVD Update Repository** and then select **OK**.
6. After the new package updates are listed, select **OK**.
The **Back up database?** message appears. If you took a VM snapshot, select **No**. You are notified that a new update package was installed.
7. Select **Restart**. At the command prompt, enter `su util`.
8. Select **Update This Node** again and then select **OK**.
The **Choose Update Repositories** window opens.
9. Select **DVD** and then select **OK**.
The available packages are listed again.
10. After the new package updates are listed, select **OK**.
The **Back up database?** message appears. If you took a VM snapshot, select **No**. You are notified that a new update package was installed.

Several informational messages appear, indicating that the repository is being created, what new package updates are available, that update transactions are being created, and that the system is restarted.
11. To complete the upgrade, you must apply some additional security update:
 - a) Log out from the system (as `util` user) and log back in as `root` user.
 - b) Run the following command:

```
/usr/local/storediq/bin/tools/postupgrade.sh
```

When the script completes, the message `Update complete` is displayed.

Updating the application stack

After you update the gateway and data servers, you can proceed to update the application stack.

Verify that the gateway and data servers are already updated.

Important: The upgrade to version 7.6.0.17 entails an upgrade of PostgreSQL. Therefore, it is strongly recommended that you back up your AppStack data before starting the upgrade. Otherwise, you will not be able to recover your data if the upgrade fails. For more information, see the topic about backing up the IBM StoredIQ image.

- Keep the services running in the AppStack before you start an upgrade.
 - Check the `/var/bootstrap/appstack.log` file whenever an upgrade is running.
1. Use an SSH tool to log in to the gateway and complete these steps.
 - a) Open port 80 by running this command:

```
iptables -A PROD-web -p tcp -m state --state NEW -m tcp --dport 80 -j ACCEPT
```

- b) Restart the services on the gateway using the **service deepfiler restart** command.
 - c) In a browser, you can go to `http://gatewayIP/products` to check whether the product package is available.
2. Use an SSH tool to log in to the application stack as `root` user and run these commands.
 - a) Run `bootstrap product.set appstack repository "pkgsource=\"http://gatewayIP/products\""`
 - b) Run `bootstrap product.upgrade appstack HEAD`

c) Required: Monitor the upgrade process by tailing the `appstack.log` file:

```
tail -f /var/bootstrap/appstack.log
```

The process stops if the available space on the Postgres data volume is not sufficient for the upgrade. In this case, fix the issue as indicated and restart the upgrade.

The process stops in any case at some point to give you the opportunity to back up your AppStack data before upgrading in case you haven't already done so.

Tip: Copy the command shown in the message.

Exit the log file. Back up your data if required. To proceed with the upgrade, paste the command that you copied earlier to the command line or enter the following command manually:

```
nohup /etc/bootstrap/upgrade-appstack.sh /var/bootstrap/products/appstack/repository HEAD  
production Continue &>/dev/null &
```

Tail the log file again to monitor the remaining upgrade process.

When the upgrade is complete, you are notified that the production mode is activated.

- d) Examine the log file for any traceback or any keywords such as `ProgrammingError` or `OperationalError`. Contact IBM Support immediately if you come across such issues.
3. Apply additional security updates by running the following command:

```
/siq/bin/postupgrade.sh
```

When the script completes, the message `Update complete` is displayed.

4. Check whether the vault service is running by using this command:

```
service vault status
```

If the vault service is not running after the upgrade (which is indicated by the message `Vault server is not running`), start it manually by using this command:

```
service vault start
```

Tip: After the upgrade, you can run these commands as `root` or as `siqadmin` user.

5. After the update of the application stack is complete, log in to the gateway as `root` by using an SSH tool and complete these steps.
- a) Close port 80 by running this command:

```
iptables -D PROD-web -p tcp -m state --state NEW -m tcp --dport 80 -j ACCEPT
```

- b) Restart services on the gateway by running the **`service deepfiler restart`** command.

Installing the Livelink APIs

The Livelink API requires the installation of the `lapi.jar`, which is provided by its vendor, OpenText, under license to their customers. IBM cannot redistribute this file. At installation, you must copy this file to each data server in the `/usr/local/IBM/ICI/vendor` directory. It is likely to be found on the Livelink server in the `C:\OPENTEXT\application\WEB-INF\lib` directory. However, Livelink installations can vary; you must consult the customer's documentation to locate the `lapi.jar` file.

1. Locate the `lapi.jar` file by going to `C:\OPENTEXT\application\WEB-INF\lib`.

Note: Installations vary by customer. If the `lapi.jar` file is not in the `C:\OPENTEXT\application\WEB-INF\lib` directory, consult the customer's documentation.

