

IBM StoredIQ Platform
Version 7.6.0.13

Deployment and Configuration Guide



IBM StoredIQ Platform
Version 7.6.0.13

Deployment and Configuration Guide



Note

Before using this information and the product it supports, read the information in Notices.

This edition applies to Version 7.6.0.13 of product number 5724M86 and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2001, 2017.**

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

Contents

About this publication	v	Configuring the application stack to receive reports and notifications.	25
IBM StoredIQ Platform product library	v	Installing IBM StoredIQ Desktop Data Collector	25
Contacting IBM StoredIQ customer support	v		
Contacting IBM	v	Chapter 8. Backing up the IBM StoredIQ image	27
How to send your comments	vi		
Chapter 1. IBM StoredIQ components	1	Appendix A. Application data migration 29	
Components of the IBM StoredIQ solution	1	Migrating the data server data	29
IBM StoredIQ products	1	Migrating the gateway data	30
		Migrating the application-stack data	30
Chapter 2. Deployment preparation	3		
Deployment planning	3	Appendix B. Upgrades for IBM StoredIQ Platform	33
Open Virtual Appliance (OVA) configuration requirements	3	Gateway upgrades from 7.6.0.2 to 7.6.0.4	33
Network and port requirements	4	Gateway and data server upgrades from 7.6.0.3 to 7.6.0.4	33
Environment sizing guidelines	5	Downloading the upgrade ISO	33
Stack-provisioning prerequisites	6	Updating the application stack	34
		Best practices for updating the application stack	35
Chapter 3. Deployment of IBM StoredIQ Platform.	9		
Open Virtualization Format (OVF)	9	Appendix C. Customizable reports.	37
Deploying the virtual appliances.	9	Installing the BIRT Report Designer	38
		Installing the JDBC driver	38
Chapter 4. Installation of IBM StoredIQ Platform on Microsoft Hyper-V	11	Adding data sources	38
Prerequisites	11	Setting report parameters.	39
Extracting the OVA	11	Adding data sets	40
Converting VMWare .vmdk to Microsoft Hyper-V .vhd	11	Testing report designs	40
Building the Virtual Machine	11	Uploading the report design.	41
Adding more disks	12	Deploying fonts	41
		Changing the reportuser password	42
Chapter 5. Installation of IBM StoredIQ Platform	13		
Deploying a three-node Elasticsearch cluster	13	Appendix D. Installing the Livelink APIs 43	
Installing the gateway	14		
Installing the data server	15	Appendix E. Configuring authenticated users for SMTP notifications	45
Installing and configuring the application stack	16	Creating an SMTP Receive Connector.	45
		Authenticating users	45
Chapter 6. Logging in to the application stack	21	Configuring the application stack to send notification emails	46
Chapter 7. Configuration of IBM StoredIQ Platform	23	Notices	49
Default open ports	23	Privacy policy considerations	51
Configuring the gateway	24	Terms and conditions for product documentation.	52
Configuring the IBM StoredIQ Platform data server	24	Trademarks	52
		Index	55

About this publication

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

IBM StoredIQ Platform product library

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

- *IBM StoredIQ Platform Deployment and Configuration Guide*, SC27-6386
- *IBM StoredIQ Platform Overview Guide*, GC27-6398
- *IBM StoredIQ Platform Data Server Administration Guide*, SC27-5692
- *IBM StoredIQ Administrator Administration Guide*, SC27-5688
- *IBM StoredIQ Data Workbench User Guide*, SC27-5691
- *IBM StoredIQ eDiscovery User Guide*, SC27-5693
- *IBM StoredIQ Policy Manager User Guide*, SC27-5694

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:

Support and assistance

- 1-866-227-2068

To e-mail IBM StoredIQ customer support, use this email address:

- storediqsupport@us.ibm.com

For information about how to contact IBM, see the Contact IBM web site at <http://www.ibm.com/contact/us/>

IBM Knowledge Center

The IBM StoredIQ publications can be found from IBM Knowledge Center.

PDF publications

The IBM Publication Center site offers customized search functions to help you find all the IBM publications you need.

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/>.

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

Send your comments by using the online reader comment form at https://www14.software.ibm.com/webapp/iwm/web/signup.do?lang=en_US&source=swg-rcf.

Chapter 1. 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.

Components of the IBM StoredIQ solution

The three components of the IBM StoredIQ solution are the application stack (often referred to as the appstack), the gateway, and data servers.

Application stack

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

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 Platform Data Server helps you to 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, and file owners.

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 an administrator user interface, administrators can deploy the IBM StoredIQ Desktop Data Collector and index desktops from the data server.

IBM StoredIQ products

The IBM StoredIQ products include IBM StoredIQ Administrator, IBM StoredIQ Data Workbench, IBM StoredIQ eDiscovery, 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 Platform interface and the applications and facilitates the transfer and communication of information. IBM StoredIQ Administrator understands and manages IBM StoredIQ Platform 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.

IBM StoredIQ eDiscovery

IBM StoredIQ eDiscovery is integral to the eDiscovery process. It helps discover, identify, refine, preserve, and collect data that is relevant to legal matters. With IBM StoredIQ eDiscovery, you can conduct discovery work more efficiently and capture the appropriate data.

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 Platform 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.

Chapter 2. Deployment preparation

To prevent any deployment problems, conduct the following planning activities that are appropriate for your environment.

Deployment planning

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

Open Virtual Appliance (OVA) configuration requirements

IBM StoredIQ Platform is deployed as virtual appliances and currently only supported in VMware ESX(i) 5.x or greater environments. You must have a virtual infrastructure that meets the IBM StoredIQ Platform hardware requirements.

