

IBM Security Verify Access
Version 10.0.2
June 2021

Administration topics



Contents

Figures.....	ix
Tables.....	xi
Chapter 1. Overview.....	1
Activation level overview.....	1
Tips on using the appliance.....	3
Chapter 2. Getting Started.....	5
Hardware appliance tasks.....	5
Connecting cables and starting the appliance.....	5
Options to configure the hardware appliance.....	5
Connecting a serial console to the appliance.....	5
Determining the system IP address.....	6
Virtual appliance tasks.....	6
Setting up the virtual network.....	6
Installing the virtual appliance by using VMware.....	6
Installing the virtual appliance by using the OVA file.....	7
Installing the virtual appliance by using the vSphere API.....	8
Installing the virtual appliance by using KVM.....	8
Installing the virtual appliance by using Red Hat Enterprise Virtualization (RHEV).....	9
Installing the virtual appliance by using Microsoft Hyper-V.....	10
XenServer support.....	11
Amazon EC2 support.....	13
Microsoft Azure support.....	17
Calculating license usage.....	21
Setting up Cloud Orchestrator support.....	22
USB support on virtual appliances.....	23
License Metric Tool support.....	23
Common tasks.....	24
Command-line interface initial appliance settings wizard.....	24
IBM Security Verify Access Appliance setup wizard by using the local management interface.....	25
Activating the product and buying support.....	25
Silent configuration.....	27
Chapter 3. Initial configuration.....	29
Chapter 4. Managing the appliance.....	31
Local management interface.....	31
Command-line interface.....	31
Web service.....	35
Required header for calling a web service.....	35
Web service error responses.....	35
Configuration changes commit process.....	36
Chapter 5. IBM Security Verify Access Appliance Dashboard.....	39
Viewing system notifications.....	39
Viewing disk usage.....	39
Viewing IP addresses.....	40

Viewing certificate expiry.....	40
Viewing partition information.....	41
Viewing network traffic.....	41
Viewing the status of the appliance in Docker.....	41
Configuring the dashboard.....	42
Chapter 6. Monitoring.....	43
Viewing the event log.....	43
Forwarding logs to a remote syslog server.....	43
Viewing memory statistics.....	44
Viewing CPU utilization.....	45
Viewing storage utilization.....	46
Viewing application interface statistics.....	46
Viewing application log files.....	47
Managing tracing specification.....	48
Chapter 7. System.....	51
Updates and licensing.....	51
Viewing the update and licensing overview.....	51
Installing updates.....	51
Configuring the update schedule.....	52
Configuring update server settings.....	52
Viewing update history.....	54
Installing a fix pack.....	54
Installing a license.....	54
Managing firmware settings.....	55
Managing trial settings.....	55
Installing an extension.....	56
Network Settings.....	57
Configuring general networking settings.....	57
Configuring DNS.....	57
Configuring interfaces.....	57
Appliance port usage.....	59
Configuring aggregated network interfaces.....	62
Configuring static routes.....	63
Testing the connection.....	64
Managing hosts file.....	65
Managing the shared volume.....	66
Managing packet tracing.....	67
Creating a cluster.....	68
Managing cluster configuration.....	69
Managing Distributed Session Cache in Docker.....	85
Managing database configuration in Docker.....	86
System settings.....	87
Configuring date and time settings.....	87
Configuring administrator settings.....	87
Configuring tracing for the local management interface.....	88
Configuring management authentication.....	90
Managing roles of users and groups.....	92
Viewing and updating management SSL certificates.....	95
Managing users and groups.....	95
Managing advanced tuning parameters.....	96
Managing snapshots.....	100
Managing support files.....	100
Configuring system alerts.....	101
Restarting or shutting down the appliance.....	104
Configuring application database settings.....	105

Setting the locale of application log files.....	106
Configuring SNMP monitoring.....	106
Configuring password quality.....	107
Auxiliary Configuration Files.....	108
Secure settings.....	109
Managing SSL certificates.....	109
Managing file downloads.....	116
Chapter 8. Docker support.....	117
Docker image for Security Verify Access.....	120
Docker image for PostgreSQL support.....	126
Docker image for OpenLDAP support.....	128
Docker image for Verify Access Runtime.....	130
Configuration.....	132
Service.....	132
Shared configuration data.....	132
Logging.....	132
Docker image for Verify Access Web Reverse Proxy.....	133
Configuration.....	134
Service.....	135
Shared configuration data.....	135
Logging.....	135
Docker image for Verify Access Distributed Session Cache.....	136
Configuration.....	137
Service.....	137
Shared Configuration Data.....	138
Logging.....	138
Scenarios.....	138
Scenario - Initial configuration.....	138
Scenario - Configuration update.....	138
Scenario - Replicated services.....	139
Scenario - Upgrade.....	139
Scenario - Support files.....	139
Scenario - AAC/Federation runtime configuration.....	140
Orchestration.....	140
Kubernetes support.....	140
Helm Charts.....	152
Docker Compose support.....	153
CLI in a Docker environment.....	157
Distributed Session Cache in Docker environment.....	158
License usage with IBM Security Verify Access deployed on Kubernetes.....	158
Deploy iLMT operator.....	158
Deploy IBM Security Verify Access.....	159
Verify the license server metrics.....	159
Chapter 9. Deployment methodologies.....	161
Cluster support.....	161
Cluster support overview.....	161
Roles and services in a cluster.....	161
Data replication in a cluster.....	163
High availability of cluster services.....	164
High availability for the policy server.....	166
Cluster failure management.....	168
Promoting a node to master.....	168
Removing an unreachable master node from a cluster.....	170
Restricted nodes in a cluster.....	171
Cluster maintenance.....	171

Cluster configuration rules.....	172
Data loss considerations.....	174
Deployment pattern.....	175
Cluster-less AAC Deployment.....	176
AAC Specific Configuration.....	178
FIDO2 Specific Configuration.....	178
Federation Specific Configuration.....	178
Limitations.....	179
Chapter 10. WebSEAL functionality on the appliance.....	181
Chapter 11. Migration.....	183
Migrating an existing WebSEAL instance to the appliance.....	183
Migrating an existing Access Manager software environment to the appliance.....	186
Chapter 12. Configuration changes commit process.....	191
Chapter 13. Runtime environment.....	195
Stopping, starting, or restarting the runtime environment.....	195
Configuring the runtime environment.....	195
Unconfiguring the runtime environment.....	197
Managing runtime configuration files.....	198
Configuring JVM debugging for the runtime profile.....	198
Chapter 14. Users and user registries.....	199
Configuring the runtime to authenticate basic users.....	199
Embedded LDAP server management.....	201
SSL support.....	201
Managing passwords.....	201
Managing suffixes.....	202
Setting debug log level.....	202
Managing federated directories.....	203
Chapter 15. Reverse proxy instance management.....	205
Stopping, starting, or restarting an instance.....	205
Configuring an instance.....	206
Unconfiguring an instance.....	207
Managing web reverse proxy configuration entries.....	207
Managing web reverse proxy configuration files.....	215
Exporting WebSEAL configuration.....	215
Configuring Web Application Firewall	216
Managing administration pages.....	222
Renewing web reverse proxy management certificates.....	223
Configuring Mobile Multi-Factor Authentication.....	224
Chapter 16. Reverse proxy status.....	225
Showing the current state of all instances.....	225
Modifying the statistics settings for a component.....	225
Managing statistics log files.....	225
Archiving and deleting reverse proxy log files with the command-line interface.....	226
Viewing reverse proxy traffic.....	227
Viewing reverse proxy throughput.....	227
Viewing reverse proxy health status.....	228
Viewing front-end load balancer health status.....	229
Viewing average response time statistics.....	229
Viewing security action statistics.....	229

Chapter 17. Junctions.....	231
Creating virtual junctions.....	231
Creating standard junctions.....	232
Managing standard and virtual junctions	234
Chapter 18. Federation management.....	235
Adding a federation.....	235
Removing a federation.....	236
Chapter 19. Authorization servers.....	237
Cleaning up authorization servers.....	237
Creating an authorization server instance.....	237
Deleting an authorization server instance.....	238
Stopping, starting, or restarting an authorization server instance.....	239
Editing an authorization server instance advanced configuration file.....	239
Editing an authorization server instance tracing configuration file.....	239
Renewing authorization server management certificates.....	240
Chapter 20. Clusters.....	241
Replicating runtime settings across the cluster.....	241
Managing Distributed Session Cache.....	241
Chapter 21. Policy management with Web Portal Manager.....	243
Chapter 22. Global settings.....	245
Managing dynamic URL configuration files.....	245
Managing junction mapping JMT configuration files.....	246
Managing client certificate CDAS files.....	247
Managing user mapping CDAS files.....	248
Managing password strength rule files.....	249
Managing forms based single sign-on files.....	250
Managing HTTP transformation files.....	251
Managing RSA SecurID configuration.....	252
Managing the Redis configuration.....	253
Chapter 23. Global keys.....	257
Managing SSO keys.....	257
Managing LTPA keys.....	257
Kerberos configuration.....	258
Managing the default values used by Kerberos.....	259
Managing realms.....	260
Managing domain realm properties.....	261
Managing CA paths.....	262
Managing keytab files.....	263
Chapter 24. Trace data.....	265
Modifying the tracing settings for a component.....	265
Managing the trace files for a component.....	265
Editing the tracing configuration file for the runtime environment.....	266
Updating a tracing configuration file	267
Chapter 25. Logging.....	269
Listing the names of all log files and file sizes.....	269
Viewing a snippet of or export a log file.....	269
Clearing a log file.....	270

Managing transaction logging components and data files	270
Managing reverse proxy log files.....	271
Managing authorization server log files.....	272
Chapter 26. Front-end load balancer.....	273
Scheduling.....	274
Load balancing layer.....	275
Persistence.....	275
Network termination.....	276
Benefits of layer 7 load balancing.....	276
Configuring front-end load balancer.....	277
Front-end load balancer advanced tuning parameters.....	281
Chapter 27. dscadmin command.....	283
replica set show.....	283
replica set list.....	283
session terminate all_sessions.....	284
session terminate session.....	284
session list.....	285
exit or quit.....	285
Chapter 28. API Access Control.....	287
Overview of the API Access Control.....	287
Components.....	287
Authorization.....	288
Resource Servers.....	289
Resources.....	290
Resource Documentation.....	292
Resource Access Policies.....	292
Cross-Origin Resource Sharing (CORS) Policies.....	294
Configuration Auditing.....	294
Storing the IBM Security Verify Access operations for managing Access Control Policies.....	295
Auditing the Verify Access operations that are performed when managing API Access Control components.....	295
Manage Access Control Policies.....	295
Create a new Access Control Policy.....	296
Modify an existing Access Control Policy.....	296
Delete one or more existing Access Control Policies.....	297
Manage Resource Servers and Resources.....	297
Resource Servers.....	297
Resources.....	303
Manage Cross-Origin Resource Sharing (CORS) Policies.....	307
Creating a CORS policy.....	307
Modifying an existing CORS policy.....	309
Deleting one or all existing CORS policies.....	310
Accessibility features for Security Verify Access.....	311
Index.....	313

Figures

- 1. Product activation levels for the IBM Security Verify Access product..... 2
- 2. Services architecture..... 162
- 3. Example cluster architecture.....165
- 4. Communication in a cluster using port 22..... 174
- 5. Sample cluster environment.....175
- 6. Front-end load balancer..... 273
- 7. Example high availability environment.....274
- 8. Network termination..... 276

Tables

1. Valid keys.....	28
2. HTTP error response codes.....	35
3. Valid trace levels.....	48
4. Ports used on the appliance (listen ports).....	60
5. Bonding modes.....	62
6. SSL additional parameters.....	64
7. Configuration database deployment scripts.....	77
8. Runtime database deployment scripts.....	82
9. LMI tracing components.....	88
10. LMI trace levels.....	89
11. Advanced tuning parameters.....	96
12. Security Verify Access Docker image sources.....	120
13. Logs directory structure.....	125
14. Additional environment variables of the ibmcom/verify-access-postgresql image.....	126
15. Supported tags.....	127
16. Additional environment variables of the ibmcom/verify-access-openldap image.....	128
17. Files in the /container/service/slapd/assets/certs directory.....	129
18. Supported tags.....	130
19. Logs directory structure.....	132
20. Logs Directory Structure.....	136
21. Example commands for some common Docker Compose tasks.....	157
22. Possible architectures for clusters that contain multiple nodes.....	164
23. WebSEAL features that the appliance does not support.....	181

24. Directory structure.....	184
25. Supported suffix elements.....	202
26.	203
27. Manage Kerberos configuration settings.....	258

Chapter 1. Overview

The IBM Security Verify Access Appliance is a network appliance-based security solution that provides both access control and protection from web-based threats.

The main features of the appliance include:

- A dashboard for viewing system status such as system notifications and disk usage.
- Analysis and diagnostics tools such as event logs, memory statistics, and CPU utilization.
- Centralized management of settings such as runtime components configuration files, and SSL certificates.
- Control of system settings such as updates, licenses, and network settings.

Most of the features are configurable by using the local management interface (LMI).

The hardware appliance consists of the hardware and preinstalled IBM Security Verify Access Appliance firmware. The preinstalled firmware software can also be obtained separately as a virtual appliance image that you can deploy in a hypervisor environment.

For information about specifications for both the hardware appliance and virtual appliance, see [System Requirements](#).

Activation level overview

Each activation level on the IBM Security Verify Access appliance offers different features. Consider the needs of your environment to determine which activation levels you require.

Security Verify Access Supporting Components: No activation key is required

The Supporting Components provide:

- Appliance management: local management interface and REST APIs
- Policy Server
- Embedded LDAP server
- Authorization Server

Security Verify Access Platform: Activation key is required

The IBM® Security Verify Access Platform secures web applications. To use the web security features, you must activate the Security Verify Access Platform. This activation level includes the following key components:

Web Reverse Proxy

Web Reverse Proxy is a high performance, multi-threaded Web server that applies fine-grained security policy to the IBM Security Verify Access protected web object space. Web Reverse Proxy can provide single sign-on solutions and incorporate back-end web application server resources into its security policy.

For more information, see [Web Reverse Proxy administration](#).

Front-end load balancer

Optimizes resource use and ensures high availability of services. The front-end load balancer accepts requests from clients and determines which backend server is the most suitable to handle the request. It forwards each request to the appropriate server. The front-end load balancer provides persistence for existing sessions.

For more information, see [Chapter 26, “Front-end load balancer,” on page 273](#).

Web application firewall

Helps protect your web servers from malicious traffic and blocks attempts to compromise the system. See [“Configuring Web Application Firewall” on page 216](#).

Distributed session cache

Maintains session state in clustered server environments. See [Distributed session cache overview](#).

Advanced Access Control Module: Activation key is required

The Advanced Access Control Module secures mobile transactions. This activation level includes features, such as:

Context-based access and an authentication service framework

Provides enhanced authentication assurance, context-based access control, and protection from web-based threats.

API protection

Uses the OAuth protocol, which provides API protection for native mobile and other API-based applications.

Device fingerprinting and registration

Stores the device fingerprint of the user in the context-based access database.

To activate this module, you must first activate the IBM Security Verify Access Platform offering.

Federation Module: Activation key is required

The Federation Module provides support for the SAML 2.0 and OpenID Connect protocols.

To activate this module, you must first activate the IBM Security Verify Access Platform offering.

[Figure 1 on page 2](#) summarizes the key features and product activation levels.

Add-on modules

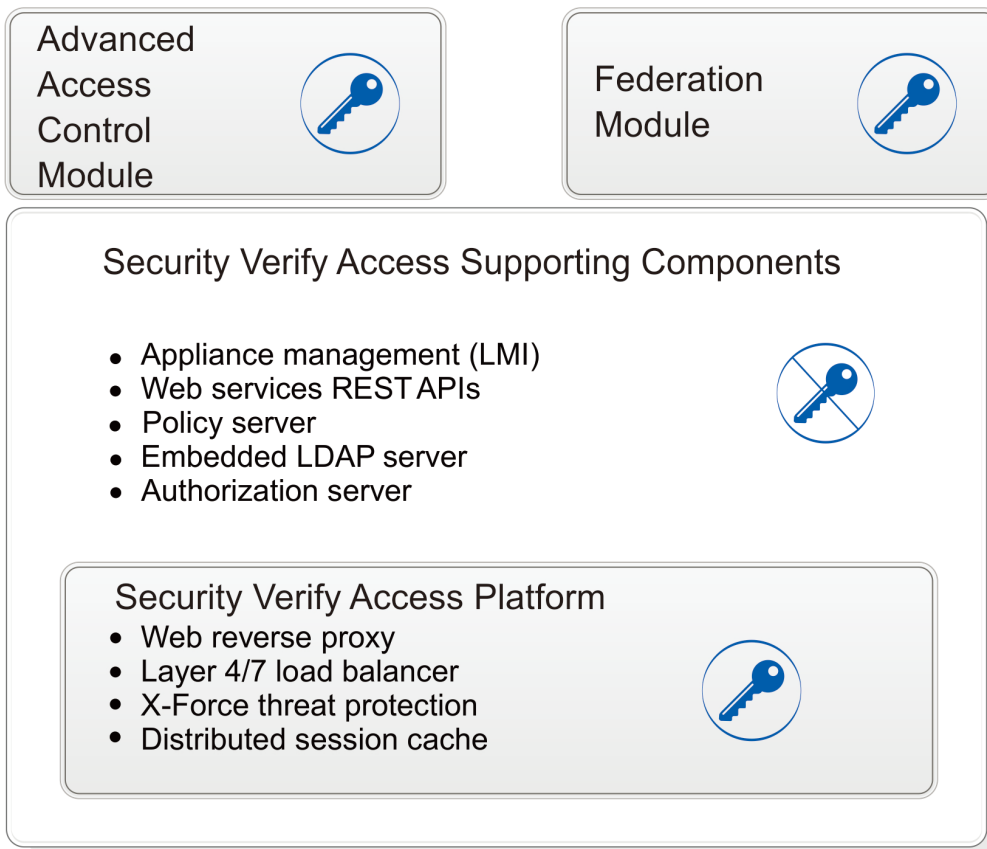


Figure 1. Product activation levels for the IBM Security Verify Access product

Tips on using the appliance

These tips might be useful during the administration of the appliance.

Backup

It is important to back up your appliance frequently. To back up the appliance, use the snapshot facility that is provided by the appliance.

A *snapshot* is a copy of the state of the appliance at a certain time. By using snapshot files, you can back up your appliance and restore the appliance later. It is a good practice to take snapshots regularly and download them from the appliance to serve as backups. However, snapshots can consume much disk space and as such it is best to clean up the old snapshots regularly.

For details about working with snapshots, see [“Managing snapshots” on page 100](#).

Session timeouts

Save your configuration updates in the local management interface (LMI) regularly to avoid any data loss in the event of a session timeout.

LMI sessions expire after the duration of time that is specified by the **Session Timeout** field on the **Administrator Settings** page. When a session timeout occurs, any unsaved data on the current page is lost.

Disk space usage

The disk space in a hardware appliance is limited by the capacity of the installed hard disk. Certain files can use up a significant amount of disk space over time. Such files typically include:

Support files

Support files are used by IBM support personnel to troubleshoot problems with the appliance. The support files contain all log files, temporary and intermediate files, and command output that is needed to diagnose customer support problems. The size of these files can grow large over time. To reduce the disk space that is occupied by these files, download unused support files to an external drive. Then, delete the support files from the appliance. For detailed instructions, see [“Managing support files” on page 100](#).

Snapshot files

Snapshot files record the state that the appliance is in at a certain time. They can be used to restore the appliance to a previous state. The snapshot files are stored on the appliance by default. To reduce the disk space that is used, you can download the snapshot files to an external drive and then delete them from the appliance. For detailed instructions, see [“Managing snapshots” on page 100](#).

The administrator must monitor the remaining free disk space, and take the necessary actions to ensure that there is adequate disk space. The appliance provides a Disk Usage dashboard widget for administrators to monitor the current disk usage. For more information about managing disk space, see [“Viewing disk usage” on page 39](#).

Chapter 2. Getting Started

Complete the following tasks that apply to your appliance format.

Hardware appliance tasks

For the hardware appliance, after you determine where to place the appliance in your network, complete the following tasks.

- Install the network cabling.
- Connect to the local management interface (LMI) or a serial console.
- Configure the initial appliance settings.

Connecting cables and starting the appliance

Connect the appliance to your network after you determine where you want to place it on the network.

Procedure

1. Connect the power cable to the appliance.
2. Connect Management Interface 1 to the network you want to use to manage the appliance.
3. Connect the network cables to the application interfaces.
4. Turn on the appliance.

Options to configure the hardware appliance

You can use either a serial console device that is connected to the appliance or the LMI to configure the hardware appliance.

The LMI is the preferable option as it offers more advanced configuration options.

To use a serial console device, you must connect the console device to the hardware appliance with a serial cable. For instructions, see [“Connecting a serial console to the appliance”](#) on page 5.

To use the LMI to configure the appliance, you must browse to the IP address of the appliance. If you do not know the IP address of the appliance, follow instructions in [“Determining the system IP address”](#) on page 6.

Connecting a serial console to the appliance

You must connect a serial console to the hardware appliance before you can proceed with the configuration tasks through the command-line interface (CLI).

Procedure

1. Connect the console device to the hardware appliance with a serial cable.

Note: Your appliance package might contain a USB serial console cable and a DB-9 serial console cable, or the package might contain only a DB-9 serial console cable. If you use the USB serial console cable and your PC does not recognize the cable, you might need to install the device driver.

The device drivers are available for download from http://public.dhe.ibm.com/software/security/products/infrastructure_protection/USBDeviceDrivers or from the driver supplier at http://www.prolific.com.tw/US/ShowProduct.aspx?p_id=225&pcid=41.

2. If you use a computer as the console device, connect to the appliance with Microsoft Hyperterminal or another terminal emulation program by using the following settings:

Option	Description
Communication Port	Typically COM1
Emulation	VT100
Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

3. Follow the instructions in [“Common tasks”](#) on page 24 to configure initial appliance settings.

Determining the system IP address

If you want to use the LMI to configure the appliance, use one of the following methods to determine the assigned appliance IP address so that you can access the LMI.

- **Method 1:** Use the LCD panel to determine the IP address of the appliance.

1. Press **OK** on the LCD panel to view the main menu.

Note: The **OK** button is labeled with an arrow.

2. Use the arrows to select **IP Address**.

3. Press **OK**.

The LCD panel displays the IP address of the appliance. Take note of the address.

- **Method 2:** Use zero-configuration networking to discover the appliance on your network.

Because the appliance uses a set of industry standard IP protocols, it can be discovered automatically when it is physically connected to your network.

Virtual appliance tasks

You must correctly configure the virtual environment before you install the appliance. Connect to the local management interface or the virtual console to configure the initial appliance settings.

Setting up the virtual network

The administrator who is installing the appliance must be familiar with virtual networking concepts.

The virtual appliance installation does not support scripts. To install multiple virtual appliances, you can install the first appliance manually and make copies of it.

Installing the virtual appliance by using VMware

Use the provided .iso image to install the virtual appliance.

Procedure

1. Create a new virtual machine with your VMware ESX or vSphere.

Note:

- The instructions for creating a virtual machine might differ depending on your VMware ESX or vSphere version. See the VMware documentation that suits your version for specific instructions.
- Ensure that the virtual machine has enough allocated disk space to store the configuration and log file data for the appliance. Allocate at least 100 GB of disk space for the appliance.

- Specify **Virtual Machine Version 7 or later** as your virtual machine version.
 - Specify **Linux** as the guest operating system and **Other 3.x Linux (64-bit)** as the guest operating system version.
 - The memory size has influence over how many WebSEAL instances can be created and how many sessions can be active at a single point in time. The minimum memory size is 4096 MB.
 - A virtual appliance must have a minimum of one and a maximum of eight network adapters.
 - Each network adapter must be of the type **E1000** or **VMXNET 3**. Use **VMXNET 3** for better performance.
 - For SCSI controller, select **LSI Logic Parallel**.
 - For Virtual Device Node, select **SCSI (0:0)**.
 - Diskette, COM ports, and LPT port must be enabled in the BIOS settings of the VM.
 - VMware Tools on the appliance provide the following enhancements:
 - VMware commands for graceful shutdown
 - Improved monitoring
 - Time synchronization with the host operating system unless an NTP server is configured
2. Configure the virtual machine to boot from the .iso file and then start the virtual machine.

Note: If the hard disk that is attached to the virtual machine already contains a Linux partition, the installer always runs in interactive mode.

 - To run the installer in silent mode, wait 10 seconds or press Enter. After the silent installation completes, the virtual machine is shut down automatically. If you want to continue with setting up the appliance, restart the virtual machine.
 - To run the installer in interactive mode, enter `interactive` and then press Enter.
 - a. Enter YES to proceed with the installation. Alternatively if you do not want to proceed with the installation, enter NO to move to the reboot prompt.
 - b. Examine the installation messages to ensure that the installation was successful. After the installation process is complete, unmount the installation media and then press Enter to reboot the appliance.
 3. When the reboot operation is complete, you can start the console-based appliance setup wizard by logging on as the `admin` user with a password of `admin`. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

Installing the virtual appliance by using the OVA file

Use the provided Open Virtual Appliance (OVA) file to install the virtual appliance with VMware.

About this task

The provided OVA file contains a pre-installed appliance image for the VMware hypervisor. This pre-installed appliance image can be used as an alternative installation mechanism to the provided installation ISO. The pre-allocated hard drive is 800 GB. The virtual machine definition includes six network interfaces and 8 GB memory. You can customize these settings after you import the OVA file into VMware.

Procedure

1. Import the provided OVA file into VMware.

Note: The instructions for importing an OVA file might differ depending on your VMware product version. See the VMware documentation that suits your version for specific instructions.

2. Start the virtual machine.

3. When the reboot operation is complete, you can start the console-based appliance setup wizard by logging on as the `admin` user with a password of `admin`. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

Installing the virtual appliance by using the vSphere API

The virtual appliance can be installed by using the vSphere API.

About this task

Security Verify Access provides a sample Python script that utilizes the vSphere API to deploy the appliance. This script can be obtained from the appliance **File Downloads** page. You can examine this script to determine the steps for deploying the appliance.

At a high level, the script has two main functions:

- Create a template on the vSphere server. This step involves uploading an ISO image of the appliance, performing a silent install, and converting this new VM to a template.
- Deploy a template. This step involves cloning this template into a new VM. This step can make use of a silent configuration ISO to configure the networking. This ISO can also be generated by this script. After the silent configuration ISO image has been generated, the script instructs the user to manually upload the image to a datastore of the vSphere environment.

Procedure

1. In the local management interface of the appliance, select **System > Secure Settings > File Downloads**.
2. Expand **Common > Sample > Deploy**.
3. Select the `deploy_isva_to_vsphere.py` file.
4. Click **Export** to save the file to your local drive.
5. Examine the script to determine the steps to deploy and run the virtual appliance. Help on the script can be obtained by running the following command:

```
python deploy_isva_to_vsphere.py --help
```

Note:

- Supported Python versions are 2.7 and 3.4.
 - In Python versions 2.7.9, 3.4.3, or later, unverified SSL connections are disabled. Ensure that the vSphere server certificate is present in the keystore that Python uses.
 - Supported vSphere versions are 4.1, 5.0, 5.1, and 5.5.
 - The **pyVmomi** library is required. It can be installed from the **pip** tool or from <https://github.com/vmware/pyvmomi>.
 - To run the script, you must have the **genisoimage** or **mkisofs** tools in your path.
6. Modify the script as needed.

Installing the virtual appliance by using KVM

The use of Kernel-based virtual machine or KVM is supported. You can use KVM with the provided `.iso` image so that you can run the virtual appliance.

Procedure

1. Create a new virtual machine.

Note:

- The instructions for creating a virtual machine might differ based on the utility that you are using to manage your virtual machines. See the KVM documentation that suits your version for specific instructions.
 - Ensure that the virtual machine has enough disk space that is allocated to store the configuration and log file data for the appliance. Allocate at least 100 GB of disk space for the appliance.
 - The memory size has influence over how many WebSEAL instances can be created and how many sessions can be active at a single point in time. The minimum memory size is 4096 MB.
 - A virtual appliance must have a minimum of one and a maximum of eight network adapters.
 - Each network adapter must be of the type **E1000** or **Virtio**. Use **Virtio** for better performance.
 - The hard disk drive must be configured as a Virtio disk device.
 - If you use certain versions of the Virtual Machine Manager (virt-manager) software to create your virtual machines, it might by default add some CPU definitions that are incompatible with the appliance and thus cause deployment errors. To fix this issue, you can use one of the following methods:
 - From the Virtual Machine Manager console, open the VM definition. Go to **Processor**. Expand the **Configuration** option and then change the value of the **Model** field to **Clear CPU configuration**. Click **Apply**.
 - From the **virsh** shell, edit the virtual machine definition (for example, edit `isva_appliance`). Locate and then remove the `<cpu> . . . </cpu>` entry. Save the file.
2. Configure the virtual machine to start from the .iso file and then start the virtual machine.
 - To run the installer in silent mode, wait 10 seconds or press Enter. After the silent installation completes, the virtual machine is shut down automatically. If you want to continue with setting up the appliance, restart the virtual machine.
 - To run the installer in interactive mode, enter `interactive` and then press Enter.