2. Copy the `lapi.jar` file to each data server in the `/usr/local/IBM/ICI/vendor` directory.

Configuring authenticated users for SMTP notifications

The IBM StoredIQ application-stack configuration requires that the username and password be blank for the default configuration in order to send email notifications from the application stack. In order to use an authenticated user with the IBM StoredIQ application-stack configuration, you must provide authorization in both Microsoft Exchange and Active Directory. While it is possible to use the default SMTP Receive Connector, it is recommended that a custom SMTP Receive Connector be created and configured to work with the IBM StoredIQ application stack. If a custom SMTP Receive Connector is created and defined with the specific IP address of the IBM StoredIQ application stack, network traffic from the application stack will utilize that connector.

Creating an SMTP Receive Connector

Using an SMTP Receive Connector, email messages are transmitted to the Exchange Server for processing.

1. In the Exchange Management Console navigate, to the **Server Configuration > Hub Transport** section. Select **New Receive Connector** from the right-side menu.
The New Receive Connector dialog box appears.
2. In the **Name** text box, enter a name for the new SMTP Receive Connector, and then click **Next**.
3. Configure your local network settings. By default, you can utilize port 25, which is listed by default port. For better security, you can utilize a different port, but you will need to verify that port is open and available.
4. Specify the **Fully Qualified Domain Name (FQDN)**. For example, you might enter `Exchange2010.ibm1ab.local`, and then click **Next**.
5. On the Remote Network settings page, click **Add**, and enter the valid IP address for the IBM StoredIQ application stack. Click **OK**.
The IBM StoredIQ application stack appears within the IP addresses.
6. Select the IP address range of **0.0.0.0-255.255.255.255**, and then click the red **X** to delete it. Click **Next**.
7. Review the configuration details, and then click **New**.
If the new receive connector is valid, a green check mark and successful completion message will appear.
8. Click **Finish**. The new connector appears within the list of available receive connectors.
9. Restart the Microsoft Exchange Transport service. Go to **Start > Administrative Tools > Services**, and select **Microsoft Exchange Transport** service. Right-click, and select **Restart**.

Authenticating users

There are two methods for authenticating users for SMTP notification: authentication against existing domain users and a single, specific domain user. Both methods are provided herein.

1. Log into your Exchange Server.
2. Launch the Active Directory Services Interfaces (ADSI) editor. Click **Start > Administrative Tools > ADSI Edit**.
3. Within ADSI, navigate to **CN=Configuration > CN=Services > CN=First Organization > CN=Administrative Groups > CN=Exchange Administrative Group > CN=Servers > CN=Exchange Server > CN=Protocols > CN=SMTP Receive Connectors**, and select **CN=SMTP Receive Connectors**.
The newly created SMTP Receive Connector is displayed as one of the available options.

4. Right-click the newly created SMTP Receive Connector and select **Properties**. The **Properties** window opens.
5. Click **Security**, click **Add**, and the **Select Users, Computers, Service Accounts, or Groups** dialog box appears.
6. In the **Enter the object names to select** text box, enter either of the following:
 - Enter **Authenticated Users**, and then click **Check Names**. Because **Authenticated Users** is a valid entry, a line appears underneath it.
 - Enter a valid, existing user ID and then click **Check Names**. In order to enter a specific user, that individual must already exist within Active Directory. Because that user exists within Active Directory, it enters the fully qualified user name and domain. For example, if a user ID of **storediqsupport** was created in Active Directory, you would enter the user ID **storediqsupport**.
7. Authenticate the user, either existing domain users or a single, specific domain user.
 - If you are authenticating existing domain users, in the **Group or user names** area, select **Authenticated Users**. In the **Permissions for Authenticated Users**, select the **Allow** check box for **Accept Any Sender** and **Accept Authoritative Domain Sender**.
 - If you are authenticating a single, specific domain user, in the **Group or user names** area, select the specific user ID. In the **Permissions for Authenticated Users**, select the **Allow** check box for **Accept Any Sender** and **Accept Authoritative Domain Sender**.
8. Click **OK**.

Configuring email notification with an authenticated user

Configure the IBM StoredIQ application stack to use an authenticated user for email notifications.

1. Using an SSH tool, log into the application stack as **siqadmin**.
2. Launch the Appstack Configuration utility using this command: **appstackcfg**
3. In the **Domain name** field, enter the fully qualified domain name (FQDN) or IP address of the application stack.
4. In the **StoredIQ Gateway** field, enter the IP address for the IBM StoredIQ gateway server.
5. In the **Server** field, provide the fully qualified hostname or IP address of the Exchange server.
6. In the **Port** field, leave the default entry of 25.
7. In the **Username** field, provide the user name of the user you with which to authenticate into the Exchange server.
 - If you authorized any Authenticated User, you can use any user name and password as long as that individual is valid member of the domain.
 - If you used a specific user, you must use the user name of the single user for which you granted permissions. This must be a fully qualified user name.
8. In the **Password** field, enter the password for the user selected in the previous step
9. Select **Enable TLS** to enable transport-layer security.
10. Select **Restart appstack services**. This is an alternative to restarting the application stack services from the command line.
11. Select **Save and exit** and wait for all IBM StoredIQ services to restart.