DataServer - Distributed size

- vCPU: 4
- Memory: 16 GB
- Storage:
 - Primary disk (vmdisk1): 100 GB
 - Data disk (vmdisk2): 1 TB

Note: The required memory depends on the index size that is handled by the data node. If the index size is 12 GB, 16 GB memory is needed. If the index size is larger, higher memory is required.

Gateway server size

- vCPU: 2
- Memory: 8 GB
- Storage:
 - Primary disk (vmdisk1): 100 GB
 - Data disk (vmdisk2): 75 GB
 - Swap disk (vmdisk3): 40 GB

Data server size

- vCPU: 4

Note: 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

Note: While the minimum value works under light-load condition, as the load increases, the data server quickly starts consuming 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

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

Note: 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

Note: 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. The default data disk size is therefore targeted for managing 10 TB of source information.

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

Note: When under load, the data server can use a lot of 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.

Application stack size

- vCPU: 1
- Memory: 4 GB
- Storage:
 - Primary disk (vmdisk1): 21 GB
 - Data disk (vmdisk2): 10 GB

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:

- One data server per 30 TBs of file shares, which can vary based on the number of volumes, objects per volume, and object types.
- 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.x or greater.
- VMware virtual machine version 8.0 or later.
- 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.

Chapter 3. Deployment of IBM StoredIQ Platform

IBM StoredIQ Platform 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.

Before you begin

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

About 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.

Procedure

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.

Chapter 4. Installation of IBM StoredIQ Platform on Microsoft Hyper-V

IBM StoredIQ Platform can be installed on Microsoft Hyper-V.

Prerequisites

To install IBM StoredIQ Platform 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 `vmdk.gz` file.
 4. `qemu-img` (Linux) or Microsoft Virtual Machine Converter 3.0 (Windows) must be installed to convert the files from VMWare `.vmdk` files to Hyper-V `.vhdx` files.
-

Extracting the OVA

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

About this task

Follow these steps to extract the OVA.

Procedure

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 `vmdk.gz` by using 7-zip for Windows or gunzip for Linux.
 4. Delete everything except the `vmdk` file when the OVA extraction is complete.
-

Converting VMWare .vmdk to Microsoft Hyper-V .vhdx

For instructions about converting the `.vmdk` file to `.vhdx` 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.

Procedure

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.

Procedure

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.

Chapter 5. Installation of IBM StoredIQ Platform

After the OVF deployment, proceed to install IBM StoredIQ Platform.

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 25

Deploying a three-node Elasticsearch cluster

If you are deploying a DataServer - Distributed type of the data server, you must first deploy the Elasticsearch cluster. Follow these steps to use the Elasticsearch OVA in deploying a three-node Elasticsearch cluster. If you are deploying a DataServer - Classic type of the data server, proceed to the next topic “Installing the gateway” on page 14.

About this task

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

Procedure

1. Download the Elasticsearch OVA from 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. Reboot the node. sshd login for builder is available.
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.
8. Copy `/etc/siq/cluster-setup.properties.sample` to `~/cluster-setup.properties` builder@localhost ~]\$ `cp /etc/siq/cluster-setup.properties.sample cluster-setup.properties`

<code>~/cluster-setup.properties</code>	Note
<code>CLUSTER_NAME=Elasticsearch</code>	Cluster name string.
<code>ES_HOSTS=10.9.9.96,10.9.9.138,10.9.9.139</code>	Required. First host becomes node1.
<code>STUNNEL_ENABLED=False</code>	Set to True for intra-node transport over stunnel.
<code>COUNTRY_ABBR=US</code>	Used for x509 certificate generation.
<code>STATE_PROVINCE=Texas</code>	Used for x509 certificate generation.
<code>ES_USER=elasticsearch</code>	Do not modify.
<code>CITY=Austin</code>	Used for x509 certificate generation.
<code>COMPANY_NAME=IBM</code>	Used for x509 certificate generation.
<code>DEPT_NAME=StoredIQ</code>	Used for x509 certificate generation.
<code>CONTACT_EMAIL=storediqsupport@us.ibm.com</code>	Used for x509 certificate generation.
<code>EXPIRATION_DAY=365</code>	Used for x509 certificate generation.

9. Edit the file:

```
builder@localhost ~]$ vi cluster-setup.properties
```

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

```
[builder@localhost ~]$ /siq/bin/cluster-setup.sh cluster-setup.properties
```

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

11. Type yes when it is prompted to accept the server certificate.

12. Enter the password.

13. Repeat Step 9 - 11 for each additional node.

Location for data is `/siq/var/data/elasticsearch`. Location for logs is `/siq/var/log/elasticsearch`.

14. Test the setup by using `curl -X GET 'http://<your first node IP>:9200'` and `curl -X GET 'http://<your first 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.

Before you begin

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

Procedure

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**. Select **Done**. The Appstack IP prompt appears.

Note: Normally, you must leave this option disabled unless otherwise specified. Contact IBM StoredIQ support and your technical team for more information.

If you need to enable it, you must understand that options to enable Secure Gateway Communications exist on both gateway and data server and they are independent of each other. You must ensure that they match. If this option is selected, see Configuring DA Gateway settings for further guidance.

9. 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.

Before you begin

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

Procedure

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 in the **Primary DNS Host** field 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. 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. If the Dataserver type that is deployed is DataServer - Distributed, select **Enable elasticsearch storage** and enter the IP address of the first node of the Elasticsearch cluster where the cluster was set up.
9. Select **Done**.
10. 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.

Installing and configuring the application stack

Follow these steps to install and configure the application stack.

Before you begin

It is useful to know both the application stack's static IP and gateway IP address. For the gateway's IP address, it is helpful to know this information prior to installation, but it is not mandatory. If you do not have this information, the gateway's configurations can be edited to add this information at a later time.