Note: If the hard disk that is attached to the virtual machine already contains a Linux partition, the installer always runs in interactive mode.

 - a. Enter YES to proceed with the installation. Alternatively if you do not want to proceed with the installation, enter NO to move to the reboot prompt.
 - b. Examine the installation messages to ensure that the installation was successful. After the installation process is complete, unmount the installation media and then press Enter to reboot the appliance.
 3. When the restart operation is complete, you can start the console-based appliance setup wizard by logging on as the `admin` user with a password of `admin`. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

Installing the virtual appliance by using Red Hat Enterprise Virtualization (RHEV)

Use the provided .iso image to install the virtual appliance in a Red Hat Enterprise Virtualization (RHEV) environment.

Procedure

1. Create a new virtual machine.

Note:

- Ensure that the virtual machine has enough disk space that is allocated to store the configuration and log file data for the appliance. Allocate at least 100 GB of disk space for the appliance.
- The memory size has influence over how many Web Reverse Proxy instances can be created and how many sessions can be active at a single point in time. The minimum memory size is 4096 MB.
- A virtual appliance must have a minimum of one and a maximum of eight network adapters.

- Each network adapter must be of the type **E1000** or **Virtio**. Use **Virtio** for better performance.
 - The hard disk drive must be configured as a Virtio disk device.
2. Configure the virtual machine to start from the .iso file and then start the virtual machine.
 - To run the installer in silent mode, wait 10 seconds or press Enter. After the silent installation completes, the virtual machine is shut down automatically. If you want to continue with setting up the appliance, restart the virtual machine.
 - To run the installer in interactive mode, enter `interactive` and then press Enter.

Note: If the hard disk that is attached to the virtual machine already contains a Linux partition, the installer always runs in interactive mode.

 - a. Enter YES to proceed with the installation. Alternatively if you do not want to proceed with the installation, enter NO to move to the reboot prompt.
 - b. Examine the installation messages to ensure that the installation was successful. After the installation process is complete, unmount the installation media and then press Enter to reboot the appliance.
 3. When the restart operation is complete, you can start the console-based appliance setup wizard by logging on as the `admin` user with a password of `admin`. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

Installing the virtual appliance by using Microsoft Hyper-V

Use the provided .iso image to install the virtual appliance.

Procedure

1. Create a new virtual machine with Microsoft Hyper-V.

Note:

- The instructions for creating a virtual machine might differ depending on your Windows version. See the Hyper-V documentation that suits your version for specific instructions.
 - Ensure that the virtual machine has enough allocated disk space to store the configuration and log file data for the appliance. Allocate at least 100 GB of disk space for the appliance.
 - Specify **Generation 1** as the virtual machine generation. The virtual appliance must be run as **Generation 1** virtual machine, **Generation 2** virtual machines are not supported.
 - The memory size has influence over how many Web Reverse Proxy instances can be created and how many sessions can be active at a single point in time. The minimum memory size is 4096 MB.
 - A virtual appliance must have a minimum of one and a maximum of eight network adapters.
 - Each network adapter must be of the type **Network Adapter**. The **Legacy Network Adapter** type is not supported.
 - The Hard Drive and DVD Drive must be attached to IDE Controller 0 and IDE Controller 1, respectively.
 - The following Integration Services are supported:
 - Operating system shutdown
 - Time synchronization
 - Heartbeat
2. Configure the virtual machine to boot from the .iso file and then start the virtual machine.

Note: If the hard disk that is attached to the virtual machine already contains a Linux partition, the installer always runs in interactive mode.

 - To run the installer in silent mode, wait 10 seconds or press Enter. After the silent installation completes, the virtual machine is shut down automatically. If you want to continue with setting up the appliance, restart the virtual machine.

- To run the installer in interactive mode, enter `interactive` and then press Enter.
 - a. Enter YES to proceed with the installation. Alternatively if you do not want to proceed with the installation, enter NO to move to the reboot prompt.
 - b. Examine the installation messages to ensure that the installation was successful. After the installation process is complete, unmount the installation media and then press Enter to reboot the appliance.
- 3. When the reboot operation is complete, you can start the console-based appliance setup wizard by logging on as the `admin` user with a password of `admin`. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

XenServer support

The Security Verify Access appliance can be installed on a XenServer hypervisor (version 6.2 and later).

The Security Verify Access appliance for XenServer is distributed as a pre-installed disk image of the appliance in Virtual Hard Disk (VHD) format. The disk has a fixed size of 100 GB. It is recommended to enable off-the-box logging and auditing to ensure that the disk is not consumed with log files. You can also use the standard installation ISO images to install the virtual appliance on XenServer.

To deploy the VHD appliance image to XenServer, you can use either of the following methods:

- The XenCenter console
- XenAPI or `xe` command line

To install the virtual appliance from the `.iso` image, use the XenCentre console.

Installing the virtual appliance by using the VHD image

Import the VHD image to XenServer with XenCenter to install the virtual appliance.

Before you begin

Make sure that you have the following prerequisites:

- A functional XenServer environment, which is used as the hypervisor to host the VHD image.
- A configured XenCenter installation, which is used to deploy the VHD image.

Procedure

1. In the XenCenter console, expand the XenCenter icon on the left.
2. Right-click the attached hypervisor and select **Import**.
3. In the **Import Source** window:
 - a) Click **Browse**.
 - b) Select the VHD image to be imported and click **Open**.
 - c) Click **Next**.
4. In the **VM Definition** window:
 - a) Specify the name, number of CPUs, and memory of the virtual machine.

Note: In most scenarios, assign the virtual machine at least one processor and 2 GB of memory. These settings can be adjusted after the virtual machine starts running.
 - b) Click **Next**.
5. In the **Location** window:
 - a) Select the destination hypervisor from the drop-down list on the right.
 - b) Click **Next**.
6. In the **Storage** window:
 - a) Select **Place imported virtual disks onto specified target SRs**.

- b) Click **Next**.
7. In the **Networking** window:
 - a) Select the network to be used for the first management interface.
 - b) Click **Next**.
8. In the **OS Fixup Settings** window:
 - a) Select **Don't use Operating System Fixup**.
 - b) Click **Next**.
9. In the **Transfer VM Settings** window:
 - a) Specify the settings to suit your network environment.

Note: If DHCP is not available in the network, a valid IP address, subnet, and gateway must be specified
 - b) Click **Next**.
10. In the **Finish** window, click **Finish** to start the import.

Note: The import operation might take a considerable amount of time to complete. You can click the **Logs** tab to check the progress of the import.
11. Start the imported virtual machine.

Note: At least one network interface must be configured in order for the appliance to start. Sometimes the XenCenter must be restarted before the new virtual appliance can be started correctly.

Installing the virtual appliance by using the ISO image

Use the provided ISO image to install the virtual appliance with XenCenter.

Procedure

1. Create a new virtual machine with XenCenter.

Note:

 - Ensure that the virtual machine has enough disk space that is allocated to store the configuration and log file data for the appliance. Allocate at least 100 GB of disk space for the appliance.
 - The memory size has influence over how many Web Reverse Proxy instances can be created and how many sessions can be active at a single point in time. The minimum memory size is 4096 MB.
 - A virtual appliance must have a minimum of one and a maximum of eight network adapters.
2. Configure the virtual machine to start from the .iso file and then start the virtual machine.
 - To run the installer in silent mode, wait 10 seconds or press **Enter**. After the silent installation completes, the virtual machine is shut down automatically. If you want to continue with setting up the appliance, restart the virtual machine.
 - To run the installer in interactive mode, enter `interactive` and then press **Enter**.

Note: If the hard disk that is attached to the virtual machine already contains a Linux partition, the installer always runs in interactive mode.

 - a. Enter YES to proceed with the installation. Alternatively if you do not want to proceed with the installation, enter NO to move to the reboot prompt.
 - b. Examine the installation messages to ensure that the installation was successful. After the installation process is complete, unmount the installation media and then press **Enter** to reboot the appliance.
3. When the restart operation is complete, you can start the console-based appliance setup wizard by logging on as the admin user with a password of admin. Alternatively, the Appliance Setup wizard can be accessed through the local management interface.

Installing the virtual appliance by using XenAPI or xe command line

The virtual appliance can be installed by using the XenAPI or **xe** command line.

About this task

Security Verify Access provides a sample python script that utilizes the **xe** command line utility to deploy the appliance. This script can be obtained from the appliance **File Downloads** page. You can examine this script to determine the steps for deploying the appliance.

At a high level, the script has two main steps:

- Create a template. This step uploads the VHD to XenServer and creates a VM template from the VHD. The Xen Web Service that is used to upload the image file requires the image to be in RAW or a more efficient proprietary XenServer chunked format. You must convert the VHD image file to a supported image format before uploading it to the XenServer. The script provides an option to perform this conversion.
- Deploy a template. This step creates an instance of the appliance from the template that was created in the previous step.

Procedure

1. In the local management interface of the appliance, select **System > Secure Settings > File Downloads**.
2. Expand **Common > Sample > Deploy**.
3. Select the `deploy_isva_to_xen.py` file.
4. Click **Export** to save the file to your local drive.
5. Examine the script to determine the steps to deploy and run the virtual appliance. Help on the script can be obtained by running the following command:

```
python deploy_isva_to_xen.py --help
```

Note:

- Supported python versions are 2.7.9 and 3.4.3.
 - This script has a dependency on the **requests** and **pexpect** modules.
 - This script is not supported on the Windows platform.
6. Modify the script as needed.

Amazon EC2 support

You can deploy Security Verify Access to the Amazon Elastic Compute Cloud (Amazon EC2) environment.

Amazon EC2 is a web service that provides:

- Scalable computing capacity in the Amazon Web Services (AWS) cloud
- Capability to deploy an Amazon Machine Image (AMI)

Deploying Security Verify Access to Amazon EC2 involves the following processes:

1. Create an Amazon Machine Image (AMI) from the appliance VHD image.
2. Launch an instance of the AMI in Amazon EC2.
3. If you want to use the Amazon Elastic Load Balancing, configure the appliance to send statistical data to Amazon CloudWatch.

Complete these processes either manually or with an automated script. Obtain the `deploy_isva_to_amazon_ec2.py` script from the **File Downloads** page in the **common > samples > deploy** directory on a running appliance. To see script help, run the following command:

```
python deploy_isva_to_amazon_ec2.py --help
```

Note:

- Supported python versions are 2.7.9 and 3.4.3.
- This script has a dependency on the **requests** module.

For details about how to use the Amazon EC2 command line interface to launch an instance, see [Launching an Instance Using the Amazon EC2 CLI](#).

Creating an Amazon Machine Image (AMI) from the Virtual Hard Disk (VHD) file

Upload the appliance VHD image to Amazon EC2 and create an AMI so that it can be deployed in Amazon EC2.

About this task

Follow these steps to manually upload an image and create an AMI with the Amazon EC2 console. If you want to automate this process, use the `deploy_isva_to_amazon_ec2.py` script from the appliance **File Downloads** page.

Procedure

1. Download and install the Amazon EC2 API Tools. You can download the tool from the [Amazon EC2 API Tools](#) page.
2. Run the following commands in the specified sequence to upload the appliance VHD for XenServer to Amazon EC2 and create an AMI.

Sequence	Command	Description
1	ec2-import-volume	Imports the appliance VHD into Amazon EC2.
2	ec2-describe-conversion-tasks	Monitors the ec2-import-volume task to show when the task is complete.
3	ec2-create-snapshot	Creates a snapshot of the imported disk image. This snapshot is required during the AMI registration process.
4	ec2-describe-snapshots	Monitors the status of the snapshot creation to show when the snapshot task is complete.

Sequence	Command	Description
5	ec2-register	Registers a snapshot as a new AMI. You must use the following parameter values when you register the AMI: architecture: x86_64 root device name: /dev/xvda virtualization type: hvm
6	ec2-delete-disk-image	Removes the uploaded disk image from the storage bucket. The image is no longer required after you finish registering an AMI from the image.

Launching the appliance AMI

Launch an instance of the appliance AMI to run the appliance in Amazon EC2.

About this task

Follow these steps to manually launch an instance of the appliance AMI with the Amazon EC2 console. If you want to automate this process, use the `deploy_isva_to_amazon_ec2.py` script that is available from the appliance **File Downloads** page.

Procedure

1. Log in to the Amazon EC2 console.
2. Go to **IMAGES > AMIs**.
3. Select the AMI that you want to launch.
4. Click **Launch**.
5. In the **Choose an Instance Type** window, select an instance type and click **Next: Configure Instance Details**.
6. In the **Configure Instance Details** window, select the options that best fit your environment and click **Next: Add Storage**.
7. In the **Add Storage** window, validate the storage and click **Next: Tag Instance**.
8. In the **Tag Instance** window, add any desired tags and then click **Click Next: Configure Security Group**.
9. In the **Configure Security Group** window, ensure that the selected security group allows inbound SSH and HTTPS access to the appliance. Restrict the access to only those IP addresses from which the appliance is administered. Click **Review and Launch**.
10. Review the details in the **Review Instance** window and click **Launch**.
11. In the **Select an existing key pair or Create a new key pair** window, you can opt to **Proceed without a key pair**. Check the acknowledgment check box. Click **Launch Instances** to proceed.

Note: You do not need to associate a key pair with the instance. If you want to log on to the console of the launched instance, log on as the **admin** user.
12. Go to **INSTANCES > Instances** to check the status of the appliance instance.

Post-installation activities

After you install the appliance in Amazon EC2, complete these activities to enable data transmission to Amazon CloudWatch or change the port on which the LMI listens.

Configuring Amazon CloudWatch support

Configure the appliance to send statistical data to Amazon CloudWatch which can then be used by AWS services to perform load balancing and scaling functions.

About this task

To install the CloudWatch agent the 'IBM Security Verify Access Extension for Amazon CloudWatch Agent' extension must be obtained from [IBM Security App Exchange](#) and installed on the appliance.

If you already have an AWS Identity and Access Management role that is associated with your instance, ensure that it has permissions to perform the Amazon CloudWatch **PutMetricData** operation. Otherwise, you must create a new IAM role with permissions to perform CloudWatch operations and associate that role before you install the extension. Additional information on the access requirements for CloudWatch can be found at [Identity and Access Management for Amazon CloudWatch](#).

After the extension has been installed the metrics collected and the frequency which the agent uploads data can be set in the `aws/amazon_cloud_watch_agent.json` configuration file in [Auxiliary Configuration Files](#). Documentation for supported JSON configuration can be found at the [Cloud Watch User Guide](#). By default the appliance reports memory, swap, and disk space usage metrics to Amazon CloudWatch.

The Access Key and Access Secret associated with this IAM role must be updated in the `aws/credentials` file in the [Auxiliary Configuration Files](#).

Procedure

1. In the appliance local management interface, go to **System > Updates and Licensing > Extensions**.
2. Click **New**.
3. Select the extension file which was obtained from IBM Security App Exchange.

Note: The file which is downloaded from IBM Security App Exchange is a zip file, and the extension file must first be extracted from this zip file.

4. Click **Next**.
5. Supply the requested information, including the Amazon CloudWatch agent RPM and signature files.
6. Click **Install**.

Configuring the local management interface port

By default, Amazon EC2 supports running an instance with a single network interface. To run the appliance with a single network interface, you might want to change the port on which the local management interface listens so that it can be used by other services on the appliance, such as the Web Reverse Proxy.

About this task

You can use the **System > Administrator Settings** page in the local management interface to configure this setting.

If you want to automate this configuration, use the `deploy_isva_to_amazon_ec2.py` script that is available from the appliance **File Downloads** page.

Procedure

1. In the appliance local management interface, go to **System > System Settings > Advanced Tuning Parameters**.

2. Click **New**.
3. Enter `lmi.https.port` in the **Key** field.
4. In the **Value** field, enter the appropriate port so that the port on which the local management interface listens can be used by other services on the appliance.
5. Click **Save Configuration**.
6. Deploy the changes.

Microsoft Azure support

You can deploy Security Verify Access to Microsoft Azure environments.

Deploying Security Verify Access to Microsoft Azure involves the following processes:

1. Create an Azure-ready VHD or obtaining the Azure-ready VHD.
2. Uploading an Azure-ready VHD to Microsoft Azure.
3. Create an image from the uploaded VHD.
4. Deploy the image as a new virtual machine using Azure Portal or the command line.

Complete these processes either manually or with an automated script. Obtain the `deploy_isva_to_azure.py` script from the **File Downloads** page in the **common > samples > deploy** directory on a running appliance. To see script help, run the following command:

```
python deploy_isva_to_azure.py --help
```

Note: The following restrictions apply to the `deploy_isam_to_azure.py` script:

- Supported Python versions are 2.7.9 and 3.4.3.
- This script has a dependency on the `requests` module.
- This script has a dependency on Microsoft Azure SDK for Python.

Creating a custom size Azure compliant Virtual Hard Disk (VHD) file

IBM provides an Azure-compliant VHD file that can be used to deploy Security Verify Access to Azure.

About this task

The size of the VHD file is 800 GB. If you want to use a size other than 800 GB, you can create a custom pre-installed Security Verify Access image for Azure manually. After the Security Verify Access installation finishes, it is not possible to resize the hard disk. This process requires a Microsoft Hyper-V environment and the Security Verify Access firmware installation ISO.

This procedure can be automated using the `generate_azure_image.ps1` Powershell Script that can be obtained from the **File Downloads** page in the **common > samples > deploy** directory on a running appliance.

These steps apply to Hyper-V Manager version 10 and similar.

Procedure

1. In the Hyper-V Manager, create a new virtual machine using the wizard. During the wizard:
 - a. When prompted to **Specify Generation**, select the **Generation 1** option.
 - b. When prompted to **Assign Memory**, enter 2048MB or more. This amount can be changed later after installation.
 - c. When prompted to **Configure Networking**, no network connection is required.
 - d. When prompted to **Connect Virtual Hard Disk**, create a new virtual hard disk. Set the size of the virtual disk to the desired custom size. This size can not be changed after installation finishes.
 - e. When prompted for **Installation Options**, attach the Security Verify Access installation ISO.

2. Start the newly created virtual machine. The virtual machine boots from the Security Verify Access installation ISO and automatically installs the Security Verify Access firmware. When this process is complete, the virtual machine shuts down automatically.
3. Wait for the firmware to install and for the virtual machine to shut down.
4. On the **Actions** tab, click **Edit Disk**. The **Edit Virtual Hard Disk Wizard** is started. During the wizard:
 - a. When prompted to **Locate Disk**, select the VHD file associated with the virtual machine created earlier.
 - b. When prompted to **Choose Action**, select the **Convert** option.
 - c. When prompted to **Choose Disk Format**, select **VHD**. Azure does not support the VHDX format.
 - d. When prompted to **Choose Disk Type**, select **Fixed** size. Azure does not support dynamically expanding or thin-provisioned disks.
 - e. When prompted to **Configure Disk**, choose a new location to save the converted disk to.
5. After the **Edit Virtual Hard Disk Wizard** is complete, the newly converted VHD is ready to be uploaded to Microsoft Azure.

Note:

- The Security Verify Access firmware must not be configured before preparing it to upload to Azure. If the machine is not in the unconfigured state when first started on Azure, it will not correctly detect the Azure environment.
- It is possible to convert the VHD using other methods, such as the Powershell extensions for Hyper-V and qemu-img.
- The firmware installation must take place in a Microsoft Hyper-V environment. For example, you can not install Security Verify Access in VMware and convert it to an Azure-appropriate VHD. The hypervisor that the Security Verify Access firmware is installed in must be the same as its intended execution environment. Microsoft Hyper-V Generation 1 is considered to be the same hardware as Microsoft Azure by the Security Verify Access firmware.
- For details about the VHD requirements, see the General Linux Installation Notes topic on the Microsoft Azure documentation website.

Uploading an Azure-compliant VHD to Azure and creating an Azure Image

To deploy a virtual machine in Microsoft Azure, an Azure-compliant VHD file that contains the Security Verify Access firmware must be uploaded to a storage account and then used to create an image. The created image artifact acts as a template and can be deployed multiple times.

About this task

These instructions demonstrate how to perform the steps using the Azure Portal (portal.azure.com). But you can also use the Azure CLI tools or any other Azure capable API to complete these steps.

Procedure

1. Upload the VHD file using the Azure Portal.
 - a. In the **Azure Portal**, select **Storage Accounts**.
 - b. Select the storage account where the Security Verify Access VHD file will be uploaded to.
 - If you do not have a storage account, click **Add** to create one.
 - Note that the selected location will dictate where the image can be created and subsequently deployed to.
 - c. Under **BLOB SERVICE**, select **Containers**.
 - d. Select a container to upload the Security Verify Access VHD file to.
 - If you do not have a storage container, click **Add Container** to create one.

e. Click **Upload** and select the Azure-compliant Security Verify Access VHD file to upload.

- Ensure that the Blob type is set to **Page Blob**.

This process might take a long time depending on your network connection and the location of your Azure storage account.

2. Create an image using the Azure Portal.

a. In the Azure Portal, select **Images**.

b. Click **Add** to create a new image.

i) Give the image a name. Remember that this image is a template that will later be deployed to a virtual machine with a different name.

ii) Ensure that the location is the same as the location of your storage account.

iii) In the OS disk section:

a) Select Linux and the OS type.

b) Click **Browse** on the **Storage Blob** field. A new panel will list your storage accounts. Using this panel, navigate through the storage account and container to locate the Security Verify Access VHD that was uploaded.

iv) Click **Create** to begin the image creation process. This process typically takes minutes to complete.

c. When the process has completed, return to the **Images** panel and verify that the new image was created.

This image can now be used to deploy new Security Verify Access virtual machines in Azure.

Creating a Security Verify Access virtual machine from an image in Azure

An image artifact in Azure can be used to create a new virtual machine in Azure. The same image can be deployed multiple times to create multiple Security Verify Access virtual machines.

About this task

These instructions demonstrate how to perform the steps using the Azure Portal (portal.azure.com). But you can also use the Azure CLI tools or any other Azure capable API to complete these steps.

Procedure

1. In the **Azure Portal**, select **Images**.

2. Select the previously created Security Verify Access image.

3. On the **Overview** panel, click **Create VM**.

a) On the **Basics** page:

i) Enter a name for the new virtual machine.

ii) Enter a user name, select the **Password Authentication** type and provide a password. You must provide a user name and password for accessing the management console. When running on Microsoft Azure, the default admin account is not created.

iii) Complete the form and click **OK**.

b) On the **Choose a size** page:

i) Select an appropriate size for the new virtual machine, keeping in mind that the recommended minimums are 4 GB of memory and 4 CPU cores.

ii) Click **Select** to continue.

c) On the **Settings** page:

i) Configure the network settings.

Note: It is not possible to configure more than one network interface from the Azure Portal. Additional interfaces can be added using the Azure CLI 2.0 or equivalent.

- ii) Click **OK** to continue
- d) On the **Summary** page, revise the configuration and click **OK** to create the Security Verify Access virtual machine.

Unsupported functionality for Security Verify Access in Microsoft Azure

Verify Access virtual machines can be deployed on Microsoft Azure with the support for basic Infrastructure as a Service capability. Microsoft Azure specific guest extensions and functionality that depends on the fabric layer are not supported.

Verify Access can be deployed into Microsoft Azure as a Linux virtual machine. Microsoft provides many operating system level capabilities for Linux virtual machines that are running on selected Linux distributions. These capabilities do not support Verify Access virtual machines.

Microsoft Azure can provide hypervisor and network level metrics about an Verify Access virtual machine such as CPU, disk, and network utilization.

Microsoft Azure runtime features which require the installation of Microsoft VM extensions. Verify Access is not a supported guest operating system for any of Microsoft Linux VM extensions.

The Windows Azure Agent provided in Verify Access is capable of bootstrapping Verify Access in Azure and reporting a heartbeat signal to the Azure fabric.

The list of unsupported Microsoft Azure Runtime features:

Settings

- Networking
 - Attaching additional network interfaces is not supported.
- Disks
 - Adding additional data disks is not supported.
- Extensions
 - Installing extensions is not supported in Verify Access. This includes Microsoft's standard extensions such as `enablevmaccess`, `LinuxDiagnostic`.
- Identity

Operations

- Backup
 - Use the Verify Access snapshots functionality for back up or restore capabilities.
- Update management
- Inventory
- Change tracking
- Configuration Management
- Run Command

Monitoring

- Insights
- Diagnostic settings
 - Boot diagnostics can be used to view the Serial log which displays the Microsoft Azure agent log
- Logs

Support and troubleshooting

- Reset password

Running Security Verify Access in Microsoft Azure

When a Security Verify Access virtual machine is deployed in Microsoft Azure, by default interface 1.1 will be configured with a single DHCP IP address of the management type, which can be used to access the LMI and SSH. The Azure fabric will assign the networks private IP address specified during deployment to this adapter using DHCP.

By default, no ports are forwarded from the public IP address to the private IP address.

Additional interfaces can be configured using the Azure command line tools. The Azure Portal does not provide the capability of creating a virtual machine with more than one interface or for adding additional interfaces to an existing virtual machine.

Addresses other than the first private IP address on 1.1 must be manually configured within Security Verify Access. Configure Security Verify Access's network settings to match the private IP addresses configured on each adapter in Azure.

The Security Verify Access virtual machine runs the Windows Azure Agent daemon to communicate with the Azure fabric.

- The log file can be viewed on the application log files page under `azure/waagent.log` or by viewing the **Boot Diagnostics** panel in the Azure Portal.
- The Windows Azure Agent will periodically make requests to an internal Azure endpoint (typically within 168.0.0.0/8 169.0.0.0/8) to report deployment and heartbeat status.

Calculating license usage

The IBM Security Verify Access virtual appliance can generate scan results which can be uploaded to the IBM License Metric Tool. See [License metric tool support](#) for more information. To calculate license usage manually, create a Processor Value Unit (PVU) report.

About this task

To manually create the Processor Value Unit (PVU) report you must determine the number and speed of the central processing units (CPUs) on the virtual machine (VM).

Procedure

1. Select the hypervisor that you are using.

VMware

- a. Open the vSphere Client and connect to the IBM Security Verify Access appliance.
- b. Supply the host name and the user name and password.
- c. Select the IBM Security Verify Access appliance from the list of VMs.
- d. Select the **Summary** tab to view the number of CPUs assigned. In the **General** section of the tab there is a line similar to the following entry.

```
CPU: 1 vCPU
```

- e. Select the **Resource Allocation** tab to view the speed of the processors. The **CPU** section of the tab displays information similar to the following entry:

```
Host CPU 0 MHz ---> 2800 MHz  
Consumed: 52.00 MHz
```

- f. Exit the vSphere Client. Retain this information for use in the next steps.

KVM

For more information, see the KVM documentation.

2. Consult the following document for specific instructions on how to calculate the PVUs for the target application (the virtual appliance). See page 8 of the document:

http://public.dhe.ibm.com/software/passportadvantage/SubCapacity/x86_Scenarios.pdf

3. Use the data that you collect to place entries in the following spreadsheet. See the instructions within the spreadsheet.

http://public.dhe.ibm.com/software/passportadvantage/SubCapacity/Manual_Calculation_of_Virtualization_Capacity_Apr_2012.xls

4. Retain the spreadsheet and data in the event of a license compliance audit.

Setting up Cloud Orchestrator support

The virtual appliance has basic support for Cloud Orchestrator as imported KVM virtual images. It is possible to run the appliance within a Cloud Orchestrator environment and use it to perform basic virtual image management tasks.

About this task

Consider these limitations before you set up Cloud Orchestrator support:

- The virtual appliance does not have full support for the orchestration and pattern building capabilities of Cloud Orchestrator.
- The appliance can be run only in KVM regions.
- The appliance can be imported and managed in the Virtual Image Library, but cannot be extended with the IBM Image Construction and Composition Tool.
- The appliance cannot be used in pattern-based deployments.
- The basic operations that are provided by Cloud Orchestrator for the imported appliance image include the ability to stop, start, or delete the virtual machine.
- The appliance must be run with a single network interface.

To use the virtual appliance as an imported virtual image within a Cloud environment, use the following high-level procedure. For more information, see the [IBM Cloud Orchestrator Knowledge Center](#).

Procedure

1. Create a KVM virtual machine image and install the appliance firmware.

Note: After the installation is complete, remove the installation media and shutdown the machine. Do not go through the first steps wizard before you deploy the image in the cloud.