You can now send email notifications from the IBM StoredIQ application stack or when generating reports.

Notices

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

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

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

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

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

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: 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 states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

IBM Corporation
J46A/G4
555 Bailey Avenue

San Jose, CA 95141-1003
U.S.A.

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.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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.

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2004, 2010. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

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.

[Depending upon the configurations deployed, this Software Offering may use [session] [and] [persistent] cookies that collect each user's [name,] [user name,] [password,] [profile name,] [position,] or [other personally identifiable information] for purposes of [session management,] [authentication,] [enhanced user usability,] [single sign-on configuration] [or other usage tracking or functional purposes.] These

cookies [cannot be disabled] [can be disabled, but disabling them will also [likely] eliminate the functionality they enable].

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

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

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.

Trademarks

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

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

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

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

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

Index

A

- application data
 - disk encryption [37](#)
- application stack
 - update [52](#)
- appstack [5](#), [41](#)
- AppStack [1](#)

B

- back up [47](#)

C

- certificates
 - replacing [34](#)
 - secure gateway communication [34](#)
 - stunnel [34](#)
- components of the IBM StoredIQ solution [1](#)
- configuration
 - IBM StoredIQ [30](#)
- configure
 - application stack [21](#), [43](#)
 - gateway [41](#)
 - IBM StoredIQ data server [42](#)
 - synchronization [43](#)

D

- data at rest
 - encryption [33](#), [37](#)
- data server
 - installation of [20](#)
- data sources [5](#)
- default open ports [30](#)
- deploy
 - OVA [9](#)
 - OVF [9](#)
 - virtual appliance [9](#)
- deployment
 - deployment planning [3](#)
- Desktop Agent [1](#)
- disk encryption
 - application data [37](#)
 - AppStack [37](#)
 - Elasticsearch index [33](#)
- download
 - gateway and data server [51](#)

E

- Elasticsearch
 - deploy three-node Elasticsearch cluster [13](#)
 - single node cluster [15](#)
 - three node cluster [13](#)

- Elasticsearch cluster [50](#), [51](#)
- Elasticsearch index
 - disk encryption [33](#)
- Elasticsearch nodes [40](#)
- email notification [55](#)
- encryption
 - data at rest [33](#), [37](#)
- environment sizing [5](#)
- ESX server [9](#)

F

- FIPS [39](#)

G

- gateway
 - installation of [18](#)
- gateway server [7](#)

H

- hardware requirements [3](#)
- harvest frequency [5](#)

I

- IBM StoredIQ
 - components [1](#)
 - IBM StoredIQ Administrator [1](#)
 - IBM StoredIQ Data Server [1](#)
 - IBM StoredIQ Data Workbench [1](#)
 - IBM StoredIQ Policy Manager [1](#)
- IBM StoredIQ data server [42](#)
- IBM StoredIQ Desktop Data Collector [46](#)
- IBM StoredIQ image [47](#)
- install
 - application stack [21](#)
 - IBM StoredIQ Desktop Data Collector [46](#)

K

- key management [7](#)

L

- legal
 - notices [57](#)
 - trademarks [59](#)
- lifecycle
 - certificates [32](#)
 - keys [32](#)
- Livelihood APIs [54](#)
- logging in
 - application stack UI [29](#)
- LUKS [7](#), [33](#), [37](#), [50](#)

M

metadata [5](#)
metric [5](#)
Microsoft Hyper-V [11](#)

N

network [5](#)
non-direct upgrades [48](#)
notices
 legal [57](#)

O

open ports [30](#)
Open Virtual Appliance (OVA) [3](#)
Open Virtualization Format (OVF) [9](#)
OVA [3](#), [9](#)
OVF [9](#)

P

port [5](#), [7](#)
port ranges [5](#)
ports [30](#)
Prerequisites for installing IBM StoredIQ on Microsoft Hyper-V [11](#)

R

receive reports and notifications [43](#)

S

security [7](#), [33](#), [37](#), [40](#)
SMTP [55](#)
SMTP notification [55](#)
SSH tool [42](#)
stack-provisioning prerequisites [7](#)
stunnel [7](#)
synchronization configuration [43](#)

T

TCP
 port ranges [5](#)
TLS [7](#)
trademarks [59](#)

U

update the Elasticsearch cluster [51](#)
upgrade [50](#)
upgrade ISO [48](#), [51](#)
upgrades [48](#)
user authentication [55](#)

V

vCenter server [9](#)
virtual appliances [3](#), [9](#)