Procedure

1. Power on the virtual machine.
2. Within the vCenter console, click **Launch Console** to start the virtual-machine console.
3. Within the CentOS EULA screen, select **I Agree** to accept the CENTOS license agreement, and then press **Enter**.
4. Within the IBM EULA screen, select **I Agree** to accept the IBM license agreement, and then press **Enter**.
5. In the Password Utility screen, enter the new password twice and press **Enter**.
6. 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 the IPv4 address of the application stack 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 is the network gateway, not the IBM StoredIQ gateway.
 - f. Enter the IP address for the domain name server in the **Primary DNS Host** field and then press **Enter**.
 - g. Select the **Restart services** option to commit the IP and restart services.
 - 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.
 - c. In the **Primary DNS Host** field, enter the primary DNS IP address.
7. Click **Next**.
8. Within the Appstack configuration screen, set these options. These options set the application stack's capability to send notification email.
 - a. In the **Domain name** field, enter the domain name for generated URLs.
 - b. In the **StoredIQ Gateway** field, enter the IP address of the StoredIQ Gateway server.
 - c. In the **SMTP Server** field, enter the mail server's fully qualified domain name.

Note: All SMTP server 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 25.

- d. In the **SMTP Port** field, enter the SMTP port.
- e. In the **SMTP User** field, enter the login username, if it is supported by the mail server.
- f. In the **SMTP Password** field, enter the login password, if it is supported by the mail server.
- g. In the **Enable TLS** option, select this option to enable TLS encryption, if it is supported by the mail server.
- h. In the **Enable FIPS mode at boot** time option, select this option to enable running your system in FIPS-compliant security mode. By default, FIPS mode is not enabled.

For additional information about SMTP notification, see Appendix E, “Configuring authenticated users for SMTP notifications,” on page 45.

- 9. Click **Next**.
- 10. Within the Certificate configuration screen, perform the procedure in its entirety to generate a certificate.

Note: Certificate configuration is an optional step. To skip certificate configuration, tab to **Finish** and select **Enter**. The certificate can be configured at a later time by logging in as the root user and running this command:
/siq/bin/certcfg

- a. Generate a self-signed root certificate. The resulting certificate can be used as a certificate authority (CA).

Note: Required settings are denoted by an asterisk.

Table 1. Creating self-signed root certificate

Setting	Notes
Common Name *	This is the name of the certificate.
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	This is the length of the key to be created. The default value is 2048.
Days to expiry *	This is the number of days before the certificate's expiration. The default value is 3650.
Root key location *	This is the location of the key. If the location does not exist, one is created.
Root certificate location	This is the output location of the root certificate.

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

Note: Required settings are denoted by an asterisk.

Table 2. Creating certificate signing request

Setting	Notes
Common Name *	This is the host name of the appstack. It must match the domain of the URL that you use. Note: This common name must be different from the common name of Step a of this procedure.
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	This is the length of the key to be created. The default value is 2048.
Key location *	This is the location of the key. If the location does not exist, one is created.
Certificate request location *	This is the output location of the CSR.

- c. Sign the request with the root certificate. Sign a provided certificate-signing request with a certificate-authority-eligible root certificate.

Note: Required settings are denoted by an asterisk.

Table 3. Signing request with root certificate

Setting	Notes
Days to expire *	This is the number of days before the certificate expires. The default value is 3650.
Certificate request location *	This is the location of the CSR to sign.
Certificate location *	This is the output location of the signed certificate.
Root key location *	This is the location of the root key.
Root certificate location *	This is the location of the root certificate.

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

Note: Required settings are denoted by an asterisk.

Table 4. Updating appstack HTTPS certificate

Setting	Notes
Key location *	The location of the key.

Table 4. Updating appstack HTTPS certificate (continued)

Setting	Notes
Certificate location *	The location of the certificate.

At this point, 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.

What to do next

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

Chapter 6. Logging in to the application stack

Follow these steps to log in to the application stack.

Before you begin

The administrator must ensure that these tasks are completed:

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

Procedure

1. Open a browser and enter the IP address or host name of the application stack server.
2. Hit **Enter** to open the login window.
3. Enter superadmin in the email address text field and admin in the password text field.
4. Click **Log in** to open IBM StoredIQ Administrator.
5. If you have several IBM StoredIQ applications that are installed, click the triangle arrow icon and select the application that you need from the list of the available applications.

Chapter 7. Configuration of IBM StoredIQ Platform

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

Default open ports

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

Default open ports on the data server

```
22 (SSH)
-A PROD-ssh -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
443 (UI and Web Services APIs)
-A PROD-https -p tcp -m state --state NEW -m tcp --dport 443 -j ACCEPT
(sigtransport services- inter-communication b/n GW & DS)
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 11103 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 11104 -j ACCEPT
```

Default open ports on the IBM StoredIQ gateway

```
22 (SSH)
-A PROD-ssh -p tcp -m state --state NEW -m tcp --dport 22 -j ACCEPT
443 (UI and Web Services APIs)
-A PROD-https -p tcp -m state --state NEW -m tcp --dport 443 -j ACCEPT
(sigtransport services- inter-communication b/n GW & DS)
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 5434 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 8765 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 7766 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 11102 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 11103 -j ACCEPT
-A PROD-transport -p tcp -m state --state NEW -m tcp --dport 11104 -j ACCEPT
```

Note: When you upgrade your existing IBM StoredIQ data server or gateway to the most recent version, script in iptables does not change, which means that no default minimized port is available. If you need them to be minimized open ports, run this command: `/usr/local/storediq/bin/util/fix_iptables_postupgrade.sh`

Supported chain and rules

```
'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'],
```

Enable or disable ports or service

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

To list current iptables list, run `/usr/local/storediq/bin/util/port_handler.pyc -s`.

To list supported services, run `python /usr/local/storediq/bin/util/port_handler.pyc -l`.