See [Installing the virtual appliance by using KVM](#) for more detailed instructions.

2. Import the virtual image to the Cloud Orchestrator Virtual Image Library.
3. In the Virtual Image Library, check out the image to an operational repository in the KVM region where you plan to deploy the appliance.
4. Use OpenStack to deploy the virtual image within this KVM region.

For example, to deploy the virtual image from the command line, perform these steps on the KVM region server:

- a. Set the environment variables for running OpenStack nova.

```
# source ~/openrc
```

- b. Verify that the appliance image is available in the image repository.

```
# nova image-list
+-----+-----+-----+-----+
| ID                | Name          | Status | Server |
+-----+-----+-----+-----+
| 9ec1d9ec-2df9-44f6-938c-2533a4d48859 | isva          | ACTIVE |        |
+-----+-----+-----+-----+
```

- c. Issue the nova boot command to start a new instance of the appliance image.

```
nova boot --image isva --flavor m1.medium isva
```

- d. Monitor the status of the new instance using the nova list command.

```
# nova list
+-----+-----+-----+-----+
| ID                | Name                | Status | Networks |
+-----+-----+-----+-----+
| 43f3e09c-a64d-4e11-8827-2d354be3d625 | my-isva-appliance  | ACTIVE | public=172.20.96.1 |
+-----+-----+-----+-----+
```

- e. The appliance is now started and the local management interface and web services interfaces are listening on the given IP address.
5. After the machine is running in the OpenStack KVM environment, you can import it into Cloud Orchestrator.
 - a. Log in to the Cloud Orchestrator management web UI.
 - b. Go to **Configuration > Hypervisors**.
 - c. Locate the hypervisor where the appliance is running.
 - d. Expand the virtual machines section and locate the appliance image.
 - e. Select **Manage > Import the Virtual Machine**.

Results

The appliance virtual machine is now visible in the Cloud Orchestrator UI on the **Instances > Virtual Machines** page.

Related tasks

[“Installing the virtual appliance by using KVM” on page 8](#)

The use of Kernel-based virtual machine or KVM is supported. You can use KVM with the provided .iso image so that you can run the virtual appliance.

USB support on virtual appliances

Administrators of virtual appliances can use a physical USB drive for tasks such as uploading a new firmware, uploading or downloading a snapshot file, and downloading a support file.

To use a USB drive with the virtual appliance, complete these steps:

1. Format the USB drive with the FAT32 file system.
2. Attach the USB drive to the appropriate location.
 - If you use XenServer or KVM as your hypervisor, attach the USB drive to the machine that is running the hypervisor.
 - If you use VMWare vSphere as your hypervisor, attach the USB drive to the machine that is running the vSphere client.
3. Update the virtual machine definition to reference the USB device. The required steps are specific to the hypervisor that you use.

License Metric Tool support

In version 10.0.0.0 and later, the appliance firmware contains the License Metric Tool Disconnected Scanner which can be used to initiate software and capacity scans, and prepare scan results that can be uploaded to the License Metric Tool.

The embedded scanner automatically performs the following scans on virtual appliances:

1. Capacity scans are performed every 30 minutes
2. Software scans are performed once every week

Accessing the scan results

The scan results can be found on the LMI Application Log Files page within the lmt directory. The appliance will keep only the latest scan result package.

See [View application log files](#) for further details about accessing the scan results.

Manually performing a scan

A scan can be manually initiated using the command-line interface.

Connect to the appliance command-line interface and execute the following command:

```
lmt scan
```

When the scan is completed, the scan results are made available from the appliance **Application Logfiles** page.

See [Command-line interface](#) for further details about accessing the command-line interface.

Common tasks

These tasks are common for both the hardware appliance and the virtual appliance.

You can choose either of the following methods to configure initial appliance settings.

- Command-line interface (CLI)
- Local management interface (LMI)

The LMI method offers more advanced configuration options.

Command-line interface initial appliance settings wizard

The initial appliance settings wizard runs the first time that an administrator logs in to the command-line interface (CLI) of an unconfigured appliance.

Navigation

You can move between screens in the wizard using the following options:

- p: Previous Screen
- n: Next Screen

To cancel the setup process at any time, use the exit command.

Modules

You must configure the following modules to set up your appliance:

Module	Description
Welcome	Describes the appliance settings that you can configure using the wizard.
Software License Agreement	Describes the appliance license agreement, IBM terms, and non-IBM terms.
FIPS 140-2 Mode Configuration	Enable this option to turn on compliance for NIST SP800-131a. If you enable this option, the appliance is automatically restarted before it continues on with the rest of the setup.

Module	Description
	Note: Enable this option only if you must comply with the NIST SP800-131a requirements. There is no advantage to enabling this option if your installation does not require it. To disable NIST SP800-131a compliance, you must reinstall the appliance.
Password Configuration	Changes your password.
Host Configuration	Changes the host name.
Management Interface Settings	Configures the management network interfaces. Displays device settings and the current working-set policy for the primary and secondary interfaces.
DNS Configuration	Configures the DNS servers that are used by the appliance.
Time Configuration	Configures the time, date, and time zone on the appliance.

Local management interface Appliance Setup wizard

The Appliance Setup wizard runs the first time that an administrator logs in to the local management interface (LMI) of an unconfigured appliance.

After you log in to the LMI for the first time, follow the Appliance Setup wizard to complete the initial configuration of the appliance. The tasks that you must complete for the initial configuration include:

- Read and accept the License Agreement.
- Download and install the license file. You must install the license to download the firmware for the hardware appliance updates.
- Depending on your requirements, choose whether to enable the FIPS option to turn on compliance for NIST SP800-131a. If you enable this option, the appliance is automatically restarted before it continues on with the rest of the setup.

Note: Enable this option only if you must comply with the NIST SP800-131a requirements. There is no advantage to enabling this option if your installation does not require it. To disable NIST SP800-131a compliance, you must reinstall the appliance.

- Set the appliance password.
- Configure the networking, which includes the host name, management interface settings, and DNS configuration.
- Configure the application interface settings.
- Configure the date and time settings.

When you complete the basic configuration, a summary screen displays. Review the details on the completion page and click **Complete Setup**.

Activating the product and buying support

Activate the product after installation so you can use all available features. You can optionally import a support license file to receive updates to the appliance.

Before you begin

Obtain your activation key and support license:

- Download your activation key from your account on [Passport Advantage](https://www-112.ibm.com/software/howtobuy/softwareandservices) at <https://www-112.ibm.com/software/howtobuy/softwareandservices>.
- Obtain your support license by following the instructions in the Welcome email that was sent by IBM.

Note: If you cannot locate your Welcome email, go to the [IBM Security Systems License Key Center](https://ibmss.flexnetoperations.com) at <https://ibmss.flexnetoperations.com>. Review the FAQs to find out how to obtain support.

About this task

You can complete the following actions from the **Support License and Product Activation** panel:

- Import the activation key, which is required.
- Import the support license, which is optional. Import the license if you want to install service release updates.

The activation key is a permanent activation for the product. Activation keys have no expiration date.

Entitlement for X-Force updates for the database is provided automatically with the product. A third-party geolocation database is provided with sample data. You must purchase separately the full set of geolocation data.

You can review activation and support license information for your installed product packages, including specific product activations, service agreements, and expiration dates from this panel.

Procedure

1. Log in to the local management interface.
2. Click **System > Updates and Licensing > Licensing and Activation**.
3. Perform one or more of the following actions:
 - Import the activation key and deploy the changes:
 - a. In the **Licensing and Activation** window, click **Import** under **Activated products**.
 - b. Browse to the activation key file that you downloaded from Passport Advantage.
 - c. Select the activation file.
 - d. Click **Open**.
 - e. Click **Save Configuration**.
 - f. Deploy the changes:

Note: You do not need to deploy changes immediately after you install the activation key. However, you must deploy changes before you can take a snapshot of the product.

 - i) In the undeployed change message, click **Click here to review the changes or apply them to the system**.
 - ii) Click **Deploy**.
 - g. The activated product name and version are displayed in the Products table. To view the software license agreement, click: **View Service Agreement**.
 - Optional: Import the product support license so that you can update the appliance:
 - a. In the **Licensing and Activation** window, under **Support license**, click **Select License**.
 - b. Browse to the support license file that you downloaded from IBM Security Systems License Key Center.
 - c. Select the license file.
 - d. Click **Save Configuration**.

Results

The menu in the local management interface refreshes to show the menu for the activated product.

If you imported a Support license, you can update the product with service releases. See [“Installing updates”](#) on page 51.



Attention: Ensure that the activation is completed before attempting any other activities using the local management interface.

Silent configuration

You can configure an appliance silently after installation with the web service interfaces by providing a metadata image that contains essential configuration data.

After the appliance firmware has been installed, shut down the machine. The ISO image that contains the configuration meta-data can then be attached to the appliance in preparation for the initial boot of the installed firmware. Once the appliance has successfully booted, it will mount the ISO image and then use the configuration meta-data to automatically configure the network.

The metadata image can be created with the local management interface or manually with a text editor.

If you use a manually created metadata image for the initial configuration of an appliance, the appliance boots up with the configured network settings automatically, but the first-steps wizard must be completed manually. You can use the local management interface or the web service interfaces to perform the first-steps configuration. To silently configure the appliance without the need to complete the first-steps wizard manually, you must use a metadata image that contains the system policy. Such metadata images can be created only through the local management interface.

See the `isva_config_sample.py` script available from the **File Downloads** page of the local management interface as an example for silent configuration with scripts.

Creating a metadata image with the local management interface

You can create a metadata image that contains essential configuration data for the initial setup of an appliance with the local management interface. This image can later be used for the silent configuration of a new virtual appliance.

About this task

A metadata image that is created with the local management interface provides more information than a manually created metadata image does. For example, you can choose to include the system policy when you create the image through the local management interface. If the system policy is included in the image, it is possible to accept the license agreement silently and complete the first-steps wizard automatically on first boot.

Procedure

1. From the top menu, select **System > Secure Settings > Silent Configuration**.
2. Enter the hostname to be configured on the new appliance.
3. Select the **IPv4**, **IPv6**, or both check boxes to specify static IP addresses.
 - If the **IPv4** check box is selected, complete the IPv4 section.
 - If the **IPv6** check box is selected, complete the IPv6 section.
4. To include the system policy, select the **Include system policy** check box.

Note:

- The system policy excludes the following configuration:
 - Management interfaces
 - Application interfaces
 - Cluster configuration
 - Advanced Access Control runtime configuration
- Besides the previously mentioned policy exclusions, there are also a few non-policy exclusions. The support license, Security Verify Access runtime environment, reverse proxy instances, authorization server instances, local LDAP server, policy databases, custom pages, and other files that are uploaded to the appliance are not included in the system policy.

- If the system policy is included, when the silent configuration takes place on a new appliance, the license agreement is automatically accepted and the first-steps wizard is automatically completed.
5. At the bottom of the page, click **Generate CDRom Image** or **Generate USB Image** to download an image that contains this metadata.
 - To use the USB image, write the IMG file to a partition on a USB device by using the dd tool.
 - To use the CD-ROM image, attach the ISO file to a CD-ROM device on your virtual appliance.

Creating a metadata image manually

You can create a metadata image that contains the initial network configuration for interface 1.1 manually with a text editor.

About this task

The metadata file is a plain text file that contains a list of key-value pairs. The file must be named app-metadata and at the root of the file system of the ISO image to be mounted.

Procedure

1. Create a text file with the name app-metadata at the root of the file system of the attached device.
2. Edit the content of the text file as needed.

The valid keys are as follows:

<i>Table 1. Valid keys</i>	
Key	Description
network.hostname	The appliance hostname
network.1.1.ipv4.address	The initial IPv4 management IP address on interface 1.1
network.1.1.ipv4.netmask	The netmask for interface 1.1
network.1.1.ipv4.gateway	The gateway for interface 1.1
network.1.1.ipv6.address	The initial IPv6 management IP address on interface 1.1
network.1.1.ipv6.prefix	The prefix length for interface 1.1
network.1.1.ipv6.gateway	The gateway for interface 1.1

You can include both ipv4 and ipv6 settings in the same file. If you include ipv4 or ipv6 settings, all associated keys (address, netmask, and gateway) must be present.

The following example initially configures an IPv4 address for interface 1.1 and the appliance hostname.

```
network.hostname = isva-appliance.ibm.com
network.1.1.ipv4.address = 10.20.0.11
network.1.1.ipv4.netmask = 255.255.0.0
network.1.1.ipv4.gateway = 10.20.0.1
```

Related tasks

[“Creating a metadata image with the local management interface” on page 27](#)

You can create a metadata image that contains essential configuration data for the initial setup of an appliance with the local management interface. This image can later be used for the silent configuration of a new virtual appliance.

Chapter 3. Initial configuration

Several initial configuration tasks are required for your IBM Security Verify Access environment.

After you complete the [Getting started](#) tasks, including activating the product, continue with these steps:

1. [Manage application interfaces](#).
2. Configure your environment based on your needs:

Stand-alone Web Reverse Proxy

- a. [Configure the runtime environment](#).
- b. [Configure Web Reverse Proxy instances](#).

Member of a cluster of appliances

Primary master in a cluster:

- a. [Manage cluster configuration](#) and specify an appliance to be the primary master.
- b. [Configure the runtime environment](#).
- c. [Configure Web Reverse Proxy instances](#).

Member of a cluster:

- a. [Manage cluster configuration](#) and join the current appliance to the cluster.
- b. [Configure Web Reverse Proxy instances](#).

Front-end load balancer:

- a. [Configure the front-end load balancer](#).

Set up communication between appliances

Complete the following task if you have one appliance with Advanced Access Control activated and one without it: [Adding runtime listening interfaces](#).

Configure the Administrative User Registry:

If you want to enforce password policies for the administrative users, configure an external user registry that implements the desired password policies as the administrative user registry. See [“Configuring management authentication”](#) on page 90.

Proceed with additional configuration tasks as your environment requires.

Note: Do not connect the IBM Security Verify Access appliance to public telecommunications network interfaces. Further certification might be required by law before you make any such connections. Do not use the appliance in Public Services Networks. Contact IBM at [IBM Customer Support](#) for more information.

Chapter 4. Managing the appliance

The appliance provides three mechanisms by which it can be managed: the local management interface (LMI), the command-line interface (CLI), and web services interface.

Local management interface

The appliance offers a browser-based graphical user interface for local, single appliance management.

The following paragraphs are general notes about the usage of the local management interface (LMI). Examples of specific commands using the LMI are provided through the remainder of this document.

To log in to the LMI, type the IP address or host name of your appliance into your web browser. The following web browsers are supported:

- Windows
 - Google Chrome, version 27 or later
 - Microsoft Internet Explorer, version 11 or later
 - Mozilla Firefox, version 17 or later
- Linux®/AIX®/Solaris
 - Mozilla Firefox, version 17 or later

Use the default credentials to log in to the local management interface for the first time:

- **User Name:** admin
- **Password:** admin

After you log in for the first time, use the first-time configuration pages to change your password.

To log out of the local management interface, click **Logout**.

A customizable access banner can be presented on the local management interface login page. Use the **Login Screen Header** and **Login Screen Message** properties on the [Administrator Settings page](#) to set the access banner content.

Command-line interface

Access the command-line interface (CLI) of the appliance by using either an ssh session or the console.

The following example shows the transcript of using an ssh session to access the appliance:

```
usernameA@example.ibm.com>ssh -l admin webapp.vwasp.gc.au.ibm.com
admin@webapp.vwasp.gc.au.ibm.com's password:
Welcome to the IBM Security Verify Access appliance
Enter "help" for a list of available commands
webapp.vwasp.gc.au.ibm.com> isam
webapp.vwasp.gc.au.ibm.com:isam> help
Current mode commands:
aac          Work with the Advanced Access Control settings.
admin       Start an administration session which can be used to
            administer the Verify Access security policy.
ca          Work with the Policy server CA update operations.
cluster     Work with the Verify Access cluster.
dscadmin    Start an administration session which can be used to
            administer the Distributed Session Cache.
logs        Work with the Verify Access log files.
policy_db_dump  Validate and maintain the Security Verify Access policy
            database.
runtime_dump  Generate a core dump of the Verify Access runtime.
Global commands:
back        Return to the previous command mode.
exit        Log off from the appliance.
help        Display information for using the specified command.
reboot      Reboot the appliance.
```

```
shutdown      End system operation and turn off the power.
top           Return to the top level.
```

Tip: Use the **help** command to display usage notes about a specific command.

The following example shows the options available under the **lmi > accounts > locked** menu.

```
webapp.vwasp.gc.au.ibm.com:locked> help
Current mode commands:
list          List all of the locked accounts and the amount of time before each
              of the accounts will be automatically unlocked.
unlock_all    Unlock all of the locked accounts.
unlock <account> Unlock a specific account.
```

The following example shows the options available under the **logs** menu.

```
webapp.vwasp.gc.au.ibm.com:logs> help
Current mode commands:
archive       Archive the log files to a USB device.
delete        Delete the log files which have been rolled over by the system.
delete_trace  Delete the trace files (trace, stats, translog) from the system.
monitor       Monitor log files on the system.
ssl           Works with the Verify Access SSL log files.
```

The following example shows the options available under the **network** menu.

```
webapp.vwasp.gc.au.ibm.com:network> help
Current mode commands:
arp           Work with the ARP cache.
defgw         Work with the default gateway.
dns           Work with the appliance DNS settings.
hostname      Work with the appliance host name.
interfaces    Work with interface settings.
routes        Work with the static routes.
```

The following example shows the options available under the **routes** menu.

```
webapp.vwasp.gc.au.ibm.com:routes> help
Current mode commands:
add           Add a static route.
delete        Delete a static route.
edit          Edit a static route.
reset         Reset all the routing tables.
show          Show the static routes including both Active and Configured.
```

The usage of the **policy_db_dump** command is as follows:

```
policy_db_dump {-f <db_name>} {-l [1|2]} {-g} {-n} {-q} {-s} {-r}
{-d <find-entry-name> [-c <replace-entry-name>[:<hostname>[:<principal>]]}
-f <db_name> : Specifies the name of the policy database. This argument is optional
              if there is only a single Verify Access domain.
-l [1|2] :    The validation check level (2 is the default).
-g :         Display the glossary information only.
-n :         Display the object names only.
-q :         Display the sequence number of the policy database.
-s :         Display statistical information from the policy database.
-r :         Validate and repair the policy database. The policy server will be
              restarted as a result of this command.
-d :         Locate an entry in the database. If the -c flag is also specified the
              located entry is replaced with the new entry, otherwise the located
              entry is deleted from the database. The policy server will be restarted
              as a result of this command.
-c :         Replace the located entry in the database. This flag can only be used
              in conjunction with the -d flag. The policy server will be restarted
              as a result of this command.
```

The following example shows the options available under the **aac** menu.

```
webapp.vwasp.gc.au.ibm.com:aac> help
Current mode commands:
config        Start a session which can be used to configure a Web Reverse
              Proxy instance so that it can act as a point of contact for
              Advanced Access Control.
restart       Restart the Advanced Access Control runtime.
```

```
unconfig      Start a session which can be used to unconfigure a Web Reverse
              Proxy instance so that it can no longer act as a point of
              contact for Advanced Access Control.
```

The following example shows the options available under the **tools** menu:

```
webapp.vwasp.gc.au.ibm.com:tools> help
Current mode commands:
connect       Test network connection to a certain port on a specified host.
connections   Display the network connections for the appliance.
curl          Test the connection to a particular Web server using curl.
database      Get the connections currently open to the database.
nslookup      Query internet domain name servers.
ping          Send an ICMP ECHO_REQUEST to network hosts.
traceroute    Trace a packet from a computer to a remote destination, showing
              how many hops the packet required to reach the destination and
              how long each hop took.
telnet        Connect to telnet server.
session       Test network sessions with TCP or SSL.
```

The following example shows the options available under the **support** menu:

```
webapp.vwasp.gc.au.ibm.com:support> help
Current mode commands:
create        Create a support information file.
delete        Delete a support information file.
download      Download a support information file to a USB flash drive.
get_comment   View the comment associated with a support information file.
list          List the support information files.
list_categories List the categories registered for the support information file.
list_instances List the instances for a specific registered category.
purge         Purge the support files from the hard drive.
set_comment   Replace the comment associated with a support information file.
```

Note: The **purge** command deletes all core files, crashmap files, and support files from the `/var/support/` directory.

The following example shows the options available under the **pending_changes** menu:

```
webapp.vwasp.gc.au.ibm.com:pending_changes> help
Current mode commands:
discard       Discard the pending changes for a particular user or all users.
list          List all users who have outstanding pending changes.
```

The following example shows the options available under the **diagnostics** menu:

```
webapp.vwasp.gc.au.ibm.com:diagnostics> help
Current mode commands:
java_dump     Generate {heap|system|thread} java dump for
              {default|runtime} profile.
kill          Kill the specified process. This command will wait for the
              process to be terminated before returning.
list          List the contents of the local filesystem.
monitor       Generate monitoring output.
monitor_list  Print a list of all available monitor items.
ps            List the processes which are running on the system.
```

The method to access the console differs between the hardware appliance and the virtual appliance:

- For the hardware appliance, a serial console device must be used. For more information about attaching a serial console device to the hardware, see [“Connecting a serial console to the appliance”](#) on page 5.
- For the virtual appliance, you can access the console by using the appropriate VMWare software.

For example, VMWare vSphere Client.

Note: The CLI contains only a subset of the functions available from the local management interface. The following list gives a high-level overview of the functions available from the command-line interface. To see a list of the options for these commands, type the command name followed by **-help**.

diagnostics

Work with the IBM Security Verify Access diagnostics.

firmware

Work with firmware images.

fixpacks

Work with fix packs.

hardware

Work with the hardware settings.

isam

Work with the IBM Security Verify Access settings.

license

Work with licenses.

lmi

Work with the local management interface.

lmt

Work with the license metric tool.

management

Work with management settings.

network

Work with network settings.

pending_changes

Work with the IBM Security Verify Access pending changes.

snapshots

Work with policy snapshot files.

support

Work with support information files.

tools

Work with network diagnostic tools.

updates

Work with firmware and security updates.

You can also use a web service call to run most CLI commands. The web service URL is `https:<appliance>/core/cli`. For details about the usage of this web service, see the REST API documentation.

Note: The following CLI commands cannot be run via the web service:

- **isam > admin**
- **isam > dscadmin**
- **isam > logs > monitor**
- **isam > thales > rocs**
- **isam > thales > hsconfig**
- **isam > thales > cknfastrc**
- **isam > thales > nfdiag**
- **isam > thales > ckcheckinst**
- **hardware > ipmitool**
- **management > set_password**

A customizable access banner can be presented on the command line interface. Use the **Login Screen Header** and **Login Screen Message** properties on the [Administrator Settings](#) page to set the access banner content.

Web service

The appliance can also be managed by sending RESTful web service requests to the appliance.

Only one user can remain logged in to the appliance at the same time. Each web service request automatically displaces any existing sessions.

The following paragraphs are general notes about the usage of the web service interface. The content and format of these web service requests are explained through the remainder of this document.

Required header for calling a web service

All web service requests must include these two headers.

Accept:application/json

The accept header must be present and the value must be set to `application/json`. If the header is missing, or set as a different value, the web service request fails.

BA header

Each request must contain a BA header with a valid user name and password. If this header is missing, the request fails.

The following example is the valid request format for retrieving the list of reverse proxy instances by using curl.

```
curl -k -H "Accept:application/json" --user username:password
https://{appliance_hostname}/reverseproxy
```

Note: The previous list contains only two headers that are mandatory for all web service requests. It is not an extensive list of headers that are required for all request actions. The previous example shows a curl GET request on a resource URI. This request requires only the two mandatory headers that are listed. Other HTTP methods, such as POST or PUT, require more headers. The following example is a valid request for starting a reverse proxy instance called `inst1` using curl:

```
curl -k -H "Accept:application/json" -H "Content-type:application/json"
--user username:password --data-binary '{ 'operation':'start' }'
-X PUT https://{appliance_hostname}/reverseproxy/inst1
```

Notice the additional required header **Content-type** for the PUT operation.

Other HTTP clients, such as Java, might require more headers. For required headers for RESTful web services, check the HTTP client documentation.

Web service responses

The response to a web service call is composed of two components: HTTP response code and JSON message.

The response to a successful web service request includes a 200 status code, and JSON data that contains context-specific information about the request processing. The response to an unsuccessful web service request includes an HTTP error response code, and JSON data that contains the error message.

HTTP response codes

Code	Description
200	Success.
400	There is a problem with the request. The JSON message describes the problem.
404	The resource that is specified in the request does not exist. The JSON message indicates which resource.

Code	Description
500	An internal error was encountered while the request is processed. The JSON message indicates the problem.

JSON error response format

```
{"message": "The error message"}
```

Configuration changes commit process

The LMI uses a two-stage commit process when you make changes to the appliance.

Stage 1

Changes are made by using the LMI and saved to a staging area.

Stage 2

The user explicitly deploys the changes into production.

Multiple changes can exist in a pending state at the same time. They are committed or rolled back together when a user deploys or rolls back these changes.

Pending changes are managed on a per user identity basis. This means that changes made by one user identity will not be visible to another user identity until the changes are deployed.

Note: As there is no validation or merging of changes that are made by different user identities to the same component, changes that are made by one user can potentially overwrite changes that are made by another user.

Any changes that affect running reverse proxy instances require a restart of the effected instances before the changes can take effect.

Certain appliance updates require either the appliance or the web server to be restarted before the changes can take effect. When one or more of these updates are made alongside other reverse proxy updates, an additional step is required to deploy the reverse proxy updates. You must:

1. Deploy all updates.
2. Restart the appliance or the web server.
3. Deploy all remaining updates.

If there are conflicts between the pending changes and the production files, then all pending changes are automatically rolled back and the production files remain unchanged.

Web service

Deploy the pending configuration changes

URL

```
https://{appliance_hostname}/isam/pending_changes/deploy
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON error response where applicable.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/deploy
```

Response:

```
200 ok
```

Roll back the pending configuration changes

URL

```
https://{appliance_hostname}/isam/pending_changes/forget
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON error response where applicable.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/forget
```

Response:

```
200 ok
```

Retrieve the number of outstanding changes

URL

```
https://{appliance_hostname}/isam/pending_changes/count
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON data that represents the number of pending changes.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/count
```

Response:

```
{"count": 3}
```

Retrieve the list of outstanding changes

URL

```
https://{appliance_hostname}/isam/pending_changes
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON data that represents the list of pending changes.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes
```

Response:

```
200 ok

[{"id": 0,
 "policy": "SSL Certificates",
 "user": "admin",
 "date": "2012-11-05T11:22:20+10:00"
}]
```

Local management interface

When there are pending changes, a warning message is displayed at the top of the main pane. To deploy or roll back the pending changes:

1. Click the **Click here to review the changes or apply them to the system** link within the warning message.
2. In the **Deploy Pending Changes** page:
 - To view the details of changes that are made to a particular module, click the link to that module.
 - To deploy the changes, click **Deploy**.
 - To abandon the changes, click **Roll Back**.
 - To close the pop-up page without any actions against the changes, click **Cancel**.

Chapter 5. IBM Security Verify Access Appliance Dashboard

The appliance provides a series of dashboard widgets in its local management interface. You can use these widgets to view commonly used system information.

These widgets are displayed right after you log in. You can also access them by clicking **IBM Security Verify Access** on the menu bar.

Viewing system notifications

You can view warning information about potential problems with the Notification dashboard widget.

Procedure

1. From the dashboard, locate the Notification widget.

Warning messages about the following potential problems are displayed:

- Certificates that are due to expire.
- The disk space utilization has exceeded the warning threshold.
- The CPU utilization has exceeded the warning threshold.
- There are pending changes, which have not been deployed.
- The external configuration database is not accessible.
- The external runtime database is not accessible.
- Reverse proxy instances that are not currently running. (This notification is not available when the appliance is running in a Docker environment.)
- The database size has reached the warning threshold, which is 80% capacity. (This notification is not available when the appliance is running in a Docker environment.)
- The time is not synced to the NTP server. (This notification is not available when the appliance is running in a Docker environment.)

2. Take appropriate actions as required.

Viewing disk usage

You can view the disk space status and remaining disk life information with the Disk Usage dashboard widget.

About this task

This widget is not available when the appliance runs in a Docker environment.

Procedure

1. From the dashboard, locate the Disk Usage widget.

Disk Space Pie Chart

Information about used disk space and free disk space is visualized in the pie chart.

Consumed Disk Space

How much space (in GB) is already used.

Note: Most of the disk space is typically used by log files and trace files. To minimize the disk footprint, set the appliance to store log and trace files on a remote server. It is also a good practice to clear unused log and trace files on a periodic basis.