To delete, run `python /usr/local/storediq/bin/util/port_handler.pyc -d <port_number/port_range>`.

For example, to enable https services, run this command: `python /usr/local/storediq/bin/util/port_handler.pyc -e 'PROD-https'`

To delete a service `PROD-objlistsvc`, run on port `21200:21299` with this command: `python /usr/local/storediq/bin/util/port_handler.pyc -d '21200:21299' .`

To open ports for desktop client access to the data server for the 7.6.0.8 OVA deployed systems, follow these steps:

1. Log in as root and run `python /usr/local/storediq/bin/util/port_handler.pyc -e desktop`
2. Run `iptables -L INPUT`. From the output of the command, look for the position of the rule that is named `PROD-reject` from the list, for example, the 6th position on the list.
3. Run `iptables -A INPUT -j PROD-reject`
4. Run `iptables -D INPUT 6`

Note: 6 is the number that the `PROD-reject` rule is positioned on the list like the example of Step 2.

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

Note: These steps are required only on a StoredIQ 7.6.0.8 OVA deployed system. The correct ports are open on an upgraded system.

Configuring the gateway

Use the SSH tool to configure the gateway.

Procedure

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.

Configuring the IBM StoredIQ Platform data server

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

About this task

You must complete this task for every data server.

Procedure

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 if you want to receive reports and notifications through email.

About this task

If the email server was not configured during deployment to receive notifications from the application stack, follow these steps.

Procedure

1. Create the file `/etc/siq/local.yaml` and add this configuration to it:

```
configuration:
  truss.email.backends.SMTPInterface:
    host: <required, SMTP mail server host name, no default>
    port: <optional, TCP port number used by SMTP mail server, default=25>
    username: <optional, user name needed by SMTP mail server, default="">
    password: <optional, password needed by SMTP mail server, default="">
    use_tls: <optional, use secure connection to SMTP mail server, default=False>
  truss.filehandlers.filehandlerinterface.FileHandlerInterface:
    handler:
      hostname: <hostname or ip address>
  truss.controllers.externalurl.ExternalUrlController:
    hostname: <hostname or ip address of appstack>
```

Note: You must follow the exactly same format with the `truss.email` and `truss.filehandlers` indented underneath configuration.

2. Restart the application stack by running the application stack file:
`/siq/bin/monit restart uwsgi`.

What to do next

For additional information about SMTP notification, see Appendix E, “Configuring authenticated users for SMTP notifications,” on page 45.

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.

Before you begin

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

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

Procedure

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.

Chapter 8. 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.

About this task

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:

Note: This process backs up the volume definitions, discovery export records, and data server settings. It does not back up infosets, data maps, or indexes. For more information, see Setting system backup configurations.

Procedure

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: `bootstrap server.mode production`

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.

Appendix A. Application data migration

Application data can be moved from an old system to a new one in different operation systems. Migrations are supported from IBM StoredIQ 7.4 and 7.5.0.x to 7.5.1.0.

You must have administrative privileges within your virtual machine environment to perform the following procedures.

Migrating the data server data

You can move all data by swapping the virtual machine disk.

Procedure

1. Log on to your old data server(s).
2. Using an SSH tool, stop all services with this command: `service siqinit stop` Verify that all services are stopped using this command: `service deepfiler status` It may take some time for all services to stop completely.
3. Using an SSH tool, run this command: `pg_dumpall -U dfuser -c -f db_file.out` This step creates a backup of your database within the `/root` directory. Wait for the dump file to be created. This file can now be copied off this machine.
4. Copy this newly created file off the machine to a location of your choice. Depending on how many infosets and filters were on the gateway, this step may take quite a bit of time. The greater the number of infosets and filters, the longer this process takes.
5. Copy these files off the machine:
 - `/etc/deepfile/siqkwd`
 - `/usr/local/iqdb-base/last-upgrade-file`
6. Shut down the old system.
7. Bring up the new virtual machine and verify that it is running.
8. Using an SSH tool, log into your new VM. Copy these files to your new system:
 - The database backup file, which is named `db_file.out` can be copied to `/ROOT`
 - `/etc/deepfile/siqkwd`
 - `/usr/local/iqdb-base/last-upgrade-file`
9. Restore the data on your new system by issuing these commands.
 - a. Stop all services on the new virtual machine: `service siqinit stop`
 - b. Verify that services are not running: `service deepfiler status`. All services must be stopped completely before you can continue with the migration.
 - c. Start the restore process: `psql -U postgres -d postgres -f db_file.out`
 - d. Wait for the restore process to complete.
10. Run the database upgrade script using this command: `sh /usr/local/iqdb-update/bin/dbversion.sh -skipbackup`
11. Restart the services using this command: `service siqinit restart`
12. Verify that the services are running: `service deepfiler status`

Migrating the gateway data

Follow this procedure to migrate data in the gateway.

Procedure

1. Log on to your old gateway server.
2. Using an SSH tool, stop all services with this command: `service siqinit stop` Verify that all services are stopped using this command: `service deepfiler status`. It may take some time for all services to stop completely.
3. Using an SSH tool, run this command: `pg_dumpall -U dfuser -c -f db_file.out` This step creates a backup of your database in the `/root` directory. Wait for the dump file to be created. This file can now be copied off this machine.
4. Copy this newly created file off the machine to a location of your choice. Depending on how many infosets and filters are on the gateway, this step may take quite a bit of time. The greater the number of infosets and filters, the longer this process takes.
5. Copy these files off the machine:
 - `/etc/deepfile/siqkwd`
 - `/usr/local/iqdb-base/last-upgrade-file`
6. Shut down the old system.
7. Bring up the new virtual machine and verify that it is running.
8. Using an SSH tool, log into your new VM. Copy these files to your new system:
 - The database backup file, which is named `db_file.out` to `/ROOT`
 - `/etc/deepfile/siqkwd`
 - `/usr/local/iqdb-base/last-upgrade-file`
9. Restore the data on your new system by issuing these commands.
 - a. Stop all services on the new virtual machine: `service siqinit stop`
 - b. Verify that services are not running: `service deepfiler status` All services must be stopped completely before you can continue with the migration.
 - c. Start the restore process: `psql -U postgres -d postgres -f db_file.out`
 - d. Wait for the restore process to complete.
10. Run the database upgrade script using this command: `sh /usr/local/iqdb-update/bin/dbversion.sh -skipbackup`
11. Restart the services using this command: `service siqinit restart`
12. Verify that the services are running using this command: `service deepfiler status`

Migrating the application-stack data

Follow these steps to migrate the application stack data.

Procedure

1. Log on to the old application stack server.
2. Using an SSH tool, stop all services with this command: `/siq/bin/monit stop all` Verify that all services are stopped using this command: `/siq/bin/monit status` It may take some time for all services to stop completely.

3. Using an SSH tool, start the PostgreSQL service using this command:
`/siq/bin/monit start postgresql` Then, verify that PostgreSQL has started:
`/siq/bin/monit status`
4. In a screen window, run this command: `/siq/env/postgresql/bin/pg_dumpall -U postgres -c -f db_file.out` It is important to run this command in a screen window so that you can see when the database dump is complete. This step creates a backup of your database in the `/root` directory. Wait for the file to be created. It will take quite a bit of time for this command to run to completion.
5. Using an SSH tool, copy `db_file.out` from the original machine to a temporary location.
6. Shut down the old machine.
7. Start your newly deployed VM.
8. Using an SSH tool, complete these commands:
 - a. Stop all services with this command: `/siq/bin/monit stop all`
 - b. Verify that all services are stopped using this command: `/siq/bin/monit status`. It may take some time for all services to stop completely.
 - c. Start PostgreSQL using this command: `/siq/bin/monit start postgresql`
 - d. Verify that PostgreSQL has started: `/siq/bin/monit status`
9. Using an SSH tool, copy the `db_file.out` file from its temporary location to `/ROOT`.
10. In a screen window, start the database restore process: `/siq/env/postgresql/bin/psql -U postgres -d postgres -f db_file.out`. Please note that it will take some time before this process will run to completion.
11. Using an SSH tool, complete these commands:
 - a. Stop all services with this command: `/siq/bin/monit stop all`
 - b. Verify that all services are stopped using this command: `/siq/bin/monit status`. It may take some time for all services to stop completely.
 - c. Start PostgreSQL using this command: `/siq/bin/monit start postgresql`
 - d. Verify that PostgreSQL has started: `/siq/bin/monit status`
12. Using an SSH tool, issue these commands. Note that a summary completion appears after an issued command is complete. Enter a command only once you see the summary for the previous command.
 - a. `export BAKE_MODULES="spire.tasks mesh.tasks"`
 - b. `/siq/bin/bake spire.migrations.upgrade path=/siq/env/python/lib/python2.7/site-packages/enamel/migrations config=/siq/svc/appstack/appstack.yaml`
 - c. `/siq/bin/bake spire.migrations.upgrade path=/siq/env/python/lib/python2.7/site-packages/bastion/migrations config=/siq/svc/bastion/bastion.yaml`
 - d. `/siq/bin/bake spire.migrations.upgrade path=/siq/env/python/lib/python2.7/site-packages/tidings/migrations config=/siq/svc/tidings/tidings.yaml`
13. Restart the database by issuing these commands.
 - a. Stop all services on the new virtual machine: `/siq/bin/monit stop all`
 - b. Verify that services have stopped: `/siq/bin/monit status`
 - c. Restart PostgreSQL: `/siq/bin/monit start postgresql`
 - d. Verify that PostgreSQL has started: `/siq/bin/monit status`

14. Using an SSH tool, issue these commands. Ensure that each step is complete before moving to the next step.
 - a. `/siq/bin/bake spire.schema.deploy drop=false schema=enamel config=/siq/svc/appstack/appstack.yaml`
 - b. `/siq/bin/bake spire.schema.deploy drop=false schema=tidings config=/siq/svc/tidings/tidings.yaml`
15. Using an SSH tool, complete these commands:
 - a. Start services with this command: `bootstrap server.mode production`
 - b. Verify that all services are started using this command: `/siq/bin/monit status`
16. Once services are up and running on the new VM, log in through the UI and verify that all data has been restored successfully.

Appendix B. Upgrades for IBM StoredIQ Platform

You can obtain the upgrade ISO from the IBM Fix Central. Upgrades consist of upgrading the gateway and data server first and then the application stack. Upgrades are supported from IBM StoredIQ 7.5.1.0 to 7.6.0.1. You must upgrade the previous release version before you upgrade to the next version. For example, you must upgrade to 7.6.0.12 before you upgrade to 7.6.0.13.

Note: It is a good practice that you take snapshots of the gateway, data servers, and application stack in the Vcenter client before you start the upgrade operation.

Go to IBM Fix Central. In the **Product selector** field, enter StoredIQ to find the upgrade ISO.

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.

Gateway and data server upgrades 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.

Procedure

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 `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 ISO: `umount /mnt/cdrom`
7. When upgrading the gateway, run this command but do not run on the data servers: `sed -i '/^icc/d' /etc/deepfile/package_excludes`
8. Proceed with the gateway and data server normal upgrade process.

Downloading the upgrade ISO

Updating the gateway and data servers is the first step in downloading the upgrade ISO.

Before you begin

Download and install either WinSCP or PuTTY to make this process go more smoothly.