Free Disk Space

How much space (in GB) is free.

Total Disk Space

How much space in total (in GB) is available to the appliance.

Note: The disk space in a hardware appliance is limited by the capacity of the hard disk drive it carries.

2. *Optional:* Click **Refresh** to refresh the data.

Viewing IP addresses

You can view a categorized list of IP addresses that the appliance is listening on with the Interfaces dashboard widget.

About this task

This widget is not available when the appliance runs in a Docker environment.

Procedure

1. From the dashboard, locate the Interfaces widget.

The IP addresses of all enabled and configured interfaces are displayed, along with the virtual IP addresses that are managed by the front-end load balancer.

Management IPs

A list of IP addresses of the management interfaces that are enabled and configured.

Application IPs

A list of IP addresses of the application interfaces that are enabled and configured.

Load Balancer IPs

A list of IP addresses of the load balancer services.

2. *Optional:* Click **Refresh** to refresh the data.

Viewing certificate expiry

You can view certificate details with the Certificate Expiry widget.

Procedure

1. From the dashboard, locate the Certificate Expiry widget.

Details about the certificates are displayed.

Certificate Label

Label of the certificate.

Expiration

The date on which the certificate expires.

Type

Type of the certificate.

Key Database

Name of the key database that the certificate belongs to.

2. *Optional:* Click **Refresh** to refresh the data.

Viewing partition information

You can view information about the active and backup partitions with the Partition Information widget.

About this task

This widget is not available when the appliance runs in a Docker environment.

Procedure

1. From the dashboard, locate the Partition Information widget.
Details about the active and backup partition are displayed.

Firmware Version

Version information of the appliance firmware

Installation Date

Date on which the appliance firmware was installed

Installation Type

Type of the appliance firmware installation

Last Boot

Time when the appliance was last booted

2. *Optional:* Click **Firmware Settings** to go the page to modify settings of the firmware.

Viewing network traffic

You can view network traffic for the past hour with the Network Traffic widget.

About this task

This widget is not available when the appliance runs in a Docker environment.

Procedure

1. From the dashboard, locate the Network Traffic widget.
The **In** and **Out** traffic details for the past hour are displayed.
2. *Optional:* Click an interface name to display the details for a specific interface.

Viewing the status of the appliance in Docker

You can view the status of the appliance that is running in a Docker environment with the Docker dashboard widget.

About this task

This widget is only available when the appliance runs in a Docker environment.

Procedure

1. From the dashboard, locate the Docker widget.
 - Deployment Model**
Indicates that the appliance is running in a Docker container.
 - Version**
The firmware version of the appliance.
 - Configuration Database**
The status of the configuration database configuration.

Runtime Database

The status of the runtime database configuration.

User Registry

The type of user registry that has been configured (local or remote LDAP).

2. *Optional:* Click **Refresh** to refresh the data.

Configuring the dashboard

You can add and arrange widgets on the dashboard to monitor traffic, events, and system health in a summary view.

About this task

The appliance includes a dashboard view for a summary of your network status. You can select and arrange the information displayed on the dashboard to meet your needs.

Procedure

1. Click **IBM Security Verify Access**.
2. To rearrange the placement of the widgets, click the banner of a widget and drag it to where you want it.

Note: Widgets snap to a grid layout on the dashboard and are automatically arranged when you move one widget to the location of another.

Chapter 6. Monitoring

You can monitor the health and statistics of the appliance.

Viewing the event log

System events are logged when the system settings are changed and when problems occur with the system. Use the Event Log management page to view system events.

Procedure

Click **Monitor** > **Logs** > **Event Log**.

The system events displayed. You can:

- Click **Pause Live Streaming** to stop the live updating of the event log.
- Click **Start Live Streaming** to resume live updating of the event log.
- Click **Export** to download the event log file.

Notes:

- a. In the exported event log file, the time occurred (occurred) field shows the seconds since Epoch (00:00:00 Universal time, 1 January 1970).
- b. When you use the table filter on the **Priority** field, the values that can be filtered are in English only (low, medium, and high). This behavior is expected on all language versions of the appliance.

Forwarding logs to a remote syslog server

Configure the appliance to forward the contents of specific log files to a remote syslog server.

About this task

The preferred logging approach for the appliance is to send the logs to an external server. This approach can also meet certain compliance requirements.

When the remote syslog forwarding capability is enabled, it monitors local log files and forwards log entries from specific log files to a remote syslog server when new log entries are written in the local log files.

Note:

- Each line in the appliance standard log file is treated as a separate remote syslog message.
- All messages from a single log file are sent to the remote syslog server using the same facility and severity, as specified in the configuration.
- The `rsyslog` forwarding mechanism implements LF based framing.

Procedure

1. Click **Monitor** > **Logs** > **Remote Syslog Forwarding**.
2. Configure the remote syslog server settings as needed.

Adding a remote syslog server definition

- a. Click **Add**.
- b. Specify the details for the remote syslog server.

Server

The IP address or hostname of the remote syslog server to which messages are to be forwarded.

Port

The port on which the remote syslog server is listening for requests.

Debug

If selected, additional debug information will be included in the log file for the remote syslog forwarder process. The log file can be accessed from the `rsyslog_forwarder` directory of the [Viewing application log files](#) page.

Protocol

The protocol which will be used to communicate with the remote syslog server.

Format

The format of syslog messages which are forwarded to the remote syslog server.

c. Click **Save**.

Specifying the log sources for a remote log server

- a. Select the remote syslog server to send logs to.
- b. Click **Sources**.
- c. Click **Add** to add a log source.
- d. Specify the details for the log source and then click **OK**.

Name

Name of the log source.

Instance Name

Name of the instance that the source log file belongs to. This field is available only if **WebSEAL** or **Azn_Server** is selected in the **Name** field.

Log file

Name of the source log file. This field is available only if **WebSEAL** or **Azn_Server** is selected in the **Name** field.

Tag

The tag to add to the sent log entries.

Facility

The facility with which to send the log entries to the remote server. All messages will be sent with the specified facility code. The available codes can be found at: <https://en.wikipedia.org/wiki/Syslog#Facility>

Severity

The severity of the sent log entries. All messages will be sent with the specified severity level.

Note: The values are not saved on the server side until you click **Save** in Step f.

- e. If you want to add multiple log sources, repeat the previous two steps
- f. Click **Save**.

Viewing memory statistics

View the memory graph to see the memory utilization of the appliance.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. Click **Monitor** > **System Graphs** > **Memory**.

2. Select a **Date Range**:

Option	Description
30 Minutes	Displays data points for every minute during the last 30 minutes.
1 Hour	Displays data points for every minute during the last 60 minutes.
12 Hours	Displays data points for every minute during the last 12 hours.
1 Day	Displays data points for every minute during the last 24 hours.
3 Days	Displays data points for every 5 minutes during the last three days. Each data point is an average of the activity that occurred in that hour.
7 Days	Displays data points every 20 minutes during the last seven days. Each data point is an average of the activity that occurred in that hour.
30 Days	Displays data points for every hour during the last 30 days. Each data point is an average of the activity that occurred in that hour.

3. In the Legend box, select **Memory Used** to review total memory utilization.

Viewing CPU utilization

View the CPU graph to see the CPU utilization of the appliance.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. Click **Monitor > System Graphs > CPU**.

2. Select a **Date Range**:

Option	Description
30 Minutes	Displays data points for every minute during the last 30 minutes.
1 Hour	Displays data points for every minute during the last 60 minutes.
12 Hours	Displays data points for every minute during the last 12 hours.
1 Day	Displays data points for every minute during the last 24 hours.
3 Days	Displays data points for every 5 minutes during the last three days. Each data point is an average of the activity that occurred in that hour.
7 Days	Displays data points every 20 minutes during the last seven days. Each data point is an average of the activity that occurred in that hour.
30 Days	Displays data points for every hour during the last 30 days. Each data point is an average of the activity that occurred in that hour.

3. In the Legend box, select the CPU utilization data that you want to review:

- User
- System
- Idle

Viewing storage utilization

View the storage graph to see the percentage of disk space that is used by the boot and root partitions of the appliance.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. Click **Monitor > System Graphs > Storage**.
2. Select a **Date Range**:

Option	Description
30 Minutes	Displays data points for every minute during the last 30 minutes.
1 Hour	Displays data points for every minute during the last 60 minutes.
12 Hours	Displays data points for every minute during the last 12 hours.
1 Day	Displays data points for every minute during the last 24 hours.
3 Days	Displays data points for every 5 minutes during the last three days. Each data point is an average of the activity that occurred in that hour.
7 Days	Displays data points every 20 minutes during the last seven days. Each data point is an average of the activity that occurred in that hour.
30 Days	Displays data points for every hour during the last 30 days. Each data point is an average of the activity that occurred in that hour.

3. In the Legend box, select which partitions you want to review:

Boot

The boot partition.

Root

The base file system, where the system user is root.

Viewing application interface statistics

To view the bandwidth and frames that are being used on your application interfaces, use the Application Interface Statistics management page.

About this task

This page is not available in the LMI when the appliance runs in a docker environment.

Procedure

1. From the top menu, select **Monitor > Network Graphs > Application Interface Statistics**.
2. In the **Date Range** field, select the period to display the statistics for.

Option	Description
30 Minutes	Displays data for every minute interval in the last 30 minutes.
1 Hour	Displays data for every minute interval in the last 60 minutes.
12 Hours	Displays data for every minute interval in the last 12 hours.
1 Day	Displays data for every 20-minute interval in one day.

Option	Description
3 Days	Displays data for every 20-minute interval during the last three days.
7 Days	Displays data for every 20-minute interval during the last seven days.
30 Days	Displays data for every day during the last 30 days.

Viewing application log files

Use the Application Log Files management page to view and download log files that are produced by IBM Security Verify Access.

Procedure

1. From the top menu, select **Monitor > Application Log Files**.

The displayed directories contain the application log files that can be viewed and downloaded:

- **access_control**: Contains log files specific to the Advanced Access Control offering. It contains subdirectories for different categories of log files, such as **auditing**, **isamcfg**, and **runtime**.
- **cluster**: Contains logs files for the cluster manager.
- **management_ui**: Contains log files for the management interface.
- **federation**: Contains logs files specific to the Federation offering.

By default, the log files are displayed in a tree view.

2. Optional: Click **Details View** to manage the log files using a more detailed view. This view shows the path, file size, and last modified time of each log file. You can also order the files by name, path, file size, or last modified time.
3. Optional: Click **Refresh** to get the most up-to-date data.
4. You can then view or download the displayed log files.

To view the log file

- a. Select the file of interest.
- b. Click **View**. The content of the log file is displayed. By default, the last 100 lines of a log file are displayed if the file is longer than 100 lines. You can define the number of lines to display by entering the number in the **Number of lines to view** field and then click **Reload**. Alternatively, you can provide a value in the **Starting from line** field to define the start of the lines. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

- c. *Optional:* Click **Export** to download the log file.

To download the log file

- a. Select the file of interest.
- b. Click **Export** to save the file to your local drive.
- c. Confirm the save operation in the browser window that pops up.

To clear or empty a log file

- a. Select the file or files of interest.
- b. Click **Clear** to clear the contents of the file.
- c. In the confirmation window, click **Yes** to confirm the clear operation.

To delete a log file

- a. Select the file or files of interest.

Note: It is the administrator's responsibility to make sure that the log file to be deleted is not in use by the system.

- b. Click **Delete** to remove the log file.
- c. In the confirmation window, click **Yes** to confirm the deletion.

Managing tracing specification

Setting trace for Oracle components "oracle.*" results in the underlying Oracle JDBC jar file being changed to a debugging jar file. This might have adverse effects on performance and as such Oracle tracing should only be enabled for debugging purposes and disabled once complete.

About this task

Procedure

1. Select the **Runtime Tracing** link from the top of this page. You can also access this panel from the top menu by selecting **Monitor > Logs > Runtime Tracing**.
2. Use one of the following ways to edit the trace level of a component.
 - Select the component name from the **Component** list. Select the ideal trace level for this component from the **Trace Level** list. Then, click **Add**. Repeat this process to modify trace levels for other components if needed. To clear all of the tracing levels, click **Clear**.

To log all events, select ALL as the trace level.

Note: This setting increases the amount of data in logs, so use this level when necessary.

```
com.tivoli.am.fim.authsvc.*
com.tivoli.am.fim.trustserver.sts.modules.*
```

Table 3. Valid trace levels. The following table contains the valid trace levels.

Level	Significance
ALL	All events are logged. If you create custom levels, ALL includes those levels and can provide a more detailed trace than FINEST.
FINEST	Detailed trace information that includes all of the details that are necessary to debug problems.
FINER	Detailed trace information.
FINE	General trace information that includes methods entry, exit, and return values.
DETAIL	General information that details sub task progress.
CONFIG	Configuration change or status.
INFO	General information that outlines the overall task progress.
AUDIT	Significant event that affects the server state or resources.

Table 3. Valid trace levels. The following table contains the valid trace levels. (continued)

Level	Significance
WARNING	Potential error or impending error. This level can also indicate a progressive failure. For example: the potential leaking of resources
SEVERE	The task cannot continue, but component, application, and server can still function. This level can also indicate an impending unrecoverable error.
FATAL	The task cannot continue, and component, application, and server cannot function.
OFF	Logging is turned off.

- Enter the name and value of the trace component in the **Trace Specification** field. To modify multiple components, separate two strings with a colon (:). Here is an example.

```
com.x.y.*=WARNING:com.a.b.*=WARNING:com.ibm.isva.*=INFO
```

3. Click **Save**.

Chapter 7. System

Information about configuring Security, Network, and System settings of your appliance.

Updates and licensing

Information about managing updates and licensing on your appliance.

Viewing the update and licensing overview

The Overview page displays current information about the appliance firmware, intrusion prevention content, IP reputation database, update servers, and licenses.

Procedure

1. Click **System > Updates and Licensing > Overview**.
2. View the updates and licensing information. Click the links on the page to make a specific update.

Installing updates

Install firmware and intrusion prevention updates to improve the appliance and the network protection that is provided by the appliance.

About this task

Important: After you install firmware updates, you must restart the appliance.

Firmware updates contain new program files, fixes or patches, enhancements, and online help.

Intrusion prevention updates contain the most recent security content that is provided by IBM X-Force research and development team.

Note: Any X-Force content update takes effect only after you restart the related Web Reverse Proxy instances.

Procedure

1. Click **System > Updates and Licensing > Available Updates**.
2. On the **Available Updates** page, use one or more of the following commands:

Option	Description
Upload	To manually add an update, click Upload . In the New Update window, click Select Update , browse to the update file, click Open , and then click Submit . Note: You can install the update after you manually add it.
Refresh	To check for updates, click Refresh .
Install	To install an update, select the update, and then click Install .

3. Restart the Web Reverse Proxy instances so that the X-Force content update can take effect.

Configuring the update schedule

Configure the update schedule to receive X-Force content updates daily, weekly, or according to a specified interval of time.

About this task

A 15-minute buffer is applied to update times so that update servers do not become overburdened. Updates are downloaded up to 15 minutes before or after the time you specify.

Procedure

1. Click **System > Updates and Licensing > Scheduled Security Updates**.
2. In the Update Schedule pane, select **Auto Update** to receive X-Force content updates.
3. Use one of the following methods to schedule updates:
 - To receive updates on a daily basis, select **Daily or Weekly**, select **Every Day** from the first list, and then select a time from the second list.
 - To receive updates on a weekly basis, select **Daily or Weekly**, select the day of the week you would like to receive updates on, and then select a time from the second list.
 - To receive updates on a schedule that ranges from 1 hour to 24 hours, select **Specified Interval**, and then select the update interval in minutes.
Range: 60 - 1440 minutes
4. Click **Save**.

Configuring update server settings

Configure your appliance to download update files from an update server.

About this task


You can configure multiple, ordered servers for failover.

Note: You cannot delete the IBM ISS Default License and Update Server. You can disable it.

Procedure

1. Click **System > Updates and Licensing > Update Servers**.
2. In the Update Servers pane, take one of the following actions:
 - To add an update server, click **New**. The **Add Server** window is displayed.
 - To edit an update server, select the server, and then click **Edit**. The **Edit Server** window is displayed.
 - To delete an update server, select the server, and then click **Delete**.
3. When you add or edit an update server, configure the following options on the General tab:

Option	Description
Order	Defines the order in which update servers are queried for appliance software updates. The appliance uses the next server on the list when a server takes more than 24 hours to respond.
Enable	Enables the update server so that it can be used by the appliance.
Name	A name that describes the update server.
Server Address	The IP address or DNS name of the update server.

Option	Description
Port	The port number that the appliance uses to communicate with the update server. Tip: The port number for the IBM ISS Download Center is 443. The default port for internal update servers is 3994.
Trust Level	<p>Defines how the appliance is authenticated with the update server.</p> <p>Explicit (user-defined) The appliance uses the local certificate that is pasted into the Certificate box to authenticate the connection to the update server. The certificate must be Base64 PEM-encoded data.</p> <p>Explicit trust is the most secure trust level. Explicit trust certificates must be Base64 PEM-encoded data.</p> <p>Explicit (xpu.iss.net) The appliance uses the local certificate for the IBM ISS update server to authenticate the connection to the update server. The IBM ISS update server certificate is installed on the appliance by default. The certificate is Base64 PEM-encoded data.</p> <p>Explicit trust is the most secure trust level. Explicit trust certificates must be Base64 PEM-encoded data.</p> <p>First Time Trust If a certificate is not on the appliance, the appliance downloads a certificate from the server when it connects to the server for the first time.</p> <p>First Time Trust is more secure than Trust All and less secure than Explicit Trust.</p> <p>Note: After the appliance downloads the certificate, it reverts to explicit-trust functionality.</p> <p>Trust All The appliance trusts the update server, and does not use SSL certificates for authentication.</p> <p>Trust all trust is the least secure trust level.</p> <p> Attention: The Trust All trust level presents a security risk because the internal update server can be spoofed and redirected to a fake server.</p>

4. Optional: If you use a proxy server, configure the following settings on the Proxy Settings tab:

Option	Description
Use Proxy	Enables the appliance to use a proxy server for update servers.
Server Address	The IP address or DNS name of the proxy server. Note: The Server Address field is displayed when you select the Use Proxy check box.
Port	The port number that the proxy server uses to communicate with the update server. Note: The Port field is displayed when you select the Use Proxy check box.
Use Authentication	Enables the appliance to authenticate to a proxy server.
User Name	User name that is required for authenticating to the proxy server. Note: The User Name field is displayed when you select the Use Authentication check box.

Option	Description
Password	Password that is required for authenticating to the proxy server. Note: The Password field is displayed when you select the Use Authentication check box.

5. Click **Submit**.

Viewing update history

View the update history to see which firmware and security content updates are downloaded, installed, and rolled back on the appliance.

About this task

After you install an update, the update package is deleted from the appliance.

Procedure

1. Click **System > Updates and Licensing > Update History**.
2. To refresh the page, click **Refresh**.

Installing a fix pack

Install a fix pack when IBM Customer Support instructs you to do so.

Before you begin

The appliance does not automatically create a backup copy of a partition when you apply a fix pack to it. If you want to back up your partition before you apply the fix pack, then you must do it manually.

Restriction: You cannot uninstall fix packs.

About this task

Fix packs are applied to the current partition. If a fix pack is installed on your appliance, you can view information about who installed it, comments, patch size, and the installation date.

Procedure

1. In the local management interface, go to **System > Updates and Licensing > Fix Packs**.
2. In the **Fix Packs** pane, click **New**.
3. In the **Add Fix Pack** window, click **Browse for fix pack:** to locate the fix pack file, and then click **Open**.
4. Click **Save Configuration** to install the fix pack.

Installing a license

You must install a current license file to receive updates to the appliance.

About this task

Contact your IBM representative to get a license registration number. You can download and register your license from the IBM Security Systems License Key Center at <https://ibmss.flexnetoperations.com>.

Procedure

1. Optional: If you are not configuring your appliance for the first time, click **System > Updates and Licensing > Licensing and Activation**.

2. On the **Licensing and Activation** page, click **Select License** and locate the license file that you want to install.
3. Select the license file that you want to install and then click **Open**.
4. Click **Save Configuration**.

Note: OCNID stands for Order Confirmation Number and ID.

Managing firmware settings

The appliance has two partitions with separate firmware on each partition. Partitions are swapped during firmware updates, so that you can roll back firmware updates.

About this task

Either partition can be active on the appliance. In the factory-installed state, partition 1 is active and contains the firmware version of the current released product. When you apply a firmware update, the update is installed on partition 2 and your policies and settings are copied from partition 1 to partition 2. The appliance restarts the system using partition 2, which is now the active partition.

Note: The appliance comes with identical firmware versions installed on both of the partitions so that you have a backup of the initial firmware configuration.

Tip: Avoid swapping partitions to restore configuration and policy settings. Use snapshots to back up and restore configuration and policy settings.

Procedure

1. Click **System > Updates and Licensing > Firmware Settings**.
2. On the **Firmware Settings** page, perform one or more of the following actions:

Option	Description
Edit	To edit the comment that is associated with a partition, select the partition and click Edit .
Create Backup	<p>Important: Create a backup of your firmware only when you are installing a fix pack that is provided by IBM Customer Support.</p> <p>Fix packs are installed on the active partition and you might not be able to uninstall the fix pack.</p> <p>Note: The backup process can take several minutes to complete.</p>
Set Active	Set a partition active when you want to use the firmware that is installed on that partition. For example, you might want to set a partition active to use firmware that does not contain a recently applied update or fix pack.

3. Click **Yes**.

If you set a partition active, the appliance restarts the system using the newly activated partition.

Managing trial settings

Use the **Trial** page to upload a trial certificate so that you can start your appliance trial.

About this task

The trial is activated by uploading a trial certificate. You can request a trial certificate by clicking the **Request a trial on the IBM Marketplace** link on the **Trial** page and following the instructions on the website. At the end of the trial request process, you will be able to download the trial license. You can then upload this trial license to your appliance to start using the appliance on a trial basis.

All offerings will be activated on a trial basis. After the trial expires, the offerings will be deactivated.

The trial can be reverted by uploading a special revocation trial certificate. When the revocation trial certificate is uploaded, the trial offerings will be deactivated. If you want any of the offerings to remain active, you must upload the activation key before reverting the trial.

When a trial period is activated, the remaining time for the trial is displayed in the title area of the LMI.

After the trial period expires, the runtime services of the appliance (for example, WebSEAL) will be disabled. If the administrator attempts to access the LMI after the appliance is disabled, the administrator will be automatically redirected to this **Trial** page.

Procedure

1. In the local management interface, go to **System > Updates and Licensing > Trial**.
2. Click **Import**.
3. Browse to the certificate and confirm the import operation.

Installing an extension

Install an IBM Security Verify Access extension in the environment.

Before you begin

1. Download any third-party dependencies for the extension that you are installing from the vendor's website.
2. Download the corresponding extension support package file from [IBM Fix Central](#).

About this task

Extensions are applied to the current partition and persists after firmware upgrade. If an extension is installed on your appliance or docker environment, you can view information about the extension and the installation date.

The following extensions are available to be installed for IBM Security Verify Access:

DynaTrace AppMon Agent

The Dynatrace (DT) AppMon is an application monitoring tool that provides all performance metrics in real time and can detect and diagnose problems automatically. It can be used to monitor the Advanced Access Control or Federation liberty runtime profiles.

Note: Installation of the DynaTrace AppMon Agent extension requires the DynaTrace AppMon agent JAR (third-party dependency file) from the DynaTrace website.

Procedure

1. From the dashboard, click **System > Updates and Licensing > Extension**.
2. In the **Extensions** pane, click **New**.
3. Upload the extension support file and click **Next**.
4. On the next dialogue box, provide the configuration parameter details and upload the third-party dependency.
5. Click **Install**.

Network Settings

Information about configuring network interfaces and information about your appliance.

Configuring general networking settings

Set the host name of the appliance.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > General**.
2. Enter the host name.

Note: Changing the appliance host name causes the security device to reset the network connection. You must reconnect after the network connection is reset. This process does not interrupt traffic through the application interfaces.

3. Click **Save Configuration**.

Configuring DNS

Define the DNS settings for your interfaces.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > DNS**.
 - To set the DNS via DHCP of an interface:
 - a. Select **Auto**.
 - b. Select the interface from the list.
 - To use manual DNS settings:
 - a. Select **Manual**.
 - b. Define the following settings:
 - Primary DNS (mandatory)
 - Secondary DNS
 - Tertiary DNS
 - DNS Search Path
2. Click **Save Configuration**.

Configuring interfaces

Create or edit your management and application interfaces.

About this task

The appliance supports the use of virtual local area network (VLAN). A VLAN is a logical association of switch ports that are based on a set of rules or criteria, such as MAC addresses, protocols, network

address, or multicast address. This concept permits the LAN to be segmented again without requiring physical rearrangement.

The interfaces with names in the format of **1.x** are real interfaces, which correspond to the network adapters on your physical appliance or the adapters that are attached to your virtual appliance. The interfaces with names in the format of **1.x:<vlanid>** are virtual interfaces.

You can add or delete virtual interfaces, but you cannot delete real interfaces. When you add an interface, you are effectively adding a VLAN to a specific interface.

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Interfaces**.

All current management and application interfaces are displayed.

2. You can add or edit interfaces and addresses that are associated with an interface.

- To add an interface:
 - a. Click **New**.
 - b. On the **General Configuration** tab:
 - i) Select the type of interface to create.

Note: For interfaces of the type **Loopback**, DHCP and bonding options are not available.
 - ii) Enter a name for the interface.
 - iii) Select the **Enabled** checkbox if you want to enable this interface at the same time when it is created.
 - iv) Enter the virtual LAN ID for the interface.
 - v) Add notes about this interface in the **Comment** field.
 - c. Click **Save Configuration** to confirm the details of this interface.
- To modify the details of an interface:
 - a. Select the interface from the table.
 - b. Click **Edit**.
 - c. Modify the details as needed.
 - d. Click **Save Configuration** to confirm the modified details.
- To delete a virtual interface:
 - a. Select the interface from the table.
 - b. Click **Delete**.
 - c. Click **Yes** to confirm the operation.
- To add an IP address to an interface:
 - a. Select the interface.
 - b. Click **Edit**.
 - c. On the **IPv4 Settings** tab:
 - i) If you want to use DHCP to assign addresses, select **Auto**.
 - a) To make this interface a management interface, select the **Management Interface** checkbox. To make this interface an application interface, leave this checkbox unchecked.
 - b) Select the **Provides Default Route** if needed.
 - ii) If you want to use static addresses, select **Manual**.
 - a) Click **New** to add an address.

- b) Enter the static address in the **Address** field in the format of `<address>/<mask>`. Masks are supported in dot-decimal and CIDR notation, for example:

```
10.0.2.38/24
10.0.2.38/255.255.255.0
```

- c) To use this address for management purposes, select the **Management Address** checkbox. To use this address for application, leave this checkbox unchecked.
 - d) By default, the appliance performs validation to ensure that overlapping subnets do not span multiple interfaces. Such validation helps prevent networking issues in certain environments. If you want to disable this validation for your environment, select the **Override the Overlapping Subnet Validation** option.
 - e) Click **Save Configuration** to confirm the details.
- d. On the **IPv6 Settings** tab:
- i) If you want to use DHCP to assign addresses, select **Auto**.
 - a) To make this interface a management interface, select the **Management Interface** checkbox. To make this interface an application interface, leave this checkbox unchecked.
 - ii) If you want to use static addresses, select **Manual**.
 - a) Click **New** to add an address.
 - b) Enter the static address in the **Address** field in the format of `<address>/<mask>`. Masks must be given in CIDR notation, for example:

```
2001:db8::38/48
```
 - c) To use this address for management purposes, select the **Management Address** checkbox. To use this address for application, leave this checkbox unchecked.
 - d) Click **Save Configuration** to confirm the details.
- To modify an IP address that is associated with an interface:
 - a. Select the interface.
 - b. Click **Edit**.
 - c. On the **IPv4 Settings** and **IPv6 Settings** tabs, select the address to modify and then click **Edit**.
 - d. Modify the settings as needed.
 - e. In the **Edit address** window, click **Save Configuration** to close the window.
 - f. Click **Save Configuration** to confirm the interface details.
 - To delete an IP address that is associated with an interface:
 - a. Select the interface.
 - b. Click **Edit**.
 - c. On the **IPv4 Settings** and **IPv6 Settings** tabs, select the address to delete and then click **Delete**.
 - d. Click **Yes** to confirm the delete operation.
 - e. Click **Save Configuration** to confirm the interface details.

Appliance port usage

The following table lists the ports that the appliance listens on and provides a description of what the port is used for and what external entities use the port.