Procedure

1. Download the release upgrade ISO from <http://www.ibm.com/support/fixcentral> to the /tmp folder on the gateway, for example, 7.6.0.13-STOREDIQ-upgrade-19.iso. You can use either WinSCP in Step 2 or PuTTY in Step 3 of the following procedure.
2. To use WinSCP, follow these steps.
 - a. Enter the gateway login details, and then click **Login**.
 - b. Browse in the left and right panes to the source and target location, placing the file in the /tmp folder on the gateway.
 - c. Click **Copy** to confirm the copy action.
3. To use PuTTY, follow these steps.
 - a. Open either PowerShell or a command-line window.
 - b. Change the directory to the location of the ISO and PuTTY.
 - c. Run **pscp** to copy the ISO to the gateway.
 - d. Log in to the gateway and verify that it is already copied to the correct location.
4. Log in to the gateway as `util util`.
 - a. Use the keyboard arrows and **Tab** to select **Manage Software Repositories**, and then select **OK**.
 - b. Use the keyboard arrows and **Tab** to select **DVD Update Repository** and then select **Edit**.
 - c. Use the keyboard arrows, spacebar, and **Tab** 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. For example, /tmp/7.6.0.13-STOREDIQ-upgrade-19.iso, and then select **Finish**.
 - e. In the **Manage Update Repositories** screen, select **Done**. You can now return to the Appliance Manager screen.
5. Select **Yes** to save the repository configuration changes.
6. In the **Appliance Manager Screen**, select **Update This Node** and select **OK**.
7. In **Choose Update Repositories**, select **DVD Update Repository** and then select **OK**.
8. 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.
9. Select **Restart**. Appliance Manager opens.
10. Select **Update This Node** again and then select **OK**. The Choose Update Repositories window opens.
11. Select **DVD** and then select **OK**. The available packages are listed again. Several informational messages appear, indicating that the repository is being created, what new package updates are available, update transactions are being created, and that the system is restarted.
12. Repeat this procedure for each of the data servers that you want to upgrade.

Updating the application stack

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

Before you begin

Verify that the gateway and data servers are already updated.

For application stack upgrades from 7.6.0.2 to 7.6.0.4, the following steps must be done before the 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.

About this task

The following procedure is for upgrading IBM StoredIQ application stack.

Procedure

1. Use the SSH tool to log in to the gateway and complete these steps.
 - a. Open the port and type `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, go to `http://gatewayIP/products`. You can see different folders with the application stack versions.
2. Use the SSH tool to log in to the application stack and run these commands.
 - a. Run `bootstrap product.set appstack repository "pkgsource=\"http://gatewayIP/products\""`
 - b. Run `bootstrap product.upgrade appstack HEAD`
 - c. Tail the `appstack.log` file while it is upgrading by entering `tail -f /var/bootstrap/appstack.log` You are notified that the production mode is activated.
3. Using the SSH tool, log in to the gateway as root and complete these steps.
 - a. After the application stack completes updating, run `iptables -D PROD-web -p tcp -m state --state NEW -m tcp --dport 80 -j ACCEPT` to close port 80.
 - b. To restart services on the gateway, enter the service `deepfiler restart` command.

Best practices for updating the application stack

The following list addresses some best practices specific for updating the AppStack.

- Keep the services running in the AppStack before you start an upgrade.
- Look into the `/var/bootstrap/appstack.log` file whenever an upgrade is running.
- After the upgrade is complete successfully, examine the log file for any Traceback or any key words such as `ProgrammingError`, `OperationalError`. Contact the IBM support immediately if you come across these issues.

Appendix C. Customizable reports

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

Note: You must enable specific configuration settings to utilize this functionality. Contact IBM Lab Services for further assistance.

You can customize reports in any of these scenarios:

- Modify reports to carry your organization's custom styles, logos, aligning them with other organization-based artifacts and documentation.
- Alter the format of the content reported in existing reports. For example, you can add more columns, switch axes in a graph, change the units for some values.
- Design reports to contain information that is not found in other, existing reports.

Using ODBC or any client that supports remote access to PostgreSQL, you can connect to either the IBM StoredIQ appserver or the gateway, locating a schema that is named `report_schema`. On both these server types, `report_schema` can be accessed remotely by logging in as a `reportuser`, which is an account that is created for the purposes of reporting on IBM StoredIQ data.

Note: Contact your administrator for `reportuser` credentials for the appstack and gateway.

Appstack

These tables are available for reporting (with SELECT access only) from the `report_schema` on the `enamel` database on the appstack.

- `infoset`
- `infoset_overlay`
- `named_filter`
- `query_filter`

These tables primarily carry information about infosets and filters that are created in IBM StoredIQ Data Workbench.

Gateway

These tables are available for reporting (with SELECT access only) from the `report_schema` on the `dfdata` database on the gateway.

- `distribution_by_accessedtime`
- `distribution_by_filetype`
- `distribution_by_size`
- `distribution_by_filesystem`
- `distribution_by_modtime`
- `distribution_by_objectclass`
- `distribution_by_objectclass_by_owners_by_age`
- `distribution_by_objectclass_by_owner_by_filetype`
- `distribution_by_objectclass_by_owner_by_size`
- `distribution_by_objectclass_by_owner`
- `distribution_by_owner_by_createdtime`

- `distribution_by_owner_by_modtime`
- `distribution_by_owner_by_accessedtime`
- `distribution_by_owner_by_filetype`
- `distribution_by_owner_by_size`
- `volumes`

The "volumes" table lists all the data sources that are defined on all data servers.

The tables that have "distribution" in their prefixes carry summarized information for various metrics for each infoSet. The infoSets are referenced by their IDs here. The mapping to their names is available from the appstack tables. Each "distribution" table reports summaries for "node counts" (number of data objects in a grouping) and "total size" (total size in bytes of data objects in a grouping).