This table can be used to decide:

- The firewall rules that are used to allow or block port access to or from the appliance
- Which ports are reserved and must be avoided by administrator configurable ports

The appliance provides two types of interface groupings: administration interface and application interface. Typically ports are assigned to one or more IP addresses from one of these groups of interfaces. In some cases, ports can be assigned to all IP addresses from both groups by providing 0.0.0.0 as the IP address to use.

<i>Table 4. Ports used on the appliance (listen ports)</i>		
Appliance port	Appliance interface type	Description
22	Administration	This port serves two roles. <ol style="list-style-type: none"> 1. Provides remote access to the CLI for the admin user. 2. Cluster inter-node communication. Each node in a cluster must have access to all other cluster nodes' SSH ports.
80	Application (The port can be assigned to both application and administration interfaces by providing 0.0.0.0 as the IP address to use.)	This port is the typical default unsecured (non-SSL) port of the first configured Web Reverse Proxy instance. This port can be configured to a different value or disabled.
443	Application (The port can be assigned to both application and administration interfaces by providing 0.0.0.0 as the IP address to use.)	This port is the typical default secured (SSL) port of the first configured Web Reverse Proxy instance. This port can be configured to a different value or disabled.
443	Administration	This port is the Local Management Interface (LMI) secure port.
636	Administration	This port is reserved for remote SSL access to the embedded user registry. The port is only active on the primary master node of the cluster when the Security Verify Access runtime is configured to use the embedded user registry.

Table 4. Ports used on the appliance (listen ports) (continued)

Appliance port	Appliance interface type	Description
2020+7	Administration	<p>This port is used by the appliance DSC servers to replicate session data between cluster master nodes. Each master node must have access to the port of its adjacent node. The primary node is adjacent to the secondary node. The secondary node is adjacent to the tertiary node. The tertiary node is adjacent to the quaternary node.</p> <p>Note: The 2020+7 value assumes that the cluster First Port is set to its default value 2020. If the cluster First Port is configured to a value other than the default, this port value must be adjusted relative to the configured First Port value (configured First Port+7).</p>
7135	Administration	<p>The policy server listens on this port if it is running on the node. Any node that is running Web Reverse Proxy servers, authorization servers, the PD.jar API, pdadmin API, or pdadmin command requires access to this port. This port can be configured to a different value.</p>
7136	Application	<p>This port is the typical first authorization server port that can be accessed by the Java or C administration or authorization APIs. This port can be configured to a different value.</p>
7137	Administration	<p>This port is the typical first authorization server admin port, which must be accessible by the machine that is running the policy server. This port can be configured to a different value.</p>
7234	Administration	<p>The Web Reverse Proxy server listens on this port if it is running on the node. This port must be accessible from the node that is running the policy server. This value is the typical port that is used for the first Web Reverse Proxy on a node. This port can be configured to a different value.</p>

Note: Many services on the appliance can be configured to access external service ports such as LDAP, SQL, DNS, NTP Web Reverse Proxy junctions, OCSP, Kerberos, and syslog server ports. The routing that is configured on the appliance determines which outgoing interface is used to access them based on the external service's IP address.

Configuring aggregated network interfaces

Set up aggregated network interfaces for high availability, increased throughput, or both.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

This capability is commonly called *bonding* in Linux environments. Use it to place multiple real network interfaces behind a virtual network interface. This feature is useful for physical appliances only, and not for virtual appliances. For virtual appliances, you can use the hypervisor to set up the NIC bonding and present a single virtual interface to the virtual appliances.

The appliance supports the following bonding modes:

Mode	Name	Description
0	balance-rr	Round-robin policy: Transmits packets in sequential order from the first available slave through the last.
1	active-backup	Active-backup policy: Only 1 slave in the bond is active. A different slave becomes active if, and only if, the active slave fails.
2	balance-xor	XOR policy: Transmits based on the selected transmit hash policy.
3	broadcast	Broadcast policy: Transmits everything on all slave interfaces.
4	802.3ad	IEEE 802.3ad Dynamic link aggregation: Creates aggregation groups that share speed and duplex settings. Uses all slaves in the active aggregator according to the 802.3ad specification.
5	balance-tlb	Adaptive transmit load balancing: Channel bonding that does not require any special switch support.
6	balance-alb	Adaptive load balancing: Includes balance-tlb plus receive load balancing (rlb) for IPV4 traffic. It does not require any special switch support. The receive load balancing is achieved by ARP negotiation.

Configuration options for these bonding modes are available through the appliance advanced tuning parameters. If set, the parameters apply to all bonding interfaces. For more details, see [“Managing advanced tuning parameters”](#) on page 96.

The bonding (enslave) order of the slaves is not configurable. The network device that is configured as the primary bonding device uses its underlying physical device as the first bonded slave.

Note: Expect interruption to any existing network traffic on the involved interfaces when the configuration changes are committed.

Procedure

1. From the top menu, select **System > Network Settings > Interfaces**.

2. Edit the appliance interface to be replaced by the virtual bonding interface behind which the aggregation of interfaces is placed. The physical network interface that is normally represented by this configuration is the first interface aggregated behind the bonding virtual interface.
 - a) Select the interface and then click **Edit**.
 - b) For this interface, set the **Bonding Mode** to something other than **None** or **Slave**.
For example, **802.3ad**.
Note: Some bonding modes, such as **802.3ad**, require equivalent support from the network switch to which they are attached.
 - c) Set the IP addresses of the interface, if not already set. This interface configuration defines the IP address of the aggregation.
 - d) Save the configuration.
3. Edit each additional interface to be added to the aggregation. For each slave:
 - a) Set the **Bonding Mode** to **Slave**.
Note: If you have an existing bonding configuration on an interface, you must set the bonding configuration back to **None** and deploy the change before you can set the interface to be a slave. That is, the **Slave** option does not appear when you list the available modes on an interface with an existing bonding configuration. You must first clear the existing bonding configuration by setting the bonding mode to **None**. After deploying the change, you can see the **Slave** option in the list.
 - b) For the **Bonded To** field, select the initial interface that was configured in previous steps.
 - c) Save the configuration.
4. Commit the changes.

Configuring static routes

Besides configuring static routes, you can also use this page to set the default gateway.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

This task is only necessary for networks that contain an additional network segment between the user segment and the appliance.

Procedure

1. From the top menu, select **System > Network Settings > Static Routes**.
2. Select the route table to edit from the **Route Table For** field.
You can use these route tables to configure routes that are specific to requests destined for a particular local IP address. Use the **Default** table if specific local IP address control is not required.
3. Take one of the following actions:
 - Click **New** to create a route.
 - Select an existing route, and then click **Edit** to change the settings.
 - Select an existing route, and then click **Delete** to remove it.
4. Define the following information in each field:
 - Enabled
 - Destination
Note: To make this route the default gateway, enter `Default` in the **Destination** field.
 - Gateway
 - Metric
 - Interface

5. Click **Save Configuration**.

Multiple routing tables

You can configure a specific set of routes for each IP that is configured on the appliance. This setting can overcome a single point of failure that occurs from having a single interface and gateway for a particular subnet, or from having a single default gateway.

Interface-specific routes might also be required to solve some firewall conflicts. In an appliance that has multiple interfaces, the return path for a particular request might be different from the request path. In certain firewall configurations, this situation is seen as a spoofing attack and the packet is discarded.

For example, if the appliance has an IP of 172.16.197.11/24 configured on Interface 1.2 and a gateway at 172.16.197.2, then select the table for IP 172.16.197.11 and add the following two static routes:

- 172.16.197.0/24 Interface 1.2
- Default via 172.16.197.2 Interface 1.2

If a set of static routes is not provided for a particular IP's table, or the static routes in the IP's table do not result in a route for the IP, then the "Default" static route table is applied.

If the ability to define different routes for different destination IP address is not required, then place all required static routes under the "Default" static route table. This is also where migrated static routes from prior releases that do not provide this feature are placed.

Testing the connection

Test a TCP or SSL connection.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Test Connection**.
2. You can test a TCP or SSL connection.

Testing a TCP connection

- a. Select the **TCP** option.
- b. Enter the server, port, and optionally the timeout value.
- c. Click **Test Connection**. Any message that is generated as output of the connection test is displayed at the bottom of the page.

Testing an SSL connection

- a. Select the **SSL** option.
- b. Enter the server, port, and optionally the timeout value.
- c. Select **Show SSL Advanced Parameters** to display additional SSL parameters that can be specified.
- d. Define any SSL additional parameters as needed.

Parameter	Description
keyfile	The keystore to use on the connection request.
label	The certificate to use on the connection request.

Table 6. SSL additional parameters (continued)

Parameter	Description
reconnect	Reconnect to the same server multiple times to ensure that session caching is working.
pause	Pause for 1 second between each read and write call.
showcerts	Show the entire certificate chain.
debug	Print more verbose debugging information.
msg	Show all protocol messages with hex dump.
nbio_test	Test non blocking IO.
state	Print the SSL session states.
nbio	Turn on non blocking IO.
crlf	Translate a line feed into CR+LF.
quiet	Inhibit the printing of session and certificate information.
tlsextddebug	Print out a hex dump of any TLS extensions received from the server.
status	Send a certificate status request to the server.
ssl2	Try to connect using SSLv2.
ssl3	Try to connect using SSLv3.
tls1_2	Try to connect using TLSv1.2.
tls1_1	Try to connect using TLSv1.1.
tls1	Try to connect using TLSv1.
dtls1	Try to connect using DTLSv1.
no_ssl2	Disable the use of SSLv2 during connect.
no_ssl3	Disable the use of SSLv3 during connect.
no_tls1_2	Disable the use of TLSv1.2 during connect.
no_tls1_1	Disable the use of TLSv1.1 during connect.
no_tls1	Disable the use of TLSv1 during connect.

- e. Click **Test Connection**. Any message that is generated as output of the connection test is displayed at the bottom of the page.

Managing hosts file

To manage hosts file with the local management interface, use the Hosts File management page.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Hosts File**.
All current host records with their IP address and host names are displayed.

2. You can then work with host records and host names.

- **Add a host record**

- a. Select the root level **Host Records** entry or do not select any entries.
- b. Click **New**.
- c. On the Create Host record page, provide IP address and host name of the host record to add.
- d. Click **Save**.

- **Add a host name to a host record**

- a. Select the host record entry to add the host name to.
- b. Click **New**.
- c. On the **Add Hostname to Host Record** page, enter the host name to add.
- d. Click **Save**.

- **Remove a host record**

- a. Select the host record entry to delete.
- b. Click **Delete**.
- c. On the confirmation page, click **Yes** to confirm the deletion.

- **Remove a host name from a host record**

- a. Select host name entry to delete.
- b. Click **Delete**.
- c. On the confirmation page, click **Yes** to confirm the deletion.

Note: If the removed host name is the only associated host name for the IP address, then the entire host record (the IP address and host name) is removed.

Managing the shared volume

In a Docker environment, you can manage the files that are stored on the shared volume (`/var/shared`) with the **Shared Volume** management page.

About this task

This page is only available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Shared Volume**.

All contents of the shared volume are displayed under the relevant directories.

fixpacks

Fix pack files.

snapshots

Snapshot files.

support

Support files.

2. You can upload, download, rename, or delete these files as needed.

3. Optional: Click **Refresh** to get the most up-to-date data.

Managing packet tracing

To manage packet tracing with the local management interface, use the Packet Tracing management page.

About this task

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Packet Tracing**.

The status of packet tracing is displayed.

Note: The top grid shows the status of the packet tracing along with the details of the current PCAP tracing file only. The bottom grid shows the details of the current PCAP tracing file along with any existing rollover PCAP tracing files.

2. Manage packet tracing settings.

- **Start packet tracing**

- a. Click **Start**.

- b. On the **Start Packet Tracing** page:

- i) Select the interface name in the **Interface** field.

Note: If no value is selected for the **Interface** field, packet tracing is enabled for all interfaces.

- ii) *Optional:* Click the **Filter** field.

- iii) *Optional:* On the **Set Filter** page, select a pre-defined filter in the **Display Filter** field, or enter the filter manually in the **Filter String** field.

- iv) Click **Save**.

- v) Define the maximum size of the packet tracing file (PCAP file) in the **Maximum File Size** field. This value is the maximum size that the packet tracing file can grow to before packet tracing is disabled.

Note: If no value is selected for the **Maximum File Size** field, the maximum file size is set to half the remaining disk size.

- vi) Define the maximum amount of data (in bytes) to be collected for each frame in the snap length field.

Note: The valid range for this field is 1 to 65535. If no value is specified, the snap length is set to 65535 bytes.

- vii) Define the maximum number of log rotation files in the maximum files field. If this number is greater than 0, the number of log files created is limited to the specified number.

When the maximum number of files is reached the capture begins overwriting files from the beginning, thus creating a 'rotating' buffer. With this 'rotating' buffer, packet capture does not stop once the log files are full. If the value is set to 0 there is a single log file and when this file is full, packet capture stops.

Note: The valid range for this field is 0 to 99. If no value is specified, the maximum files is set to 0.

- c. Click **Start**.

Note: Only a single packet tracing operation can be running at the same time. A new packet trace cannot be started until the PCAP file from the previous trace is deleted.

- **Stop packet tracing**

- a. Click **Stop**.

b. Click **Yes** to confirm the action.

- **Export the packet tracing PCAP file**

a. Select the trace file to export from the Packet Tracing Files grid.

b. Click **Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

c. Confirm the save action in the browser pop-up window.

- **Delete the packet tracing PCAP file**

a. Click **Delete**.

b. Click **Yes** to confirm the action.

Note: If packet tracing is running, the PCAP file cannot be deleted. You must stop the associated packet tracing before you delete the PCAP file.

The delete option deletes all the tracing files including the rollover files.

Creating a cluster

You can configure multiple appliances into a cluster that shares configuration information and runtime information. Use the Cluster Configuration management page to administer cluster support for the appliance.

About this task

The Cluster Configuration page is not available in the LMI when the appliance runs in a Docker environment.

In a cluster, you must designate one of the appliances as the primary master. You can designate up to three subordinate masters, which are called the secondary, tertiary, and quaternary masters. The cluster services can fail over between these masters. The remaining appliances serve as nodes.

You must activate the primary and secondary masters of the cluster at the highest level of all the nodes in the cluster. For example, if any of the nodes have been activated with the Advanced Access Control module, the primary and secondary masters must also be activated for Advanced Access Control. Activation levels are validated when:

- A node joins the cluster. Validation ensures that the primary and secondary masters are activated to at least the same level.
- A new primary or secondary master is set to ensure that the activation level of the new master is at least at the same level of the current primary master.

By default, every appliance operates as a stand-alone cluster with only a single node. You can optionally configure a group of appliances into a cluster with multiple nodes.

For detailed information about clusters, see [“Cluster support” on page 161](#).

Procedure

1. Select an appliance to be the primary master. You can choose any appliance as the primary master if it is not a member of another cluster. If the selected appliance is in another cluster, you must unregister it before you can appoint it as the primary master of a new cluster.
2. On the **General** tab of the **Cluster Configuration** page:
 - a. Select the **Multinode** option.
 - b. Click **Create Cluster**.
 - c. In the **Create Cluster** window, configure the **Cluster Identifier**, and then click **Create Cluster**.

Note: For more information about the Cluster Identifier, see [“Cluster general configuration reference” on page 72](#).

3. Save and deploy this update. The chosen appliance is configured as the primary master of a cluster that can contain multiple nodes.
4. Export the cluster signature file on the primary master. The cluster signature file contains configuration information so that cluster members can identify and communicate with the primary master.
5. Join appliances to the cluster by importing the cluster signature file on each appliance that you want to become a cluster member. The process of joining an appliance to the cluster is a *registration*.
6. Update the cluster configuration on the primary master. As part of the cluster configuration, you can define more masters from the pool of registered nodes. For more information, see [“Failover in a cluster” on page 165](#).
7. Save and deploy the configuration changes.

Note: As a rule, try to limit the number of changes that are made to the cluster configuration in a single policy update.

Related reference

[“Cluster general configuration reference” on page 72](#)

Use the Cluster Configuration management page to administer cluster support for the appliance.

[“Session cache reference” on page 73](#)

Use the Cluster Configuration management page to administer cluster support for the appliance.

[“Runtime database” on page 79](#)

You can view and update the current runtime database settings with the **Runtime Database** tab on the **Cluster Configuration** management page.

Managing cluster configuration

From the Cluster Configuration management page, administer cluster support for the appliance.

Before you begin

Configure the browser to allow pop-up windows if you want to export files.

About this task

The Cluster Configuration page is not available in the LMI when the appliance runs in a Docker environment.

About the **Stand-alone** option:

- It is the default setting on the appliance.
- You can choose it in the following situations:
 - The appliance is a primary master with no other node in the cluster.
 - The appliance is a node in a cluster, but it is in stand-alone mode for recovery purposes.
- The corresponding **Primary Master** default IP address on the appliance is 127.0.0.1.
- These initial settings indicate that by default the appliance operates as a stand-alone cluster with a single node.
- If you do not want this appliance to be the primary master, but rather a node in an existing cluster, follow the steps in [Join the current appliance to the cluster](#).
- When the **Stand-alone** option is selected, the **First Port** field is enabled and the fields under **Masters for All Services** are disabled.

About the **Multi-node** option:

- To use this appliance as the primary master of a cluster with multiple nodes, you must set the **Multi-mode** option.

- When the **Multi-node** option is selected:
 - If the appliance is the primary master, the **First Port** field is enabled.
 - If the appliance is not the primary master, the **First Port** field is disabled.

Note: Cluster configuration updates do not take effect until you deploy the changes through the local management interface.

Procedure

1. From the top menu of the local management interface, select **System > Cluster Configuration**. A list of the nodes in the cluster is under **Nodes**.
2. Take any of the following actions and click **Save**. Clicking **Save** submits all configuration changes from the **General**, **Session Cache**, and **Database** tabs.

Add a description to a cluster node

- a. Select the node.
- b. Click **Edit Description**.
- c. Enter the description for the node.

Specify an appliance to be the primary master of a cluster

- a. Select the **General** tab.
- b. To make the current node the primary master:
 - If the appliance is in stand-alone mode, select **Multi-node**.
 - If the appliance is a non-primary node in a cluster, click **Make Primary Master**.

View and update the current cluster general configuration

Note: You can perform the update operation only through the primary master local management interface.

- a. Select the **General** tab.
- b. Edit the current configuration.

View and update the current cluster session cache configuration

Note: You can perform the update operation only through the primary master local management interface.

The distributed session cache is one of the cluster services. It is used by the IBM Security Verify Access appliance to distribute session data. You must configure the distributed session cache settings for the cluster on the primary master.

- a. Select the **Session Cache** tab.
- b. Edit the current settings.

View and update the current runtime database configuration

The runtime database stores runtime data.

Note: You can perform the update operation only through the primary master local management interface.

- a. Select the **Database** tab.
- b. Edit the current settings.

If you change the location of the runtime database from **Local to the cluster** to **Remote to the cluster**, **Database Maintenance** displays the following error message:

```
System Error FBTRBA091E The retrieval failed because
the resource cannot be found.
```

Complete the following steps to restart the local management interface:

- i) Use an ssh session to access the local management interface.
- ii) Log in as the administrator.
- iii) Type `lmi`, and press Enter.
- iv) Type `restart`, and press Enter.
- v) Type `exit`, and press Enter.

Export the cluster signature file from the cluster master

The signature file contains connection and security information. A node uses this file to register with the cluster master server and participate in the cluster.

Note: You can generate the cluster signature file only on the primary master.

- a. On the **General** tab, click **Export**.

Note: If the **Stand-alone** option is selected, the cluster is a stand-alone cluster and the **Export** function is not available. To export the cluster signature file, select the **Multi-node** option.

- b. Confirm the save operation to export the cluster signature file to your local drive.

Join the current appliance to the cluster

This process is referred to as registration. To review the registration rules, see [“Cluster registration”](#) on page 173.

Note: You must perform this operation through the local management interface of the appliance that is joining the cluster.

- a. On the **General** tab, select the **Multinode** option, and then click **Join Cluster**.
- b. Set the **Cluster Identifier**.

Note: For more information about the Cluster Identifier, see [“Cluster general configuration reference”](#) on page 72.

- c. In the **Join Cluster** window, click **Browse** to select the cluster signature file, which you exported from the primary master. See [Export the cluster signature file from the cluster master](#).
- d. To join the cluster as a restricted node, check **Join as restricted node**. See [“Managing restricted nodes in a cluster”](#) on page 171.
- e. Click **Join Cluster**.

View the status of all nodes

On the **Overview** tab, all cluster nodes are displayed under **Nodes**.

- **Accessible** indicates whether the node can be contacted.
- **Synchronized** indicates whether the node is running with the current cluster configuration. If this column is empty, it means that the current configuration information cannot be obtained from the primary master.
- **Master** indicates whether the node is a cluster master.

Remove a node or a secondary master node from the cluster

This process is referred to as *unregistration*. The cluster configuration prohibits deleting a node that is acting as a master.

Note: Perform this operation through the local management interface of the primary master.

- a. Take one of the following actions:
 - To remove a node, select the node you want to remove from **Nodes** on the **Overview** tab.
 - To remove a secondary master node:
 - i) Delete the secondary master from **Master for All Services** on the **General** tab.
 - ii) Select the node you want to remove from **Nodes** on the **Overview** tab.

- b. Click **Delete**.
- c. To force the removal of the node even if the node cannot be reached, select the **Force**.
- d. Click **Yes**.

Replicate settings across the cluster

You can enable the replication of the IBM Security Verify Access runtime settings and certificate database settings. After you enable the replication option, you can no longer update runtime and certificate database settings from the non-primary nodes.

Note: Perform this operation through the local management interface of the primary master.

- a. Select the **Replication** tab and take one of the following actions:
 - For runtime settings, click **Runtime component**.
 - For certificate database settings, click **Certificate databases**.
 - b. Select **Replicate with Cluster**.
 - c. Click **Yes**.
3. Deploy the changes.

Related reference

[“Cluster general configuration reference” on page 72](#)

Use the Cluster Configuration management page to administer cluster support for the appliance.

[“Session cache reference” on page 73](#)

Use the Cluster Configuration management page to administer cluster support for the appliance.

[“Runtime database” on page 79](#)

You can view and update the current runtime database settings with the **Runtime Database** tab on the **Cluster Configuration** management page.

Cluster general configuration reference

Use the Cluster Configuration management page to administer cluster support for the appliance.

You can view and update the current cluster general configuration:

First Port

The appliance uses a range of 30 ports, starting with the assigned **First Port** value.

This field is mandatory and cannot be empty. The default value is 2020.

The following settings are available only when the **Multinode** option is selected.

Cluster Identifier

The cluster identifier is the IP address or hostname that other nodes in the cluster will use to communicate with this node. If an IP address is used, it must be a statically configured IP address on the current appliance. If a hostname is used, all appliances in the cluster must be able to resolve the hostname. Prior to the 9.0.4.0 release, the first static management IP address was automatically selected by the appliance as the cluster identifier.

Primary Master

The cluster identifier of the primary master. This field is mandatory and cannot be empty.

If you are configuring the appliance as a stand-alone cluster with only a single node, you can use the local IP address (127.0.0.1).

- If you change this value, you must save and deploy the changes before you can configure the remaining fields.
- If you want to configure other masters, you must first join the appliances to the cluster.
- Add the entries for **Primary Master**, **Secondary Master**, **Tertiary Master**, and **Quaternary Master** in order. For example, you cannot add a tertiary unless a secondary exists, and you cannot remove a secondary if a tertiary exists.

- Use the **Secondary Master**, **Tertiary Master**, and **Quaternary Master** fields to manage the supplementary masters. You can update these values at any time to demote existing masters or promote new masters.

When you configure the master nodes, you must adhere to the cluster configuration rules. For more information, see [“Cluster configuration rules”](#) on page 172.

Secondary Master

The cluster identifier of the secondary master.

Master External Reference Entity

The IP address of an external reference device that the primary and secondary masters can use to check the health of the network.

Note: This field is required if both the **Primary Master** and **Secondary Master** fields are set. Otherwise, it is disabled.

Tertiary Master

The cluster identifier of the tertiary master.

Note: You can set this field only if there is a **Secondary Master** defined.

Quaternary Master

The cluster identifier of the quaternary master.

Note: You can set this field only if there is a **Tertiary Master** defined.

Session cache reference

Use the Cluster Configuration management page to administer cluster support for the appliance.

You can view and update the current cluster session cache configuration:

Worker threads

The number of worker threads that handle the server requests. At a minimum, use a number that is greater than the maximum number of clients.

Maximum session lifetime

The maximum lifetime in seconds for each session. Use a value greater than the maximum lifetime of all clients. That is, use a value greater than the maximum **[session] timeout** value that the WebSEAL clients use.

For more information about the **[session] timeout** configuration entry, see the reference topics for the Web Reverse Stanza Proxy in the Knowledge Center.

Client grace period

The grace period in seconds that a client has available to restart and register an interest in the session again before the session is removed from the session cache. This period gives the client a chance to restart without losing the session from the server.

Use a similar value to the idle timeout value for the session on the client. That is, use a value similar to the **[session] inactive-timeout** value that is set in the client Web Reverse Proxy configuration.

For more information about the **[session] inactive-timeout** configuration entry, see the reference topics for the Web Reverse Stanza Proxy in the Knowledge Center.

Connection idle timeout

The maximum length of time that a connection from a client can remain idle before it is closed by the server. A value of 0 indicates that connections will not be reused. The default value is 0.

Support internal clients only

Indicates that only internal clients can use the distributed session cache.

Notes:

- The current version supports internal clients only.
- If this option is selected, the remaining fields are disabled.

- Clients can be turned off. For more information about failover events, search for the Options for handling session failover events topic in the Administering topics in the Knowledge Center. For more information about configuration properties, see [Advanced configuration properties](#).

Support internal and external clients

Indicates that both internal and external clients can use the distributed session cache.

Note: To share the key files across the cluster, navigate to the **SSL Certificates** page and select the **Replicate with Cluster** check box.

Session cache supports mutual TLS. Ensure that the client's certificate in the Distributed Session Cache (DSC) server's trust store and the server's certificate in the client's truststore are added.

The DSC by default supports internal client. It runs on port 2026 and 2027. If external clients support is required, use a different port.

Port

The port on which external clients can communicate with the session cache. This field is mandatory if you enable support for internal and external clients.

Enable SSL

If selected, the distributed session cache uses secure communication with its clients.

Note: If you enable SSL, you must also configure the **Keyfile**.

Keyfile

Lists the existing keyfiles on the appliance. These keyfiles are managed from the SSL certificates page. You can click the **SSL Certificates** link on the right to go to that page.

Note: If you want to share the key files across the cluster, you must go to the **SSL Certificates** page and select the **Replicate with Cluster** check box.

Label

Lists the certificate labels in the selected keyfile. This field is disabled if a keyfile is not selected.

Trace level

Specifies the trace level for the DSC with an integer (0 - 9). 0 indicates that trace is disabled. 9 indicates the maximum trace level.

Note: The trace level setting is not a part of the cluster policy. So this setting is not replicated across the cluster and is not persistent across firmware updates. The trace messages are sent to the log file for the DSC.

Configuration database

You can view and update the current configuration database settings with the **Configuration Database** tab on the **Cluster Configuration** management page.

Note: If Oracle is set as the external configuration database and either the local management interface or runtime server trace specification includes Oracle trace points (for example, `oracle.*`) the underlying Oracle JDBC jar file is changed to a debugging jar file. This might have adverse effects on performance and as such Oracle tracing should only be enabled for debugging purposes and disabled once complete.



Warning: Enabling trace for Oracle components "oracle.*" might result in the Oracle database administrator password being logged in clear text.

The configuration database stores configuration data, including policy information. This data is shared with all appliances in the cluster.

Local to the cluster

Specifies the use of the internal configuration database.

Database export

Exports the current configuration data from the internal database so that it can be imported into an external database of the chosen type. This option is useful if you want to migrate the appliance's internal configuration database to an external database. Supported external database

types are DB2, Oracle, and PostgreSQL. The exported data are compressed into a zip file. A readme file is included in the zip file to provide instructions on how to import the data into the external database.

Note: For DB2 and Oracle, the configuration database schema (table and index definitions), which is available from the **File Downloads** area of the appliance, must be applied to the database that will house the configuration data before the data can be imported. For PostgreSQL, this step is not required as the zip file also contains the database schema.

Remote to the cluster

Specifies the use of an external configuration database. Specify the following information for the external configuration database:

Use external database for internal file sharing

Enable this option to allow the configurations to be modified on non-primary nodes of the cluster.

Note: When you enable this option, the appliance will be rebooted when the change is committed. During the reboot, the files will be migrated between the local file system and the external configuration database.

Type

The database type, which is one of **DB2**, **Oracle**, or **PostgreSQL**.

Address

The IP address or hostname of the external database server.

Port

The port on which the external database server is listening.

Username

The name of the database administrator.

Password

The password for the database administrator.