Installing the BIRT Report Designer

The BIRT Report Designer provides core reporting features such as report layout, data source configuration, and scripting features that render content and layout dynamically during report generation. Installing the BIRT Report Designer is a requirement for customizing reports.

Procedure

1. Install Eclipse, which can be downloaded from <http://download.eclipse.org/>.
2. Download BIRT from <http://download.eclipse.org/birt/downloads/>.

Note: Several different packages contain BIRT. You need the all-in-one download, which must be installed on your workstation, specifically BIRT version 4.3.2.

You can also place BIRT into your existing Eclipse environment with the Eclipse Update Manager - BIRT Update Site.

Installing the JDBC driver

A JDBC driver is needed to work with remote PostgreSQL databases with BIRT.

Procedure

1. Download the `postgres_9.2` JAR file from the <http://jdbc.postgresql.org/download.html>.
2. Start Eclipse. Click **File > New > Project**, and in **Business Intelligence and Reporting Tools**, select **Report Project**.
3. Click **New > Report**, and in the Add a new Data Source dialog box, select **JDBC Data Source**.
4. Click **Next**, and then complete the text boxes within the Edit Data Source dialog box.
5. Click **Manage Drivers**.
6. Click **Add** to add the downloaded PostgreSQL JDBC .JAR file.

Adding data sources

The gateway and appstack databases act as the primary data sources for most reporting projects.

About this task

When you create a data source for either the gateway or appstack database, be certain to select **JDBC Data Source** as the data source type.

Procedure

1. In the **Driver Class** text box, enter `org.postgresql.Driver (v9.2)`
2. In the **Database URL** text box, enter `jdbc:postgresql://x.x.x.x/enamel`, where `x.x.x.x` is the IP address or DNS resolvable host name of the appstack.

Note: When the data source that is being created is for the gateway, this URL is used: `jdbc:postgresql://x.x.x.x/dfdata`, where `x.x.x.x` is the IP address or resolvable host name of the gateway.

3. In the **User Name** text box, enter `reportuser`.
4. In the **Password** text box, enter the password that is associated with this user name.

Note: Your administrator might set different passwords for the appstack and the gateway. Contact your administrator to receive this password.

5. Click **Test Connection** to ensure that a connection can be established with the remote database.

Creating a connection profile

When you create data sources repeatedly, you can create a connection profile as an alternative.

Procedure

1. When you create a new data source, click **Create from a connection profile in the profile store**.
2. From **Select a Connection Profile**, click **New**.
3. Click **New** and then select **BIRT JDBC Data Source**, providing the connection details specific to either the appstack or the gateway.
4. In the **Specify file name** text box, enter a name for your connection profile.

Results

After the connection profile is created, it can be reused when you create a new data source. Select **Create from a connection profile in the profile store** and then browse to select the profile from the available list.

Setting report parameters

Some report parameters must be provided when you render the report.

About this task

Currently, IBM StoredIQ supports only one report parameter, which must map to InfoSet ID. It is the only parameter that can be supplied when you render the report.

Procedure

1. Create a report parameter.
2. Enter a useful name such as `infoset_id`.

3. For the message that must be displayed to the user that renders the report, enter the text in the **Prompt text** text box.
4. Select the appropriate **Data Type**, **Display Type**, and **Is Required** values. For an InfoSet ID parameter, these values must be set to **String**, **Text Box**, or **True**.

Adding data sets

A report's charts and tables are built on data sets, which can either use data from the gateway or from the appstack. You can also define a joint data set that joins data sets from both the appstack and the gateway.

Procedure

1. To create an appstack data set:
 - a. Create a data set.
 - b. Select the previously named data source that references the appstack.
 - c. Enter a valid name.
 - d. Enter an SQL query that fetches some data back from the report_schema. If your data set targets infoSet data, it carries a where-clause that looks like where infoSet_id=?. The ? denotes that a report parameter must be used here.
 - e. If a report parameter is to be used in the SQL query, you must add a parameter to the data set.
 - Click **Parameters > New**.
 - Change the **Linked to Report Parameter** value from **None** to the name of your report parameter that is created during data set creation.
2. To create a gateway data set:
 - a. Create a data set.
 - b. Select the previously named data source that references the gateway.
 - c. Enter a valid name.
 - d. Enter an SQL query that fetches some data back from the report_schema. If your data set targets infoSet data, it carries a where-clause that looks like where infoSet_id=?::uuid. The ? denotes that a report parameter must be used here. The type-cast to uuid is important on the gateway.
 - e. If a report parameter is to be used in the SQL query, you must add a parameter to the data set.
 - Click **Parameters > New**.
 - Accept the default values in this dialog box; however, the **Linked to Report Parameter** value must be changed from **None** to the name of the report parameter that is created in the previous step.

Testing report designs

Report designers can test their designs in the BIRT Report Designer itself using the preview feature, which is an HTML rendering of the report.

About this task

Before you upload the customized report, you must preview the report as both HTML and as a PDF. It helps to ensure that the report renders properly on the appstack.

Procedure

1. In Eclipse, open your custom design.
2. Click **Preview** to view the HTML preview of the custom report.
3. Click **Run > Run as > PDF** to view the PDF rendering of the report.

Uploading the report design

IBM StoredIQ supports uploading a report design in a compressed file. Report designs are expected to be in the `rptdesign` format within the compressed file. Currently, only one `rptdesign` is expected to be present within a compressed file. The design can use one or more `.css` files and one or more image files, which can all be packaged within the compressed file.

About this task

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

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

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

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

Procedure

1. Log in to IBM StoredIQ Administrator.
2. Select the **Reports** tab.
3. Click **Create**.
4. Type the name of the report.
5. Optional: Enter the description of the report.
6. Select a report compressed file from your local drive and click **Create** to upload a report design to IBM StoredIQ Administrator. The report design is then also shown in IBM StoredIQ Data Workbench.