DB2

Secure

Select this check box to create a secure connection with the DB2® server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the DB2 server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external DB2 server.

Enable HADR and ACR

Select this checkbox to enable High Availability Disaster Recovery and Automatic Client Reroute.

Alternate Address

The IP address or hostname of the failover database server in the HADR configuration.

Alternate Port

The port on which the failover database server in the HADR configuration is listening.

Oracle

Secure

Select this check box to create a secure connection with the Oracle server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the Oracle server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Service name

The name of the service instance on the external Oracle server.

Driver type

Specifies the type of Oracle JDBC driver that is used to connect to the Oracle server. Available options are **Thin** and **OCI**.

PostgreSQL

Note: High availability, with an external PostgreSQL server, is achieved through the use of an external load balancer.

Secure

Select this check box to create a secure connection with the PostgreSQL server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the PostgreSQL server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external PostgreSQL server.

Enable failover support

Select this check box to enable PostgreSQL failover support. Once this box is checked, the failover servers management section is enabled.

Failover servers

Manage the PostgreSQL failover servers in this section.

Add

1. Click the **Add** button to add a new failover server. A new dialog opens.
2. Specify the new failover server address and port.
3. Click **Submit** to add the server.

Delete

1. Select the failover server in the grid.
2. Click the **Delete** button to remove the server from the list.

Move Up and Move Down

1. Select the failover server in the grid.
2. Click the **Move Up** or **Move Down** button to change the order of the server in the list.

Deploying an external configuration database

To optimize performance or increase storage capacity for the appliance, you can deploy an external configuration database. You can configure the appliance to connect to DB2, PostgreSQL, or Oracle database on an external server.

About this task

A Security Verify Access appliance with Advanced Access Control or Federation includes an internal database to store configuration data.

The appliance provides scripts to deploy the configuration database on an external DB2, PostgreSQL, or Oracle server. You can then configure the appliance to use the external database.

The **Oracle Compatibility** mode in DB2 must be turned off when you are using an external DB2 Configuration Database or HVDB with IBM Security Verify Access.

Note: IBM Security Verify Access uses the configured username as the schema name to connect to the database. Therefore, aliases might need to be created in DB2 in the event that the username does not match the schema name in the database.

Procedure

1. Use the **File Downloads** management page in the local management interface to access the configuration database deployment files for your environment.

Database type	Deployment scripts
DB2	/access_control/database/db2/config/ cluster_config_db2.sql
PostgreSQL	/access_control/database/postgresql/config/ cluster_config_postgresql.sql
Oracle	/access_control/database/oracle/config/ cluster_config_oracle.sql

2. Save the deployment script on the database server.
3. Run the DB2, PostgreSQL, or Oracle script to create the external database.

PostgreSQL script

Run the following command:

```
psql --echo-all --variable ON_ERROR_STOP=1 --file <sql file name>  
--username <username> --host <host> --port <port> <database name>
```

Oracle script

- a. Copy the downloaded `cluster_config_oracle.sql` file into the Oracle home directory. For example, `ORACLE_HOME=/opt/oracle/app/oracle/product/11.2.0/dbhome_1`
- b. Log in to SQL*Plus.
- c. At the SQL prompt, run **START cluster_config_oracle.sql**.

DB2 script

- a. Create a DB2 instance to contain the configuration database. For information about creating the DB2 instance, see the DB2 documentation.
- b. Open the `cluster_config_db2.sql` file in an editor on the DB2 server.
- c. Replace the following macros with the values specific to your environment:

&DBINSTANCE

The name of the DB2 instance.

&DBUSER

The name of the DB2 administrator.

&DBPASSWORD

The password for the DB2 administrator.

- d. Save the changes.
- e. Log in to the DB2 Command utility (Windows) or DB2 host (UNIX) as the DB2 administrator.
- f. Run the following command:

```
db2 -tsvf <fully_qualified_path_to_script>
```

The following example shows the fully qualified path to the script:

```
db2 -tsvf /tmp/cluster_config_db2.sql
```

4. Validate that the tables were successfully created.
5. Ensure that no errors were returned during the creation and log in to the database to manually check that the tables exist.
6. Populate the database with initial configuration data. Export the embedded configuration database data and then import this data into the external server.

Note: Ensure that data is exported from the appliance and validated. Ensure that this data is also imported into the external server successfully.

7. From the top menu of the local management interface, select **System > Cluster Configuration** to open the **Cluster Configuration** management page.
8. Select the **Database** tab.
9. You must enter the following JDBC connection information:

Type

The database type, which is either DB2, PostgreSQL, or Oracle.

Address

The IP address of the external database server.

Port

The port on which the external database server is listening.

Username

The name of the database administrator.

Password

The password for the database administrator.

DB2 also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external DB2 server.

Complete the following steps to identify and specify the DB2 database name when your DB2 database is remote to the cluster that you are configuring.

- a. Open the `cluster_config_db2.sql` file that was used to create the database and tables.
- b. In the **CREATE DATABASE** entry, get the name that is specified. In the following entry, `CONFIG` is the string that identifies the default database name:

```
CREATE DATABASE CONFIG ALIAS CONFIG using codeset UTF-8 territory us
PAGESIZE 8192 WITH "CONFIG Tables";
```

Note: `PAGESIZE 8192` is an example. Adjust according to your requirements.

PostgreSQL also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external PostgreSQL server.

Oracle also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the `lmi_trust_store` and also a keystore which only contains public keys that needs to be created. Use the **SSL Certificates** page to create this keystore and to import the appropriate certificate

Certificate Store

Choose the keystore which contains the certificate that will be used to communicate with the server.

Service name

Specify the name of the Oracle instance on the external server. Contact your Oracle database administrator for this information. SID will work but might show a warning in the LMI on saving configuration. This can be ignored.

10. Click **Save**.

11. Deploy the changes.

Results

The appliance is configured to use the configuration database that is deployed on the external system.

What to do next

- Tune the external database by setting the configuration parameters. See [Runtime database tuning parameters](#).
- On Oracle 12.2 check that the supported login protocol is set on the DBMS. If it is not, set the value `SQLNET.ALLOWED_LOGON_VERSION=11` in the `sqlnet.ora` file. For more information, see <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/required-tasks-complete-upgrading-oracle-database.html#GUID-12B920E9-B2DA-48A0-832C-3E07D172A011>

Runtime database

You can view and update the current runtime database settings with the **Runtime Database** tab on the **Cluster Configuration** management page.

Note: If Oracle is set as the external runtime database and either the local management interface or runtime server trace specification includes Oracle trace points (for example, `oracle.*`) the underlying Oracle JDBC jar file is changed to a debugging jar file. This might have adverse effects on performance and as such Oracle tracing should only be enabled for debugging purposes and disabled once complete.



Warning: Enabling trace for Oracle components “oracle.*” might result in the Oracle database administrator password being logged in clear text.

The runtime database contains runtime data that is used by the context-based access component. You can configure this database as an embedded database or an external database. The embedded database is suitable for small environments only. For large-scale production environments, configure an external database.

Note: Legacy OIDC was deprecated in IBM Security Verify Access v10.0.0. If a Verify Access installation prior to version 10.0.0 is upgraded to version 10.0.0 or later, use the `RemoveOauthDBSchema.sql` to remove the tables were used by legacy OIDC from the external HVDB. Go to **System > Secure Settings > File Downloads > Federation > Database > Common > RemoveOauthDBSchema.sql**.

If Oracle is set as the external database, the user who owns the schema has to be dropped. This step is not included in the script as it has to be executed with caution due to the possibility that the user can own other schemas as well.

Local to the cluster

Specifies the use of the internal runtime database.

Note: Only the **Maximum Size** field relates to the internal runtime database. If you use the internal runtime database, all other fields are disabled.

Maximum Size (% of available disk)

The size of the internal runtime database. If you select the **Local to the cluster** option, this field is mandatory. The maximum size is a percentage of the remaining disk space at the time that the policy is applied.

The valid value range is from 10% to 80%. If a change in this value results in a calculated maximum size, which is smaller than the current size of the database, the database must be re-created. In this case, all existing data from the database is lost.

To determine the percentage of available disk space to assign to the internal database, consider the following aspects of your environment:

- The current disk usage on the appliance. You can view the **Disk Usage** on the Appliance Dashboard in the LMI.
- Internal disk requirements for other utilities such as logging and snapshots.

Database export

Exports the current runtime data from the internal database so that it can be imported into an external database of the chosen type. This option is useful if you want to migrate the appliance's internal runtime database to an external database. Supported external database types are DB2, Oracle, and PostgreSQL. The exported data are compressed into a zip file. A readme file is included in the zip file to provide instructions on how to import the data into the external database.

Remote to the cluster

Specifies the use of an external runtime database. Specify the following information for the external runtime database:

Type

The database type, which is either **DB2**, **Oracle**, or **PostgreSQL**.

Address

The IP address or hostname of the external database server.

Port

The port on which the external database server is listening.

Username

The name of the database administrator.

Password

The password for the database administrator.

DB2

Secure

Select this check box to create a secure connection with the DB2 server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the DB2 server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external DB2 server.

Enable High Available Disaster Recovery and Automatic Client Reroute

Select this checkbox to enable HADR and ACR.

Alternate Address

The IP address or hostname of the failover database server in the HADR configuration.

Alternate Port

The port on which the failover database server in the HADR configuration is listening.

Oracle

Secure

Select this check box to create a secure connection with the Oracle server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the Oracle server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Service name

The name of the service instance on the external Oracle server.

Driver type

Specifies the type of Oracle JDBC driver that is used to connect to the Oracle server. Available options are **Thin** and **OCI**.

PostgreSQL

Note: High availability, with an external PostgreSQL server, is achieved through the use of an external load balancer.

Secure

Select this check box to create a secure connection with the PostgreSQL server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the PostgreSQL server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external PostgreSQL server.

Enable failover support

Select this check box to enable PostgreSQL failover support. Once this box is checked, the failover servers management section is enabled.

Failover servers

Manage the PostgreSQL failover servers in this section.

Add

1. Click the **Add** button to add a new failover server. A new dialog opens.
2. Specify the new failover server address and port.
3. Click **Submit** to add the server.

Delete

1. Select the failover server in the grid.
2. Click the **Delete** button to remove the server from the list.

Move Up and Move Down

1. Select the failover server in the grid.
2. Click the **Move Up** or **Move Down** button to change the order of the server in the list.

Deploying an external runtime database

To optimize performance or increase storage capacity for the appliance, you can deploy an external runtime database. You can configure the appliance to connect to DB2, PostgreSQL, or Oracle database on an external server.

About this task

Note: If Oracle is set as the external runtime database and either the local management interface or runtime server trace specification includes Oracle trace points (for example, `oracle.*`) the underlying Oracle JDBC jar file is changed to a debugging jar file. This might have adverse effects on performance and as such Oracle tracing should only be enabled for debugging purposes and disabled once complete.



Warning: Enabling trace for Oracle components “oracle.*” might result in the Oracle database administrator password being logged in clear text.

A Security Verify Access appliance with Advanced Access Control includes an internal database to store user data such as session attributes and device fingerprints. This embedded database is suitable for small environments. In a production environment, use an external runtime database that can handle the required volume of data.

The appliance provides scripts to deploy the runtime database on an external DB2, PostgreSQL, or Oracle server. You can then configure the appliance to use the external database.

Procedure

1. Use the **File Downloads** management page in the local management interface to access the runtime database deployment files for your environment.

Database type	Deployment scripts
DB2	<code>/access_control/database/db2/runtime/ isam_access_control_db2.sql</code>
PostgreSQL	<code>/access_control/database/postgresql/runtime/ isam_access_control_postgresql.sql</code>
Oracle	<code>/access_control/database/oracle/runtime/ isam_access_control_oracle.sql</code>

2. Save the deployment script on the database server.
3. Run the DB2, PostgreSQL, or Oracle script to create the external database.

PostgreSQL script

Run the following command:

```
psql --echo-all --variable ON_ERROR_STOP=1 --file <sql file name>  
--username <username> --host <host> --port <port> <database name>
```

Oracle script

- a. Copy the downloaded `isam_access_control_oracle.sql` file into the Oracle home directory. For example, `ORACLE_HOME=/opt/oracle/app/oracle/product/11.2.0/dbhome_1`
- b. Log in to SQL*Plus.
- c. At the SQL prompt, run **START isam_access_control_oracle.sql**.

DB2 script

- a. Create a DB2 instance to contain the runtime database. For information about creating the DB2 instance, see the DB2 documentation.

- b. Open the `isam_access_control_db2.sql` file in an editor on the DB2 server.
- c. Replace the following macros with the values specific to your environment:

&DBINSTANCE

The name of the DB2 instance.

&DBUSER

The name of the DB2 administrator.

&DBPASSWORD

The password for the DB2 administrator.

- d. Save the changes.
- e. Log in to the DB2 Command utility (Windows) or DB2 host (UNIX) as the DB2 administrator.
- f. Run the following command:

```
db2 -tsvf <fully_qualified_path_to_script>
```

The following example shows the fully qualified path to the script:

```
db2 -tsvf /tmp/isam_access_control_db2.sql
```

4. Validate that the tables were successfully created.
5. Ensure that no errors were returned during the creation and log in to the database to manually check that the tables exist.
6. From the top menu of the local management interface, select **System > Cluster Configuration** to open the **Cluster Configuration** management page.
7. Select the **Database** tab.
8. You must enter the following JDBC connection information:

Type

The database type, which is either DB2, PostgreSQL, or Oracle.

Address

The IP address of the external database server.

Port

The port on which the external database server is listening.

Username

The name of the database administrator.

Password

The password for the database administrator.

DB2 also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external DB2 server.

Complete the following steps to identify and specify the DB2 database name when your DB2 database is remote to the cluster that you are configuring.

- a. Open the `isam_access_control_db2.sql` file that was used to create the database and tables.

- b. In the **CREATE DATABASE** entry, get the name that is specified. In the following entry, HVDB is the string that identifies the default database name:

```
CREATE DATABASE HVDB ALIAS HVDB using codeset UTF-8 territory us
PAGESIZE 8192 WITH "HVDB Tables";
```

Note: *PAGESIZE 8192* is an example. Adjust according to your requirements.

PostgreSQL also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external PostgreSQL server.

Oracle also requires the following information:

Secure

Select this check box to create a secure connection with the server.

Note: Before a secure connection can be established, you must first import the certificate that the appliance uses to communicate with the server into the **lmi_trust_store** and also a keystore which only contains public keys that needs to be created. Use the **SSL Certificates** page to create this keystore and to import the appropriate certificate.

Certificate Store

Choose the keystore which contains the certificate that will be used to communicate with the server.

Service name

Specify the name of the Oracle instance on the external server. Contact your Oracle database administrator for this information.

9. Click **Save**.

10. Deploy the changes.

Results

The appliance is configured to use the runtime database that is deployed on the external system.

General Information

HVDB data is language agnostic. Character support might not be an issue with regards to the installation that is chosen.

Oracle DB_BLOCK_SIZE or PAGE_SIZE can vary based on deployments but it is suggested to have at least 16384.

The suggested character set is ALUTF8.

User permissions for database must have read-write access and the ability to execute the commands in the SQL script.

Oracle:

```
dbca -createDatabase -templateName sampletemplate.dbc -gdbname hvdb -sid hvdb -responseFile
NO_VALUE -
characterSet AL32UTF8 -memoryPercentage 20 -emConfiguration LOCAL -dbsnmpPassword
mypassword -sysPassword
mypassword -systemPassword mypassword -silent
```


DB2:

```
CREATE DATABASE HVDB ALIAS HVDB using codeset UTF-8 territory us
PAGESIZE 8192 WITH "HVDB Tables";
```

What to do next

- Tune the external database by setting the configuration parameters. See [Runtime database tuning parameters](#).
- On Oracle 12.2 check that the supported login protocol is set on the DBMS. If it is not, set the value `SQLNET.ALLOWED_LOGON_VERSION=11` in the `sqlnet.ora` file. For more information, see <https://docs.oracle.com/en/database/oracle/oracle-database/12.2/upgrd/required-tasks-complete-upgrading-oracle-database.html#GUID-12B920E9-B2DA-48A0-832C-3E07D172A011>

Managing Distributed Session Cache in Docker

Use this page to view and update the Distributed Session Cache (DSC) configuration data in a Docker environment.

About this task

This page is available only when Security Verify Access is running in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > DCS Configuration**.
2. Specify the general settings.

Worker Threads

The number of worker threads that are allocated to processing requests.

Maximum Session Lifetime

The maximum lifetime (in seconds) of any session that is stored by the DSC.

Client Grace Period

The length of time (in seconds) that a client (aka Web Reverse Proxy) has to reconnect before sessions that are owned by that client are discarded.

Connection idle timeout

The maximum length of time that a connection from a client can remain idle before it is closed by the server. A value of 0 indicates that connections will not be reused. The default value is 0.

Service Port

The port number on which the DSC will listen for requests.

Replication Port

The port number on which the DSC will listen for requests from replicated DSC servers.

3. Specify the external connection settings. This data is used when configuring the DSC clients (aka Web Reverse Proxy and administration client). It corresponds to the host identifier and port used to connect to the replication and session services of the various DSC servers. For failover purposes, up to 4 DSC servers can be configured (primary, secondary, tertiary, and quaternary).

Address

The IP address or resolvable host name over which clients can connect to the DSC.

Service Port

The port that can be used by clients to connect to the DSC for session requests. This port can be different to the configured **Service Port** under general settings due to the port mapping capability of Docker.

Replication Port

The port that a DSC server should use when connecting to a replicated DSC server. This port can be different to the configured **Replication Port** under general settings due to the port mapping capability of Docker.

4. Click **Save**.

Managing database configuration in Docker

You can view and update the current runtime database settings with the Runtime Database tab on the Database Configuration management page.

About this task

The runtime database contains runtime data that is used by the context-based access and federation components. The database must be configured prior to activating either of these components.

This page is available only when Security Verify Access is running in a Docker environment.

Procedure

1. From the top menu, select **System > Network Settings > Database Configuration**.
2. Specify the following information for the runtime database.

Type

The database type, which is either **DB2**, **Oracle**, or **PostgreSQL**.

Address

The IP address or hostname of the external database server.

Port

The port on which the external database server is listening.

Username

The name of the database administrator.

Password

The password for the database administrator.

The following fields are specific to each type of database.

DB2

Secure

Select this check box to create a secure connection with the DB2 server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the DB2 server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external DB2 server.

Enable High Available Disaster Recovery and Automatic Client Reroute

Select this checkbox to enable HADR and ACR.

Alternate Address

The IP address or hostname of the failover database server in the HADR configuration.

Alternate Port

The port on which the failover database server in the HADR configuration is listening.

Oracle

Secure

Select this check box to create a secure connection with the Oracle server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the Oracle server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Service name

The name of the service instance on the external Oracle server.

Driver type

Specifies the type of Oracle JDBC driver that is used to connect to the Oracle server. Available options are **Thin** and **OCI**.

PostgreSQL

Note: High availability, with an external PostgreSQL server, is achieved through the use of an external load balancer.

Secure

Select this check box to create a secure connection with the PostgreSQL server.

Note: Before a secure connection can be established, you must first import the certificate for the appliance to use for communication with the PostgreSQL server. The certificate must be imported into the **lmi_trust_store** and **rt_profile_keys** key files. Use the **SSL Certificates** page to import the appropriate certificate.

Database name

The name of the database instance on the external PostgreSQL server.

3. Click **Save**.

System settings

Information about managing system settings on your appliance.

Configuring date and time settings

Use the **Date/Time Configuration** page to configure the date, time, time zone, and NTP server information.

Procedure

1. Click **System > System Settings > Date/Time**
2. Configure the following options:

Option	Description
Time Zone	Specifies the time zone for the appliance.
Date/Time	Specifies the day, month, year, and time for the appliance.
NTP Server address	Lists the NTP (NIST Internet Time Service) servers the appliance uses. You can enter multiple NTP servers, separated by commas.

3. Click **Save**.

Configuring administrator settings

Use the **Administrator Settings** management page to tune the local management interface so that it can run more efficiently.

About this task

LMI Access Log

An access log of requests to the management interface can be enabled by setting the LMI Access Log Format parameter. The format of this string should conform to the HTTP access log format for IBM Liberty application server. More details about log format options can be found in the [Liberty documentation](#).

Procedure

1. Click **System > System Settings > Administrator Settings**.

The available tuning parameters are listed in a table.

2. Modify the parameters.

- To edit a parameter, select the parameter from the table and click **Edit**. In the edit window, change the parameter value.

Note: If you edit the SSHD port parameter in a clustered environment, all machines in the cluster must be configured with the same SSHD port. As the configured port will not be automatically distributed across all machines in the cluster, each machine must be updated individually.

- To delete the current settings for a parameter and change its value to unset, select the parameter from the table and click **Delete**.

Note: The administrator password cannot be reset.

3. Deploy the changes.

Configuring tracing for the local management interface

Use the **LMI Tracing** page to configure the tracing specifications for different components of the local management interface.

About this task

You can now set the tracing specifications of the local management interface for debugging purposes.

Note:

- Changing these tracing specifications might have an adverse effect on the performance of the local management interface.
- Setting trace for Oracle components "oracle.*" results in the underlying Oracle JDBC jar file being changed to a debugging jar file. This might have adverse effects on performance and as such Oracle tracing should only be enabled for debugging purposes and disabled once complete.



Warning: Enabling trace for Oracle components "oracle.*" might result in the Oracle database administrator password being logged in clear text.

Procedure

1. Select **System > System Settings > Administrator Settings**.
2. Click **LMI Tracing**.

The LMI tracing components and trace levels are displayed.

Component	Description
com.ibm.isam.*	This option enables tracing for the components of all offerings, which include the Security Verify Access Base, Advanced Access Control, and Federation offerings.
com.ibm.isam.core.*	This option enables tracing for the common components of all offerings. These common components are shared by the various offerings. For example, the Security Verify Access runtime and SSL certificates management.

<i>Table 9. LMI tracing components (continued)</i>	
Component	Description
com.ibm.isam.wga.*	This option enables tracing for the components of the Security Verify Access Base offering. For example, the management of reverse proxy instances.
com.ibm.isam.mga.*	This option enables tracing for the components of the Security Verify Access Advanced Access Control and Federation offerings. For example, the risk based analysis (RBA) configuration, the management of federations, partners, and module chains.
com.ibm.mesa.*	This option enables tracing for the underlying components that compose the LMI framework. These components are used both as a base for all of the offerings and to provide the management of most system settings. For example, updates and network configuration.
HTTP	This option enables tracing for the components of the web application server that are involved in HTTP communication.
SSL	This option enables tracing for the components of the web application server that are involved in SSL communication.
JSP	This option enables tracing for the JavaServer pages components of the web application server.
Servlet Engine	This option enables tracing for the servlet engine and web container components of the web application server.
Session Management	This option enables tracing for the components of the web application server that make up the session and session management functionality.
Configuration	This option enables tracing for the configuration of the web application server.
Native Security	This option enables tracing for the native security components of the web application server.

<i>Table 10. LMI trace levels</i>	
Level	Description
all	All events are logged. If you create custom levels, all includes those levels, and can provide a more detailed trace than finest.
finest	A more detailed trace that includes all the detail that is needed to debug problems.
finer	Detailed trace information.
fine	Trace information that includes general trace, method entry, exit, and return values.
detail	General information that details the subtask progress.
config	Configuration change or status.

<i>Table 10. LMI trace levels (continued)</i>	
Level	Description
info	General information that outlines the overall task progress.
audit	Significant event that affects the server state or resources.
warning	Potential error or impending error. This level can also indicate a progressive failure, for example, the potential leaking of resources.
severe	Task cannot continue. But component, application, and server can still function. This level can also indicate an impending unrecoverable error.
fatal	Task cannot continue. Component, application, and server cannot function.
off	Logging is turned off.

3. Define the trace specifications in either of the following methods.
 - Select a component and trace level from the table, and then click the **Add**. Repeat this procedure until all trace specifications are added.
 - Manually enter the trace specifications in the **Trace Specification** text area.
4. Click **Save**.
5. Deploy the changes.

Note: The local management interface is automatically restarted so that the changes can take effect.

Configuring management authentication

To configure management authentication with the local management interface and the command line interface, use the **Management Authentication management** page.

Procedure

1. From the top menu, select **System > System Settings > Management Authentication**.
All current management authentication settings are displayed.
2. In the Main tab:
 - Select **Local User Database** if you want to use the local user database for authentication.
 - Select **Remote LDAP User Registry** if you want to use the remote LDAP user registry for authentication.

Note: If a remote user registry is configured for management authentication, the local administrator user (admin) can continue to be referenced with the "admin@local" user name. You can use this as a fail safe in the event that the remote user registry is not reachable.

 - a. In the **LDAP** tab:
 - i) Specify the name of the LDAP server in the **Host name** field.
 - ii) Specify the port over which to communicate with the LDAP server in the **Port** field.
 - iii) Select the **Anonymous Bind** check box if the LDAP user registry supports anonymous bind.
 - iv) Specify the DN of the user that is used to bind to the registry in the **Bind DN** field.
 - v) Specify the password that is associated with the bind DN in the **Bind Password** field.

- vi) *Optional*: If you want to enable LDAP client debugging for authentication related issues, select the **Debug** check box. The LDAP debugging log can be viewed by going to **Monitor > Application Log Files** and accessing the **management_ui > ldap_debug.log** file.

b. In the **LDAP General** tab:

- i) Specify the name of the LDAP attribute that holds the supplied authentication user name of the user in the **User Attribute** field.
- ii) Specify the name of the LDAP attribute that is used to hold the members of a group in the **Group Member Attribute** field.
- iii) Specify the base DN that is used to house all administrative users in the **Base DN** field.
- iv) Specify the DN of the group to which all administrative users belong in the **Administrative Group DN** field.

Note: All administrative users must have permission to view the specified `admin_group_dn` group within the user registry.

c. In the **LDAP SSL** tab:

- i) Select the **Enable SSL** check box to define whether SSL is used when the system communicates with the LDAP server.
- ii) Select the name of the key database file in the **Key File Name** field.
- iii) Select the name of the certificate to be used if client authentication is requested by the LDAP server in the **Certificate Label** field.

d. In the **User Mapping** tab:

- i) Select the **Enable** checkbox to enable the mapping of a client certificate DN to a new format prior to the user being authenticated.
- ii) The User mapping script field contains a placeholder script that will be set as the default unless a new script is entered. The script contains a Javascript function that takes a Map as an input and returns a String that represents the new DN. The map contains the following values that can be used in the mapping:

- **cert** > The actual X509Certificate object.
 - **principal** > The X500Principal in the certificate.
 - **san** > The certificate subject alternative names.
 - **dn** > The certificate DN.
 - **baseDN** > The management authentication configuration value for base DN.
 - **userAttribute** > The management authentication configuration value for user attribute.
 - **groupMemberAttribute** > The management authentication configuration value for group member attribute.
- Each distinct rdns from the DN separately. For example the DN of `cn=testuser,o=test,c=us` would result in 3 entries in the map:
- **cn** > `testuser`
 - **o** > `test`
 - **c** > `us`

The default script will use the `cn` from the map and combine it with the `userAttribute` and `baseDN` to create the new DN:

```
function mapUser(props) {
    var user = props.get("cn");
    return props.get("userAttribute") + "=" + user + "," + props.get("baseDN");
}
```

3. Click **Save** to save your settings.

Note: For the changes to take effect, they must be deployed.

4. *Optional*: Click **Test** to test the authentication.

Note: If there have been changes made to the management authentication configuration that have not yet been deployed, this test will run using the undeployed configuration.

- a. In the **Test Authentication** window, enter the user name in the **Username** field.
- b. Enter the password in the **Password** field.
- c. Click **Test**.

If the authentication is successful, a success message is displayed. If the authentication is not successful, an error message is displayed.

Managing roles of users and groups

Assign certain roles to users and groups to control which sections of the local management interface and web services they can access.

About this task

By default, role-based authorization is disabled on the appliance. You must first enable this function from the management interface to make use of it.

With **Management Authorization**, you can perform the following tasks:

- Add or remove a role.
- Assign a role to groups or users in local or remote LDAP user registry.

Note: You can search for remote LDAP users or groups by entering a search pattern and clicking **Search**. Then, select the user or group from the search results and click **Add**.

- Edit permissions for a role.

The roles for a user session are determined when a user first logs in. If the authorization configuration is modified and deployed when a user is logged in, the changes take effect immediately.

You can customize the default roles to better suit your environment. You can also remove all default roles and create new ones from scratch.

Note: If you plan to use the default roles, you must carefully review these roles to ensure that they are appropriate for your environment.

The default roles are not updated after an appliance firmware upgrade. If the appliance firmware upgrade introduces new features, existing roles are not updated to include permission for any new features. The default roles can be manually updated in the **Management Authorization** page. See [Step 3 "Editing permissions for a role"](#).