Deploying fonts

By default, the IBM StoredIQ reporting engine supports only certain fonts that are provided with the product.

About this task

By default, the IBM StoredIQ reporting engine supports the following fonts:

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

Procedure

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

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

2. Restart IBM StoredIQ services. IBM StoredIQ services must be restarted on the application stack after new fonts are deployed.
 - a. To restart services, issue the command `bootstrap server.mode idle` and wait until `bootstrap server.mode` shows `idle`.
 - b. Run the command `bootstrap server.mode production`
 - c. To check the status, run the command `bootstrap server.mode`

Changing the reportuser password

If needed, the password for the reportuser account on both the gateway and the application stack can be changed.

Procedure

1. To change the password for the reportuser on the application stack:
 - a. Using an SSH tool, log into the application stack as `root`.
 - b. Run the `/siq/bin/change_reportuser_password` script with the new password for the reportuser as a parameter.
2. To change the password for the reportuser on the gateway:
 - a. Using an SSH tool, log into the gateway as `root`.
 - b. Copy the `/siq/bin/change_reportuser_password` from the application stack.
 - c. Run this script with the new password for the reportuser as a parameter.

Appendix D. 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.

Procedure

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.

Appendix E. 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.

Procedure

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.

About this task

Procedure

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 the application stack to send notification emails

By configuring the IBM StoredIQ application stack, you can send email notifications from the IBM StoredIQ application stack for from IBM StoredIQ eDiscovery matter dashboard during report generation.

Procedure

1. Using an SSH tool, log into the application stack as root.
2. Launch the Appstack Configuration utility using this command:
`/siq/bin/appstackcfg`
3. In the **Domain name** field, provide the fully qualified hostname or short hostname for the application stack. For example, enter `cd-7603-as.ibm-lab.local`
4. In the **IP Address** field, provide 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 username of the user you wish to authenticate into the Exchange server.
 - If you authorized any Authenticated User, you can use any username and password as long as that individual is a 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 username.
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 or when using the IBM StoredIQ eDiscovery matter dashboard.

Notices

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

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

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

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

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

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

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

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

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

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

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

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

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

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

The client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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

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

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

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 actual people or business enterprises 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: © Copyright IBM Corp. 2004, 2015. All rights reserved.

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" at <http://www.ibm.com/legal/copytrade.shtml>

Adobe, the Adobe logo, PostScript, and the PostScript logo 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, Windows, Windows NT, and the Windows logo 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.

Other company, product, and service names may be trademarks or service marks of others.

Index

A

- application stack 6, 21, 35
 - update 35
- appstack 1, 4, 24

B

- back up 27
- best practices 35
 - application stack 35
 - updating the application stack 35
- BIRT Report Designer
 - install 38

C

- components of the IBM StoredIQ solution 1
- configuration 23
 - IBM StoredIQ Platform 23
- configure 17, 24, 25
 - application stack 17, 25
 - gateway 24
 - IBM StoredIQ Platform data server 24
- connection profile
 - create 39
- CSV exception export 25
- customizable reports 37

D

- data server 1, 5, 6
 - installation of 16
- data sets
 - add 40
- data sources 4
 - add 39
- deploy 9
 - OVA 9
 - OVF 9
 - virtual appliance 9
- deployment 3
 - deployment planning 3
- Desktop Agent 1
- download 33
 - gateway and data server 33

E

- Elasticsearch 13
 - deploy three-node Elasticsearch cluster 13
- email notification 45
- environment sizing 5
- ESX server 9

F

- fonts
 - deploy 42

G

- gateway 1, 4, 24
 - installation of 14
- gateway server 6

H

- hardware requirements 3
- harvest frequency 5

I

- IBM StoredIQ 1
 - components 1
 - IBM StoredIQ Administrator 1
 - IBM StoredIQ Data Workbench 1
 - IBM StoredIQ eDiscovery 1
 - IBM StoredIQ Platform Data Server 1
 - IBM StoredIQ Policy Manager 1
- IBM StoredIQ Desktop Data
 - Collector 25
- IBM StoredIQ image 27
- IBM StoredIQ Platform 1, 3
- IBM StoredIQ Platform data server 24
- infoset creation 25
- install 17
 - application stack 17
 - IBM StoredIQ Desktop Data Collector 25

J

- JDBC driver 38
 - install 38
- joint data sets 40

L

- legal
 - notices 49
 - trademarks 52
- Livelihood APIs 43
- logging in 21

M

- metadata 5
- metric 5
- Microsoft Hyper-V 11
- migrate 29, 30
 - application stack 30
 - data server 29

- migrate (*continued*)
 - gateway 30
- migration 29
 - application data 29

N

- network 4
- notices
 - legal 49

O

- Open Virtual Appliance (OVA) 3
- Open Virtualization Format (OVF) 9
 - OVA 3, 9
 - OVF 9

P

- port 4, 6
- port ranges 5
- Prerequisites for installing IBM StoredIQ Platform on Microsoft Hyper-V 11
- PuTTY 33

R

- receive reports and notifications 25
- report design
 - requirements 41
 - upload 41
- report designs
 - test 40
- report parameters 39
- reports 37
- reportuser password
 - change 42

S

- SMTP 45
- SMTP notification 45
- SSH tool 24
- stack-provisioning prerequisites 6

T

- TCP 5
 - port ranges 5
- trademarks 52

U

- upgrade ISO 33
- upgrades 33
 - appstack 33
 - data servers 33

upgrades (*continued*)
 gateway 33
 user authentication 45, 46

V

vCenter server 9
virtual appliances 3, 9
vSphere Client 9, 24

W

WinSCP 33



Printed in USA

SC27-6386-15