The authorization settings do not affect the main system account **admin**, which always has read and write permission to all features. The **admin** account can be used for recovery.

Permissions can be set for all features in the appliance except for the **Home: Appliance Dashboard**. Any user who can authenticate can view **Home: Appliance Dashboard**, even if they are not assigned to any roles.

To ensure complete flexibility with the role configuration, the permissions for each feature are controlled separately. Some pages in the local management interface, such as the **Management Authorization** page, use multiple features. As a result, users might need permissions for more than one feature to use all of the features on a particular page of the local management interface. For example, to access all of the functions on the **Management Authorization** page, the user needs permissions for the following features:

- Account Management
- Management Authorization

If a user clicks a link or attempts to complete an action for which they do not have the appropriate permission, an error message is returned. The error message includes the details about which permission is required for the selected action.

When you search for remote LDAP users or groups, consider the following points:

- Users are assumed to be contained in the **Base DN** and are identified based on the **User Attribute** that is set on the **Management Authentication** page.
- Groups are also assumed to be contained in the **Base DN** that is defined on the **Management Authentication** page.
- Groups are identified based on **cn**.
- Groups must be among the following types: **group**, **groupofUniqueName**, or **groupOfNames**.

Authorization enforcement applies to the local management interface, web services, and client certificate authentication.

Authorization enforcement in the local management interface

When a user logs in the local management interface, the menu displays only the pages that the user has access to. When users attempt to go to a page to which they do not have access, a page is displayed that explains that the user does not have authorization to view the page. When a user views a page with read-only permission, users cannot modify the configuration or change the state of any services on the page. If a user attempts to do so, a message is displayed stating that the user does not have permission to perform the requested action.

Authorization enforcement in web services

If a user has read-permission for a feature, they can perform GET requests against the associated Web services. If a user has write-permissions on a feature, they can issue any of the associated GET, POST, PUT, and DELETE web services. When a user attempts to issue a web service request that they are not authorized to perform, they receive a response with the HTTP status code **403 Forbidden** and a message that states that they are not authorized to complete the transaction.

Authorization enforcement in client certificate authentication

If you want to use client certificates to authenticate to the local management interface, ensure that the authorization framework can map the DN of the presented client certificate to a user that exists in the registry that is used for authentication.

For example, a certificate is presented with DN: `cn=testUser,ou=qa,o=ibm,c=au`.

When you use a remote LDAP user registry for authentication, the authorization decision is made for a user that matches the entire DN in the user registry.

For example, a user that matches `cn=testUser,ou=qa,o=ibm,c=au` is searched for in the remote LDAP user registry, and the policy that is associated with that user is enforced.

When you use the local user database, the authorization decision is made for a user that matches the CN of the presented DN. For example, the user that is called `testUser` is searched for in the local user database, and the policy that is associated with that user is enforced.

Authorization enforcement in the Command Line Interface

Access to the command line interface from the console or SSH can be restricted by using the 'CLI and CLI Web Service' feature. Only those users who have 'write' access to this feature will be permitted to access the command line interface.

A user can be assigned multiple roles. In this case, the user receives the highest cumulative permission from these roles for each feature. For example, if they are assigned two roles and one role has read-permission for a feature but the second role has write-permission for the feature, the user is granted write-permission.

Note: The appliance caches authentication details to reduce load on the user registry. The authentication details might be used for up to 10 minutes after they are changed. This behavior can be changed by using an advanced tuning parameter. Add the advanced tuning parameter `lmi.authCache.baenabled` with a value of `false` to disable this caching. See [“Managing advanced tuning parameters”](#) on page 96.

A performance penalty is incurred when you use this parameter. The user registry is queried when:

- A user logs in the local management interface through the browser.
- A request to the web services API by using Basic Authentication is received.

There is some degradation of performance in environments that make heavy use of the web services API by using Basic Authentication.

Full Read and Full Write roles

The **Full Read** and **Full Write** roles do not use a permissions list when determining authorization. Instead administrators who are members of the **Full Read** role are permitted to perform HTTP GET requests to all Local Management Interface URLs. Similarly users who are members of the **Full Write** role are permitted to perform HTTP GET, POST, PUT and DELETE requests to all Local Management Interface URLs.

Procedure

1. Select **System > System Settings > Management Authorization**.
2. Under **Roles**, select the **Enable Authorization Roles** check box.
3. Follow the prompts to complete the action you want to take.

Tip: Use the quick filter to retrieve group names, user names, and features.

Adding a role

- a. In the **Roles** panel on the left, click **New**.
- b. In the **Create New Role** window, enter a name for the new role.
- c. Click **OK**.

Removing a role

- a. In the **Roles** panel on the left, select the role to delete.
- b. Click **Delete**.
- c. In the **Removing Role** window, verify that the role name to delete is correct and then click **Yes**.

Assigning a role to local groups or users

- a. In the **Roles** panel on the left, select the role to edit membership for.
- b. In the **Role Membership** panel on the right, select the **Local User Database** tab if it is not already selected.
- c. Click **Edit** above the group name table or the user name table.
- d. In the **Edit Local Members** window, select or clear the check box on the **Groups** and **Users** tabs as needed.
- e. Click **OK**.

Assigning a role to LDAP groups or users

- a. In the **Roles** panel on the left, select the role to edit membership for.
- b. In the **Role Membership** panel on the right, select the **Remote LDAP User Registry** tab if it is not already selected.
- c. In the **Edit Remote LDAP Members** window, modify LDAP groups and users on the **Groups** and **Users** tabs as needed.
 - To add an LDAP group or user, enter the details in the text field and then click **Add**.
 - To remove an LDAP group or user, select the entry and then click **Delete**.
- d. Click **OK**.

Editing permissions for a role

- a. In the **Roles** panel on the left, select the role to edit permissions for.
- b. In the **Features** panel on the right, select the permission that you want from the drop-down list in each row.

If you upgrade from a previous version of the appliance, new role membership features are set to **None** by default. Configure the permissions, if necessary.

Note: The displayed features reflect the features that are available in the activated offerings. If you deactivate a product, the features that are specific to that product are removed from any existing roles. If you reactivate the product in the future, these features and the associated permissions are added to the roles again. Any permissions from a prior activation are re-instantiated. If it is the first time that the product is activated, the product-specific features are added to each role with no assigned permissions.

c. Click **Save** to save the permission settings.

Viewing and updating management SSL certificates

View and update the management SSL certificate details in the **Management SSL Certificate** page of the local management interface.

View the details of the current management SSL certificate

1. From the top menu, select **System > System Settings > Management SSL Certificate**.
2. The details of the current management certificate are displayed.

Update the management SSL certificate

1. From the top menu, select **System > System Settings > Management SSL Certificate**.
2. Select **Update**.
3. Under **Certificate File**, click **Browse**.
4. Browse to the directory that contains the certificate container file and select the file.

Note: The certificate container file must be PKCS12 format (.p12 file) and can contain only a single certificate. You can generate this certificate on a server that hosts a certificate utility such as iKeyman. This certificate is used as the management SSL certificate.

5. Click **Open**.
6. Click **Update**. A message that indicates successful update is displayed.

Note: For the changes to take effect, they must be deployed.

Managing users and groups

You can manage administrative users and groups, change user passwords, and configure group membership with **Account Management** so that you can control their access.

About this task

With **Account Management**, you can perform the following tasks:

- Add or delete a user.
 - All current users are in the Users table.
 - You cannot change information about `admin`, the statically configured user.
- Change a user password.
 - The first and last character of the password cannot be a space character. Any leading or trailing spaces in the password are removed.
 - If the user is logged in, you
 - Can also click **Set Password** in the top banner.
 - Must enter the existing password before you can change it.
 - If you change the password while logged in as the `admin` user, the password update is automatically deployed without the need for a manual deployment step.
- Create or delete a group.

- Add a user to or remove a user from a group.
 - You can do this step either from the **Users** or **Groups** page.
 - The links in the title bars switch between **Users** and **Groups**.
- Add or change role membership. See “Managing roles of users and groups” on page 92.

Note: The authentication cache that stores the credentials for configured users refreshes every 10 minutes by default. If you just changed a user password or deleted a user, the change might not be effective immediately. It is possible for the user to continue performing web service calls with their original credentials until the authentication cache is refreshed.

Procedure

1. From the top menu, select **System > System Settings > Account Management**.
2. Select the **User** or **Group** link.
3. Follow the prompts to complete the action you want to take.

Managing advanced tuning parameters

Change the advanced tuning parameter values only under the supervision of IBM software support.

In the local management interface, select **System > System Settings > Advanced Tuning Parameters**. The following table lists the advanced tuning parameters available.


<i>Table 11. Advanced tuning parameters</i>		
Parameter	Value	Description
nist.sp800-131a.strict	The default value is false.	Specifies whether nist.sp800-131a.strict mode is enabled.  CAUTION: A value of true causes you to lose access to the appliance local management interface if your browser does not support TLS 1.2.
gw_net.tuning.downdelay	The default value is 0.	Specifies the time, in milliseconds, to wait before disabling a slave after a link failure is detected. The gw_net.tuning.downdelay value must be a multiple of the gw_net.tuning.miimon value; if not, it is rounded down to the nearest multiple. If your switches take a long time to go into backup mode, it might not be desirable to activate a backup interface immediately after a link goes down. It is possible to delay the moment at which a link is disabled by passing the module parameter downdelay.

Table 11. Advanced tuning parameters (continued)

Parameter	Value	Description
gw_net.tuning.miimon	The default value is 100.	<p>Specifies the MII link monitoring frequency in milliseconds.</p> <p>High availability is achieved by using MII status reporting. The bonding driver can regularly check all its slaves links by checking the MII status registers. This parameter determines how often the link state of each slave is inspected for link failures.</p> <p>A value of 0 disables MII link monitoring. A value of 100 is typically a suitable value. It means that a dead link will be detected 100 milliseconds at most after it goes down. The value must not come too close to 1000/HZ (10 ms on i386) because such setting might reduce the system interactivity.</p>
gw_net.tuning.updelay	The default value is 0.	<p>Specifies the time, in milliseconds, to wait before enabling a slave after a link recovery is detected.</p> <p>The gw_net.tuning.updelay value must be a multiple of the gw_net.tuning.miimon value; if not, it is rounded down to the nearest multiple.</p> <p>When a switch restarts, it is possible that its ports report "link up" status before they become usable. This behavior might cause a bond device to use some ports that are not ready yet. It is possible to delay the moment at which an active link is reused by passing the module parameter gw_net.tuning.updelay (in milliseconds, must be a multiple of gw_net.tuning.miimon).</p> <p>A similar situation can occur when a host renegotiates a lost link with the switch (in case of cable replacement).</p> <p>A special case is when a bonding interface loses all slave links. Then, the driver immediately reuses the first link that goes up, even if gw_net.tuning.updelay parameter was specified. If there are slave interfaces in the gw_net.tuning.updelay state, the interface that first went into that state is immediately reused. This setting reduces downtime if the value of gw_net.tuning.updelay was overestimated.</p>

Table 11. Advanced tuning parameters (continued)

Parameter	Value	Description
gw_net.tuning.use_carrier	0, 1(default)	<p>Specifies whether gw_net.tuning.miimon uses MII / ETHTOOL ioctls, or <code>netif_carrier_ok()</code> to determine the link status. The MII / ETHTOOL ioctls are less efficient and use a deprecated calling sequence within the kernel. The <code>netif_carrier_ok()</code> relies on the device driver to maintain its state with <code>netif_carrier_on/off</code>. Most, but not all, device drivers support this facility.</p> <p>If bonding insists that the link is up when it cannot be, the cause might be that your network device driver does not support <code>netif_carrier_on/off</code>. The default state for <code>netif_carrier</code> is "carrier on". So if a driver does not support <code>netif_carrier</code>, it appears as if the link is always up. In this case, setting gw_net.tuning.use_carrier to 0 causes bonding to revert to the MII / ETHTOOL ioctls method to determine the link state.</p> <p>A value of 1 enables the use of <code>netif_carrier_ok()</code>. A value of 0 specifies to use the deprecated MII / ETHTOOL ioctls. The default value is 1.</p>

Table 11. Advanced tuning parameters (continued)

Parameter	Value	Description
gw_net.tuning.xmit_hash_policy	layer2 (default), layer2+3, layer3+4	<p>Selects the transmit hash policy to use for slave selection in balance-xor, 802.3ad, and tlb modes. Here are the possible values:</p> <p>layer2</p> <p>Uses XOR of hardware MAC addresses and packet type ID field to generate the hash. The formula is as follows:</p> <ul style="list-style-type: none"> • hash = source MAC XOR destination MAC XOR packet type ID • slave number = hash modulo slave count <p>This algorithm places all traffic to a particular network peer on the same slave.</p> <p>This algorithm is 802.3ad compliant.</p> <p>layer2+3</p> <p>This policy uses a combination of layer2 and layer3 protocol information to generate the hash. It uses XOR of hardware MAC addresses and IP addresses to generate the hash. The formula is as follows:</p> <ul style="list-style-type: none"> • hash = source MAC XOR destination MAC XOR packet type ID • hash = hash XOR source IP XOR destination IP • hash = hash XOR (hash RSHIFT 16) • hash = hash XOR (hash RSHIFT 8) • hash = hash Modulo (bonding_slave_count) <p>If the protocol is IPv6, then the source and destination addresses are first hashed by using ipv6_addr_hash.</p> <p>This algorithm places all traffic to a particular network peer on the same slave. For non-IP traffic, the formula is the same as for the layer2 transmit hash policy.</p> <p>This policy is intended to provide a more balanced distribution of traffic than layer2 alone, especially in environments where a layer3 gateway device is required to reach most destinations.</p> <p>This algorithm is 802.3ad compliant.</p> <p>layer3+4</p> <p>This policy uses upper layer protocol information, when available, to generate the hash. This allows for</p>

Managing snapshots

Use snapshots to restore prior configuration and policy settings to the appliance. Back up the appliance on a frequent basis by downloading snapshot files.

About this task

Snapshots are stored on the appliance. However, you can download snapshots to an external drive in case of system failure.

Note: The snapshot files do not contain the internal user registry data. Use standard LDAP back-up tools, using port 636 on the appliance, to back-up and restore the data associated with the internal user registry.

Procedure

1. Click **System** > **System Settings** > **Snapshots**.
2. In the Snapshots pane, use one or more of the following commands:

Option	Description
New	To create a snapshot, click New , type a comment that describes the snapshot, and then click Save .
Edit	To edit the comment for a snapshot, select the snapshot, click Edit , type a new comment, and then click Save .
Delete	To delete snapshots, select one or more snapshots, and then click Delete .
Apply	To apply a snapshot, select the snapshot, and then click Apply . Note: The password of the 'admin' user is not contained in a snapshot. Therefore the password of the 'admin' user will remain unchanged after the application of a snapshot.
Download	To download a snapshot, select the snapshot, click Download , browse to the drive where you want to save the snapshot, and then click Save . Note: If you download multiple snapshots, the snapshots are compressed into a .zip file.
Upload	To upload snapshots, click Upload , browse to the snapshots you want to upload and select the snapshots. Wait for the Comment field of the Upload Snapshot window to be populated automatically. When the Comment field is populated, click Save Configuration . Note: You can upload only one snapshot at a time.
Refresh	To refresh the list of snapshots, click Refresh .

Managing support files

IBM Customer Support uses support files to help you troubleshoot problems with the appliance. Support files contain all log files, temporary and intermediate files, and command output that is needed to diagnose customer support problems.

About this task

Support files might contain customer-identifiable information, such as IP addresses, host names, user names, and policy files. Support files might also contain confidential information, such as passwords, certificates, and keys. The support file contents are stored as a .zip file. All files inside the support file can be inspected and censored by the customer.

Tip: You can create multiple support files to track an issue over time.

Procedure

1. Click **System > System Settings > Support Files**.
2. In the Support Files pane, use one or more of the following commands:

Option	Description
New	To create a support file, click New , select the categories and instances to include in the support file, optionally enter a comment that describes the support file, and then click Save Configuration . A new support file is created on the appliance.
Edit	To edit the comment for a support file, select the support file, click Edit , type a new comment, and then click Save .
Delete	To delete a support file, select the support file, and then click Delete .
Download	To download support files, select the support files, click Download , browse to the drive where you want to save the support files, and then click Save . Note: If you download multiple support files, the files are compressed into a .zip file.

Configuring system alerts

Configure where you want the system to send notifications about changes to system settings and problems with the system.

About this task

Available alerts include system alerts pre-defined in the system and any alert objects that you created.

Procedure

1. Click **System > System Settings > System Alerts**.
2. In the System Alerts pane, complete one or more of the following tasks:
 - To receive notifications for problems with the system, select one or more system alert objects from the Available Objects pane, and add them.
 - To create or edit alert objects, see these related topics to configure one or more of the following alert objects:
 - [“Configuring email alert objects” on page 102](#)
 - [“Configuring remote syslog alert objects” on page 103](#)
 - [“Configuring SNMP alert objects” on page 101](#)
 - To delete a system alert, select the alert and then click **Delete**.

Configuring SNMP alert objects

Configure SNMP alert objects to enable the system to send system alerts to an SNMP Manager.

Procedure

1. Click **System > System Settings > System Alerts**.
2. In the **System Alerts** page, take one of the following actions:
 - Click **New > SNMP**.
 - Select an existing object, and then click **Edit**.
3. Type a name for the alert object.
4. Select a trap version from the list.

5. In the SNMP Manager box, type the IP address, host name, or fully qualified domain name (FQDN) of the SNMP manager.

Note: The SNMP host must be accessible to the appliance to send SNMP traps.

6. Type the port number that the SNMP manager monitors for notifications.

Note: The default port number is 162.

7. Type a comment to describe the SNMP alert object.

8. For trap versions V1 or V2c, type the name of the community that is used to authenticate with the SNMP agent.

9. For trap version 3, configure the following options:

Option	Description
Name	Type the user name to be authenticated in the SNMP database.
Notification Type	On the Notification Type tab, select Inform or Trap in the SNMP Trap Version field.
Authentication	On the Authentication and Privacy tab, select Enabled to enable authentication, type the authentication passphrase, and then select an authentication type.
Privacy	Select Enabled to enable privacy, type the privacy passphrase, and then select a privacy type.

10. Click **Save**.

Configuring email alert objects

You can create email alert objects to send an email notification to specified users or to administrators when specified events occur on your network. You can also select the event parameters to include in the message so that important information about detected events is provided.

Procedure

1. Click **System > System Settings > System Alerts**.
2. In **System Alerts** page, take one of the following actions:
 - Click **New > Email**.
 - Select an existing object, and then click **Edit**.
3. Configure the following options:

Option	Description
Name	Specifies a meaningful name for the response. Note: This name displays when you select responses for events, so give the response a name that allows users to easily identify what they are selecting.
From	Specifies the email address that displays in the From field of the alert email.
To	Specifies the email address or group of addresses to receive the alert. Note: Separate individual email addresses with a comma or semicolon.
SMTP Server	Specifies the fully qualified domain name or IP address of the mail server.

Option	Description
	Note: The SMTP server must be accessible to the appliance to send email notifications.
SMTP Port	Specifies the custom port that is used to connect to the SMTP server. The default is 25.
Comment	Type a comment to identify the email alert object.

4. Click **Save**.

Configuring remote syslog alert objects

Configure remote syslog alert objects to enable the system to record system events in a remote log file.

Procedure

1. Click **System > System Settings > System Alerts**.
2. In the **System Alerts** page, do one of the following steps:
 - Click **New > Remote Syslog**.
 - Select an existing remote syslog alert object, and then click **Edit**.
3. Configure the following options:

Option	Description
Name	Specifies a meaningful name for the response.
Remote Syslog Collector	Specifies the fully qualified domain name or IP address of the host on which you want to save the log. Note: The host must be accessible to the appliance.
Remote Syslog Collector Port	Specifies the custom port that is used to connect to the syslog collector. The default is 514.
Comment	Type a comment to identify the remote syslog alert object.

4. Click **Save**.

Muting selected system alert events

You can configure a list of event IDs for events which should not generate alerts.

Events can be muted on a per-alert object type or per-alert object basis by providing a list of event IDs or patterns which should be ignored. To configure the list of ignored patterns, create a new advanced tuning parameter with the following syntax:

```
events.response.<object type>.ignored = <patterns>
events.response.<object type>.ignored.<object instance> = <patterns>
```

Object Types

<object type> corresponds to the type of System Alert object. A list of the mechanisms and their types follows:

Mechanism	<type>
Event log	logdb
SNMP	snmp

Mechanism	<type>
Email	email
Remote Syslog	syslog

Note: Event Log refers to the default mechanism, which is used to populate the System Events that is seen on the **Event Log** page in the Local Management Interface and the System Events Log REST APIs.

Object Instances

<object instance> refers to the UUID of a specific alert object instance. This can be used to mute events for just a particular System Alert object instance rather than for all objects of a given type.

The **<object instance>** can be discovered by using the REST APIs. Refer to **Manage: System Settings > System Settings > System Alerts > Event Log** within the Web Services documentation.

Patterns

<patterns> are a list of space separated event IDs or patterns which should be ignored.

Each pattern can be either one of the following:

1. A literal event ID. For example, GLGPL1002I
2. A regular expression pattern. For example, (.*)I or GLG(.*)

Examples

To mute the events GLGSY0102I and GLGPL1001I on an SNMP alert object with the UUID 2a6cc325-2d98-4747-85f0-8d7bcac4daec:

Tuning Parameter	Value
events.response.snmp.ignored. 2a6cc325-2d98-4747-85f0-8d7bcac4daec	GLGSY0102I GLGPL1001I

To mute all informational events (events which end with an I character) and GLGAU0003W (user failed to login) on all SNMP alert objects:

Tuning Parameter	Value
events.response.snmp.ignored	(.*)I GLGAU0003W

To mute all events related to snapshots (events which begin with GLGSS) on the System Event Log:

Tuning Parameter	Value
events.response.logdb.ignored	GLGSS(.*)

For more information on Event IDs and their meaning, see [Appliance Messages](#) and [Events that are generated by the events framework](#).

Restarting or shutting down the appliance

Use the **Restart or Shut down** page to restart or shut down the appliance.

About this task

Important: When the appliance is restarting or shutting down, traffic is not passed through the appliance and your network might not be protected.

This page is not available in the LMI when the appliance runs in a Docker environment.

Procedure

1. Click **System > System Settings > Restart or Shut down**
2. Perform one of the following tasks:

Option	Description
Click Restart to restart the appliance	Restarting the appliance takes it offline for several minutes.
Click Shut down to turn off the appliance	Shutting down the appliance takes it offline and makes it inaccessible over the network until you restart it.

3. Click **Yes**.

Configuring application database settings

Configure auto updating and feedback for application databases. Application databases store classifications for web applications and websites.

About this task

To receive updates to application and IP reputation databases, you must enable auto updating. You cannot manually update application and IP reputation databases.

Procedure

1. Click **System > Updates and Licensing > Application Database Settings**.
2. Enable or disable the following options for updating application databases:

- Auto Update
- Enable Feedback

The system classifies a URL as unknown if it is not listed in the application database. Enable the Feedback option to submit unknown URLs and statistics about web application matching to IBM. IBM will classify unknown URLs and include them in a subsequent database update.

IBM uses statistics about matched web applications and actions to continuously improve the classification quality and match ratio for web applications. Feedback data does not include any personal or confidential information about your network.

3. Enable or disable the following options for the IP reputation database:

- Auto Update
- Enable Feedback

Enable the feedback option to submit statistical data to IBM that can make your IP reputation classifications more accurate. This data does not include any personal or confidential information about your network.

- Include IP reputation info

Enable inclusion of IP reputation information in the security events. When disabled, the appliance does not perform IP reputation lookup for security events.

4. Optional: If you use a proxy server, configure the following proxy settings:

Option	Description
Use Proxy	Enables the appliance to use a proxy server for application databases.
Server Address	The IP address or DNS name of the proxy server. Note: The Server Address field is displayed when you select the Use Proxy check box.

Option	Description
Port	The port number that the proxy server uses to communicate with the update server. Note: The Port field is displayed when you select the Use Proxy check box.
Use Authentication	Enables the appliance to authenticate to a proxy server.
User Name	User name required for authenticating to the proxy server. Note: The User Name field is displayed when you select the Use Authentication check box.
Password	Password required for authenticating to the proxy server. Note: The Password field is displayed when you select the Use Authentication check box.

Setting the locale of application log files

Use the **Application Locale** management page to set the locale in which the application log files are written.

Procedure

1. From the top menu, select **System > System Settings > Application Locale**.
2. Select the language that you want the application log files to be written in.
3. Click **Save**.

Configuring SNMP monitoring

Configure SNMP Monitoring so that you can monitor the status of the appliance with a monitoring solution that supports Simple Network Management Protocol. You can monitor the appliance in an IBM Tivoli® Monitoring environment.

About this task

The SNMP Monitoring page is not available in the LMI when the appliance runs in a Docker environment.

Use the Agentless Monitoring for Linux OS agent to monitor the appliance with IBM Tivoli Monitoring.

For more information about configuring the IBM Tivoli Monitoring environment and the Agentless Monitoring for Linux OS agent, see the [IBM Tivoli Monitoring Knowledge Center](#).

The following management information bases, or MIBs, are used by the SNMP agent:

- SNMPv2-MIB
- TCP-MIB
- SNMPv2-SMI
- UDP-MIB
- SNMP-FRAMEWORK-MIB
- HOST-RESOURCES-MIB
- SNMP-MPD-MIB
- MTA-MIB
- SNMP-TARGET-MIB
- DISMAN-EVENT-MIB
- SNMP-USER-BASED-SM-MIB

- NOTIFICATION-LOG-MIB
- SNMP-VIEW-BASED-ACM-MIB
- UCD-SNMP-MIB
- IF-MIB
- UCD-DLMOD-MIB
- IP-MIB
- UCD-DISKIO-MIB
- IPV6-MIB
- UCD-SNMP-MIB
- IP-FORWARD-MIB
- NET-SNMP-AGENT-MIB
- NET-SNMP-VACM-MIB

Procedure

1. From the top menu, select **Manage System Settings > System Settings > SNMP Monitoring**.
2. Type the port number that the SNMP agent must listen on in the **Port** field.

Note: The default port number is 161.

3. Select the **SNMP Protocol** that the agent must use.

- **SNMPv1/SNMPv2c**

Type the name of the community that the SNMP uses to authenticate with the SNMP agent.

- **SNMPv3**

Configure the following options to describe the user who accesses the SNMP agent.

Security Level

Select the security level of the user.

User Name

Type the name of the user who accesses the SNMP agent.

Auth Protocol

Select the authentication protocol to use.

Auth Password

Type the password to use for authentication.

Confirm Auth Password

Type the password to use for authentication.

Priv Protocol

Select the privacy protocol to use.

Priv Password

Type the password to be used as a privacy passphrase.

Confirm Priv Password

Type the password to be used as a privacy passphrase.

4. Click **Save**.

Configuring password quality

IBM Security Verify Access makes use of the PAM password quality checking module (`pam_pwquality`) for accounts which are used to access the local management interface.

For IBM Security Verify Access environments established on version 10.0.0 or newer, the default password quality policy is:

Advanced tuning parameter	Value
password.policy	minlen=8 dcredit=1 ucredit=1 lcredit=1

For IBM Security Verify Access environments established on earlier versions, password quality checking is not performed unless the password.policy tuning parameter is added manually.

When Password Quality checking is performed

Password quality checking is performed for the default admin account during any password change operation or for any System Account when the account is created or a password change operation is taking place.

Events which set a password using non-interactive methods such as silent configuration or bootstrapping processes when deploying in cloud environments are not subject to the password quality checking.

Configuring Password Strength Rules

The password quality policy is configured by setting or modifying the Advanced Tuning Parameter password.policy. The expected format of this parameter is a series of key-value pairs corresponding to pam_pw quality options.

To disable password quality checking, remove the Advanced Tuning Parameter password.policy.

Supported options

The following options from the pam_pw quality module can be used when authoring a password policy:

- minlen
- dcredit
- ucredit
- lcredit
- ocredit
- minclass
- maxrepeat
- maxclassrepeat

Note: Dictionary-based checking is not supported.

Auxiliary Configuration Files

Auxiliary configuration files allow administrators to manage files used by Verify Access to provide static configuration.

Procedure

1. From the top menu, select **Manage System Settings > System Settings > Auxiliary Configuration Files**.

The displayed directories contain the configuration files which are available.

Note: On initial install there might be no auxiliary configuration files listed. The files are automatically created and added to the auxiliary configuration files section as they are required.

2. Click **Refresh** to get the most up-to-date data.
3. Select the required file.
4. Click **Download** to save the file to your local drive.
5. Confirm the save operation in the browser window that pops up.

Secure settings

Information about managing secure settings on your appliance.

Managing SSL certificates

In the local management interface, go to **System > Secure Settings > SSL Certificates**.

The appliance local management interface supports the following authentication mechanisms:

- Forms authentication (UI only)
- Basic authentication (Web services only)
- Client certificate (UI and Web services)

The server uses the certificates that are found in the `lmi_trust_store` certificate database when it authenticates a client certificate. Therefore, to successfully authenticate against the server, the certificate database must contain either the client certificate itself, or the certificate of the CA that signed the client certificate.

Note: As a prerequisite for client certificate authentication, you must configure your browser to trust the CA for the appliance server certificate. In addition, the URL in the request must match the domain name of the appliance.

Configuring SSL connections

Configure Secure Socket Layer (SSL) connections to enable encrypted communication between the LDAP policy information point (PIP) and the LDAP Server to ensure that LDAP traffic is secure and confidential.

About this task

After you import a server certificate, the appliance can authenticate with the LDAP server. For more information, see [“Managing SSL certificates” on page 109](#).

Procedure

1. Log in to the local management interface.
2. Select **System > Secure Settings > SSL Certificates**.
3. Import the LDAP server certificate into the trust store of the runtime profile.
For example: `rt_profile_keys`.

Listing current certificate database names

To list all current certificate database names with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. You can view all current certificate database names and their last modified time information.

Adding description to a certificate database

To add a description to a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database that you want to describe.

3. Select **Manage > Describe**.
4. In the Describe SSL Certificates Database window, enter the description of the certificate database.
5. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Creating a certificate database

To create a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. From the menu bar, click **New**.
3. On the **Create SSL Certificate Database** page, enter the name of the certificate database that you want to create. The name of the certificate database name must be unique.
4. Select the type of the certificate database.
 - If you select **Local** as the type, you can go to Step 5.
 - If you select **Network** as the type, complete the following fields:
 - a. On the **Main** tab, fill in the **Token Label** and **Passcode** fields.
 - b. Select the HSM type.
 - If you select **nCipher nShield Connect** as the HSM type, complete the following fields:
 - i) On the **HSM** tab, the **HSM IP Address** field for the primary HSM device is required. The rest of the fields are optional. You can also provide details of a secondary HSM device. The secondary device can be used for load balancing and failover.
 - ii) On the **RFS** tab, if you select **Automatic**, enter the address of the remote file system that stores the key files. The rest of the fields are optional. If you select **Manual Upload**, click **Browse** to select the zip file that contains the required key files. The contents of the zip file will be extracted and stored on the local file system.

Note:

- The nCipher nShield Connect integration is only available if you first install the 'IBM Security Verify Access nCipher nShield Connect HSM Extension'. This extension is available for download from the IBM Security App Exchange (<https://exchange.xforce.ibmcloud.com/hub/IdentityandAccess>).
- If the files in the remote file system are changed and you selected the **Manual Upload** option, you must manually upload an updated zip file. The updated zip file overwrites existing file entries but does not delete "missing" file entries.
- If you select **SafeNet Luna SA** as the HSM type, complete the **IP Address** and **Admin Password** fields on the SafeNet tab.

Note: The SafeNet integration is only available if you first install the 'IBM Security Verify Access SafeNet Luna Network HSM Extension'. This extension is available for download from the IBM Security App Exchange (<https://exchange.xforce.ibmcloud.com/hub/IdentityandAccess>). You can then use the appliance to manage the certificates that are contained on the HSM device. However, some operations, such as certificate extract, are not supported.

5. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Renaming a certificate database

To rename a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database that you want to rename.
3. Select **Manage > Rename**
4. In the Rename SSL Certificates Database window, enter the new name of the certificate database. The new name of the certificate database name must be unique.
5. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Importing a certificate database

To import a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select **Manage > Import**.
3. Click **Browse** under **Certificate Database File**.
4. Browse to the directory that contains the file to be imported and select the file. Click **Open**.
5. Click **Browse** under **Stash File**.
6. Browse to the directory that contains the file to be imported and select the file. Click **Open**.
7. Click **Import**.

A message that indicates successful import is displayed.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Exporting a certificate database

To export a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database that you want to export.
3. Select **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

4. Confirm the save operation when the browser prompts you to save the .zip file.

Deleting a certificate database

To delete a certificate database with the local management interface, you can use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database that you want to delete.
3. Select **Delete**
4. In the window that pops up, click **Yes**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Replicating the certificate databases across the cluster

If your appliance is the primary master of a cluster environment, you can replicate the certificate databases across the cluster with the **SSL certificate** management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Click **Replicate with Cluster** to have the certificate databases automatically replicated across the cluster.

Note: This option is available only if the current appliance is the primary master of a cluster. If this option is selected, you cannot modify the certificate databases on any appliance other than the primary master.

Managing signer certificates in a certificate database

To manage signer certificates in a certificate database, you can use the SSL Certificates management page. In particular, you can import, export, or delete signer certificates, and list all signer certificate names.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database of interest.
3. Select **Manage > Edit SSL Certificate Database**.
4. All signer certificate names are displayed on the **Signer Certificates** tab.

Import a signer certificate

- a. Click **Manage > Import**.
- b. Click **Browse**. Then, select the signer certificate to be imported.
- c. In the **Certificate Label** field, enter what you want to label the signer certificate.
- d. Click **Import**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

View and export a signer certificate

- a. Select the signer certificate that you want to view.
- b. Click **Manage > View**. The content of the signer certificate is displayed in the browser.
- c. *Optional:* Click **Export**. Then, confirm the save operation in the window that pops up.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

Export a signer certificate

- a. Select the signer certificate that you want to export.
- b. Click **Manage > Export**.
- c. Confirm the save operation in the browser window that pops up.

Delete a signer certificate

- a. Select the signer certificate that you want to delete.
- b. Click **Delete**.
- c. In the window that pops up, click **Yes**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Load a signer certificate from a server

Use the Load function to retrieve a server certificate from the specified server and port, and then install this certificate into the keyfile as a signer certificate with a specific label.

- a. Click **Manage > Load**.
- b. In the **Load Signer Certificate** window, specify the following fields:

Server

The server name from which to load the certificate.

Port

The port from which to load the certificate.

Certificate Label

The name to give to the certificate.

- c. Click **Load**.

Managing personal certificates in a certificate database

To manage personal certificates in a certificate database with the local management interface, use the SSL Certificates management page.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database of interest.
3. Select **Manage > Edit SSL Certificate Database**.
4. Click the **Personal Certificates** tab. All personal certificate names are displayed on this tab.

Note: If the **Issuer** or **Subject** field contains characters in a language other than English, these characters might be displayed in the panel as encoded characters.

Import a personal certificate

- a. Click **Manage > Import**.
- b. Click **Browse**. Then, select the file that contains the personal certificate to import.

Note: Any PKCS 12 file to be imported must have the file extension `.p12` for the import operation to be successful.

- c. *Optional:* Specify the password for the file that contains the personal certificate to import.
- d. Click **Import**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Receive a personal certificate

Note: A personal certificate can be received only if a corresponding certificate request exists.

- a. Click **Manage > Recieve**.
- b. Click **Browse**. Then, select the personal certificate to be received.
- c. Select the **Default** check box if you want to set the personal certificate as default.
- d. Click **Receive**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

View a personal certificate

- a. Select the personal certificate you want to view.
- b. Click **Manage > View**. The content of the personal certificate is displayed in the browser.
- c. *Optional:* Click **Export**. Then, confirm the save operation in the window that pops up.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

Export a personal certificate

- a. Select the personal certificate that you want to export.
- b. Click **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

- c. Confirm the save operation in the browser window that pops up.

Extract a personal certificate

Note: The **Extract** option is used to export a single certificate and its private key (if one exists) from the current key database to a new pcks12 formatted key database.

- a. Select the personal certificate that you want to extract.
- b. Click **Manage > Extract**.
- c. In the Extract Personal Certificate window, enter a password for the extracted certificate container and confirm the password.
- d. Click **Extract**.

Note: You might want to save the certificate with the .p12 file extension for later use. Any PKCS 12 file to be imported must have the file extension .p12 for the import operation to be successful.

Delete a personal certificate

- a. Select the personal certificate that you want to delete.
- b. Click **Delete**.
- c. In the window that pops up, click **Yes**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Create a personal certificate (self-signed)

- a. Click **New**.
- b. Enter **Certificate Label**, **Certificate Distinguished Name**, **Key Size**, and **Expiration Time**. The default value for **Expiration Time** is 365 days.

A distinguished name must be in the following format:

```
CN=cName, OU=orgUnit, O=org, L=city, S=state, C=countryCode
```

S= and ST= might be used for specifying state. However, the certificate or the certificate request always uses ST= .

Any error in the distinguished name results in Error box with CTGSK3024W Invalid value for parameter "-dn" (<entered dn-value>) on save

- c. Optionally, select an entry from the **Signature Algorithm** list. If this option is not specified, the default signature algorithm is used.
- d. Select the **Default** check box if you want to set this personal certificate as the default certificate.
- e. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Set a personal certificate as default

- a. Select the personal certificate that you want to edit.
- b. Click **Edit**.
- c. Select the **Set as the Default Certificate** check box to set the personal certificate as the default certificate.
- d. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Managing certificate requests in a certificate database

To manage certificate requests in a certificate database with the local management interface, use the SSL Certificates management page. In particular, you can create, view, export, or delete certificate requests, and list all certificate request names.

Procedure

1. From the top menu, select **System > Secure Settings > SSL Certificates**.
2. Select the certificate database of interest.
3. Select **Manage > Edit SSL Certificate Database**.
4. Click the **Certificate Requests** tab. All certificate request names are displayed on this tab.

Create a certificate request

- a. Click **New**.
- b. Enter **Certificate Request Label**, **Certificate Request Distinguished Name**, and **Key Size**.

A distinguished name must be in the following format:

```
CN=cName, OU=orgUnit, O=org, L=city, S=state, C=countryCode
```

S= and ST= might be used for specifying state. However, the certificate or the certificate request always uses ST= .

Any error in the distinguished name results in Error box with CTGSK3024W Invalid value for parameter "-dn" (<entered dn-value>) on save

- c. Optionally, select an entry from the **Signature Algorithm** list. If this option is not specified, the default signature algorithm is used.
- d. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

View and export a certificate request

- a. Select the certificate request that you want to view.
- b. Click **Manage** > **View**. The content of the certificate request is displayed in the browser.
- c. *Optional:* Click **Export**. Then, confirm the save operation in the window that pops up.

Export a certificate request

- a. Select the certificate request that you want to export.
- b. Click **Manage** > **Export**. The content of the certificate request is displayed in the browser.
- c. Confirm the save operation in the window that pops up.

Delete a certificate request

- a. Select the certificate request that you want to delete.
- b. Click **Delete**.
- c. In the window that pops up, click **Yes**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Managing file downloads

Use the File Downloads management page in the local management interface to access files that are available for download from the appliance.

Procedure

1. From the top menu, select **System** > **Secure Settings** > **File Downloads**.

The displayed directories contain the files that can be downloaded. There are three parent directories:

- `access_control` contains files specific to the IBM Security Verify Access Advanced Access Control offering.

Note: This directory is shown only if Advanced Access Control has been activated.

- `common` contains files that are common across Security Verify Access.
- `isva` contains files specific to IBM Security Verify Access base offering..

Note: This directory is shown only if the base has been activated.

- `federation` contains files specific to the IBM Security Verify Access Federation offering.

Note: This directory is shown only if Federation has been activated.

These parent directories might contain sub-directories for different categories of files.

2. *Optional:* Click **Refresh** to get the most up-to-date data.
3. Select the file of interest.
4. Click **Export** to save the file to your local drive.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be downloaded.

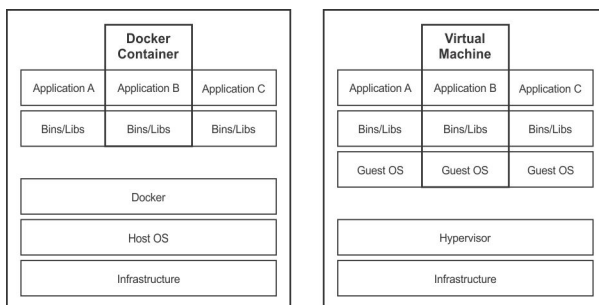
5. Confirm the save operation in the browser window that pops up.

Chapter 8. Docker support

Security Verify Access can run in a Docker environment.

Docker vs virtual machine

Compared to traditional virtual machines, Docker containers are more light-weight. Virtual machines are an abstraction of physical hardware turning one server into many servers. Each virtual machine includes a full copy of the OS, one or more applications, necessary binaries and libraries. As a result, a typical virtual machine image might take up tens of GBs and can be slow to start. Docker containers are an abstraction at the application layer that packages code and dependencies together. Multiple containers can run on the same machine and share the OS kernel with other containers, each running as isolated processes in user space. Containers take less space than virtual machines (container images are typically tens of MBs in size), and can start almost instantly.



Docker principles

Security Verify Access Docker support was implemented with the following Docker principles in mind.

- Containers are supposed to be ephemeral.

Design them in a way that you can stop and destroy an old container and build a new one with an absolute minimum of set-up and configuration.

- Minimize the images.

To reduce complexity, dependencies, file sizes, and build times, avoid installing extra or unnecessary packages. For example, do not include a text editor in a database image.

- Single service.

Decoupling applications into multiple containers makes it much easier to scale horizontally and reuse containers. For instance, a web application stack might consist of three separate containers, each with its own unique image to manage the web application, database, and an in-memory cache in a decoupled manner.

These principles are guidelines from Docker. For more information, see the [Best practices for writing Dockerfiles](#) topic on the Docker website.

Docker terms

The following paragraphs explain some of the common Docker terms used throughout this document.

Image

Docker images are the basis of containers. An Image is an ordered collection of root filesystem changes and the corresponding execution parameters for use within a container runtime. An image typically contains a union of layered file systems stacked on top of each other. An image does not have state and it never changes.

Container

A container is a runtime instance of a Docker image. A Docker container consists of:

- A Docker image
- An execution environment
- A standard set of instructions

Volume

A volume is a specially-designated directory within one or more containers that bypasses the Union File System. Volumes are designed to persist data, independent of the container's life cycle. For more details, see <https://docs.docker.com/engine/tutorials/dockervolumes/>.

For more Docker terms, see the [Docker Glossary](#) page on the Docker website.

Docker networking

The Docker host manages the networking of the Docker containers. Docker containers that reside on the same Docker host can communicate with each other using the internal Docker network. If a Docker container wishes to expose a service (or port) to machines that are not located on the same Docker host, they need to utilize the port mapping capabilities of the Docker host. This capability allows a port from the Docker container to be mapped to a port on the Docker host.

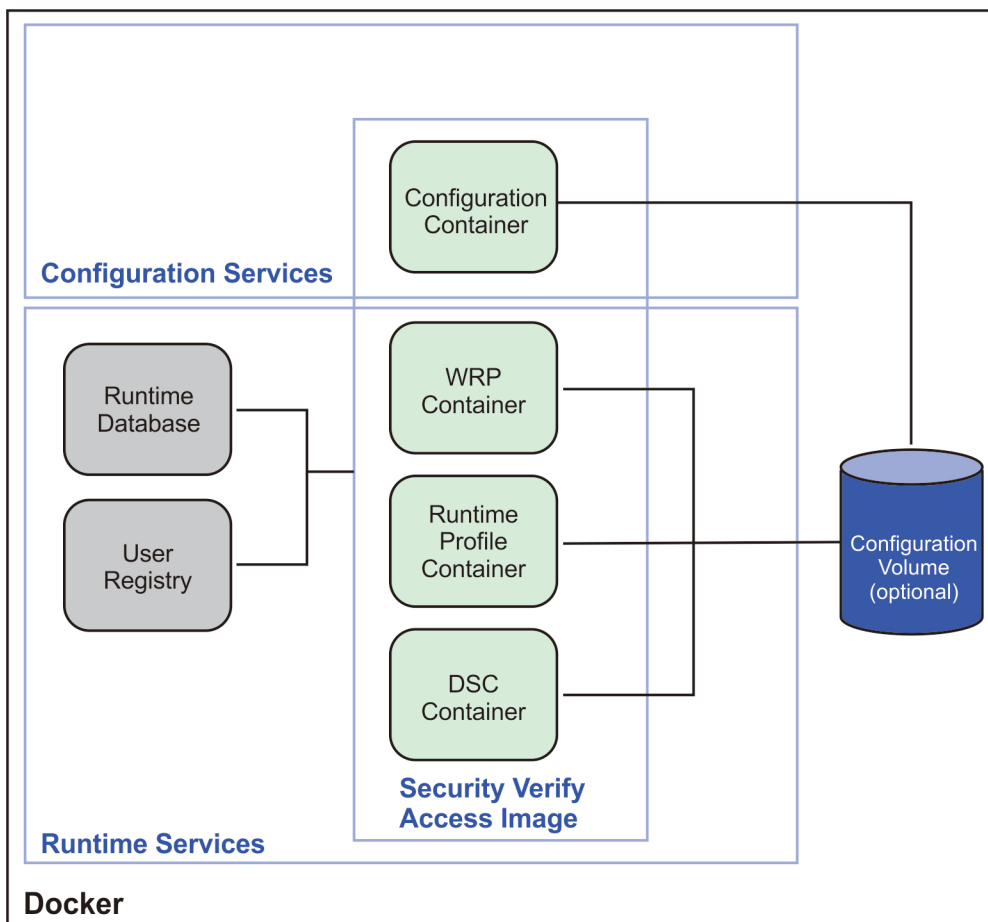
You expose ports using the EXPOSE keyword in the Dockerfile or the `--expose` flag to `docker run`. Exposing ports is a way of documenting which ports are used, but does not actually map or open any ports. Exposing ports is optional.

You publish ports using the PUBLISH keyword in the Dockerfile or the `--publish` flag to `docker run`. This tells Docker which ports to open on the container's network interface. When a port is published, it is mapped to an available high-order port (higher than 30000) on the host machine, unless you specify the port to map to on the host machine at runtime. You cannot specify the port to map to on the host machine in a Dockerfile, because there is no way to guarantee that the port will be available on the host machine where you run the image.

For more information about Docker networking, see the [Docker container networking](#) topic on the Docker website.

Security Verify Access in Docker

The following diagram shows the required elements for Security Verify Access to run in a Docker environment. Each box corresponds to a Docker container.



When Security Verify Access runs in a Docker environment, each container provides a single service, such as configuration, Web Reverse Proxy instance, runtime profile (also known as Advanced Access Control/ Federation), and Distributed Session Cache (DSC). The Security Verify Access Image can run as any one of these four containers (shown in green boxes). Smaller and more efficient images are provided as a recommended alternative to using the main Security Verify Access image for the runtime profile (see [Docker image for Verify Access Runtime](#)), the Web reverse proxy (see [Docker image for Verify Access Web Reverse Proxy](#)) and the distributed session cache (see [Docker image for Verify Access Distributed Session Cache](#)).

The environment also requires an external user registry and database for runtime (for example, DB2, Oracle). The runtime database is required only if you use the Advanced Access Control or Federation capabilities. The external user registry is always required. IBM provides some extensions to third party images that can be used to provide these services. These images (**ibmcom/verify-access-ldap** and **ibmcom/verify-access-postgresql**) are available for download from [Docker Hub](#).

The configuration container is used as a tool to generate the configuration data. The configuration data is shared with the runtime containers through one of the following methods:

- Using a shared volume that has been mounted to the `/var/shared` directory in the container.
- Manually copying the snapshot to the correct location using the Docker commands (the default snapshot file name is: `/var/shared/snapshots/isva_<release_number>_published.snapshot`, for example `/var/shared/snapshots/isva_10.0.0.0_published.snapshot`).
- Using the configuration service that has been exposed from the Security Verify Access configuration container. See [“Configuration service”](#) on page 122.

Images that include all the necessary services to run Security Verify Access in a Docker environment are provided for download.

Table 12. Security Verify Access Docker image sources

Image	Image repository	Image name
IBM Security Verify Access	Docker Hub	ibmcom/verify-access
IBM Security Verify Access Runtime	Docker Hub	ibmcom/verify-access-runtime
IBM Security Verify Access Web Reverse Proxy	Docker Hub	ibmcom/verify-access-wrp
IBM Security Verify Access Distributed Session Cache	Docker Hub	ibmcom/verify-access-dsc
OpenLDAP	Docker Hub	ibmcom/verify-access-openldap
PostgreSQL	Docker Hub	ibmcom/verify-access-postgresql

Docker image for Security Verify Access

The Security Verify Access Docker image can run as a configuration container, a Web Reverse Proxy container, a runtime profile (aka Advanced Access Control/Federation) container, or a Distributed Session Cache (DSC) container. It contains the essential components to get Security Verify Access running on Docker.

Consider the following points when you start a container.

- The docker container should be started as the 'isam' user (UID: 6000). In standard docker environment this will happen automatically but in a Kubernetes environment the security context should be set to allow the container to start as this particular user.
- The following Linux capabilities are required by the container (these capabilities are allowed by default in a standard Docker environment):
 - CHOWN
 - DAC_OVERRIDE
 - FOWNER
 - KILL
 - NET_BIND_SERVICE
 - SETFCAP
 - SETGID
 - SETUID
- The following environment variables are used by the container:

CONTAINER_TIMEZONE

The timezone that will be used by the container (this is a standard Docker environment variable). For example: "Australia/Brisbane"

INSTANCE

The service instance that the container will provide.

In a Web Reverse Proxy container, this environment variable is used to specify the name of the Web Reverse Proxy instance to start. This parameter is required when running in "webseal" mode.

In a Distributed Session Cache container, this environment variable specifies the role of the container (i.e. primary/secondary/tertiary/quaternary) in the format of **INSTANCE = '1|2|3|4'**. For example, to specify that the container acts as the primary, use **INSTANCE = '1'**. To specify that the container acts as the secondary, use **INSTANCE = '2'**.

SERVICE

The service that the container will provide. If no service is specified, the container will default to "config". Valid values are: config, webseal, dsc, and runtime.

Note: The capability of running the verify-access image as anything other than a configuration container is being deprecated and will not be available in verify-access images released after 2021. The new [verify-access-runtime](#), [verify-access-wrp](#), and [verify-access-dsc](#) images should be used instead.

SNAPSHOT

The name of the configuration data snapshot that is to be used when starting the container. This will default to the latest published configuration.

FIXPACKS

A space-separated ordered list of fix packs to be applied when starting the container. If this environment variable is not present, any fix packs present in the fixpacks directory of the configuration volume will be applied in alphanumeric order.

CONFIG_SERVICE_URL

The URL that will be used to access the published configuration/fix-pack data. If using the configuration service of the Security Verify Access configuration container, the URL would be of the format: `https://<container-ip>:<mapped port>/shared_volume`. A BA header will be supplied to handle authentication to the configuration service.

Note: In the configuration container this environment variable is used to specify the service to which a snapshot will be published. When an administrator chooses to publish a snapshot the generated snapshot file will be sent, via a HTTP POST operation, to the specified service.

CONFIG_SERVICE_USER_NAME

The user that will be used when accessing the configuration service.

Note: In the configuration container this environment variable is used to specify the service to which a snapshot will be published. When an administrator chooses to publish a snapshot the generated snapshot file will be sent, via a HTTP POST operation, to the specified service.

CONFIG_SERVICE_USER_PWD

The password for the user that will be used when accessing the configuration service.

Note: In the configuration container this environment variable is used to specify the service to which a snapshot will be published. When an administrator chooses to publish a snapshot the generated snapshot file will be sent, via a HTTP POST operation, to the specified service.

ADMIN_PWD

The initial seeded password for the built in 'admin' user that is used when accessing the configuration service. If this parameter is not specified the default password 'admin' is used.

Note: This environment variable is not available prior to version 9.0.5.0.

AUTO_RELOAD_FREQUENCY

The frequency, in seconds, that the container will check to see if the configuration has been updated. If an updated configuration is detected, the container will automatically reload the configuration data. Note that there will be a service interruption while the reload takes place. If this environment variable is missing, the container will not attempt any automatic reload of configuration data. This environment variable is ignored in the configuration container.

USE_CONTAINER_LOG_DIR

This environment variable, if set to any value, is used to indicate that the log files should be written to a container specific logging directory (underneath the '/var/application.logs' path). This allows multiple container replicas to write log information to the same persistent volume. An alternative, in a Kubernetes environment, is to deploy the containers in a 'StatefulSet'. See the official Kubernetes documentation for information on StatefulSets.

Note: This environment variable is not available prior to version 10.0.0.0

VERIFY_FILES

This environment variable, if set to any value, will cause the container to verify all binary files in the container at start-up to ensure that they have not been modified. If this variable is not set the files will not be checked during the container start-up. By electing to not verify the files the length of time that the container takes to start will be decreased, but this also means that the binary files on the file system do not get verified to ensure that they have not been tampered with.

LANG

The language in which messages which are sent to the console will be displayed. If no language is specified the messages will appear in English. The following table lists the supported languages:

Language	Environment Variable Value
Czech	cs_CZ.utf8
German	de_DE.utf8
Spanish	es_ES.utf8
French	fr_FR.utf8
Hungarian	hu_HU.utf8
Italian	it_IT.utf8
Japanese	ja_JP.utf8
Korean	ko_KR.utf8
Polish	pl_PL.utf8
Portuguese (Brazil)	pt_BR.utf8
Russian	ru_RU.utf8
Chinese (Simplified)	zh_CN.utf8
Chinese (Traditional)	zh_TW.utf8

Consider the following points regarding user registry support when you use Security Verify Access in a Docker environment.

- The embedded user registry can only be used to house the **secAuthority=Default** suffix in conjunction with basic users. If full Security Verify Access users are required, the **secAuthority=Default** suffix must be stored in an external user registry.
- An external user registry is always required for the user suffix. Configure the external user registry as a federated user registry if the embedded user registry is being used for the **secAuthority=Default** suffix.

Configuration service

All configuration must be completed using the configuration container. The configuration service supports a scaled-down version of the LMI. You can use this LMI to manage the configuration data.

Note: To make a configuration available to a runtime container, you must click the **Publish configuration** button in the LMI.

The management service (i.e. the LMI) always listens on port 9443 of the container.

Runtime Management Web Services

The Runtime management Web services are light-weight web services hosted on the runtime container. These Web Services are used to manage the runtime of the container. They cannot be used to manage the configuration data, which should be done with the configuration Web services that is provided by the

configuration container. The runtime management Web services listen on port 9443 of the runtime container.

Currently the runtime Web Services provides the ability to access the CLI and the application log files. For those APIs, the `{appliance_hostname}` parameter can be the configuration container address or the runtime container address. See the following Web Services documentation for the usage:

"Run CLI Command" Web Services

This Web Service acts as a front-end to the command line interface.

"Application Log Files" Web Services

The "Application Log Files" Web Services allows you to access the application log files.

Migrating an appliance to Docker

To migrate your appliance to the Docker environment, you can create a snapshot of the appliance in its original environment and import the snapshot into a running Security Verify Access configuration container.

You can only import a snapshot from an appliance if the following conditions are met:

- For a Security Verify Access Base only activation, the snapshot was taken on version 9.0.0.0 or later. For an Advanced Access Control or Federation activation, the snapshot was taken on version 9.0.2.0 or later.
- The appliance was configured with an embedded configuration database and an external runtime database.
- The appliance runtime environment was using an external LDAP server. Alternatively, if the appliance was running Security Verify Access 9.0.4.0, an embedded LDAP server can be used if the **"wga_rte.embedded.ldap.include.in.snapshot"** advanced tuning parameter was set to `true` before generating the snapshot.

When a snapshot from an appliance is imported to a Docker container:

- The LMI HTTPS listening port will be rewritten to 9443.
- Any reverse proxy instances will have their HTTPS and HTTP ports rewritten to 443 and 80 respectively.

Restrictions

Security Verify Access, when run in a Docker environment, has the following restrictions:

- Any configuration changes require the service containers to be reloaded. You can use the CLI to trigger a manual reload. Changes to the Federation configuration and the policy database will not result in any service down time. Changes to junction definitions and Web Reverse Proxy configuration will result in minimal service down time while the Web Reverse Proxy is restarted. See ["CLI in a Docker environment" on page 157](#).
- The authorization server (i.e. pdaclD) is not supported.
- The front-end load balancer capability of the Security Verify Access appliance is not supported.
- The IP reputation policy information point (PIP) capability of Advanced Access Control is not supported.
- Network HSM devices are not supported. All keys are stored locally.
- A sample geo-location database is not provided. If a sample geo-location database is required, it should be obtained from the downloads area of a running virtual or hardware appliance. See [Updating location attributes](#).
- Pre-installed federation partner templates are not provided. See [Managing federation partner templates](#). The connector package is available from the IBM Security App Exchange site (<https://www.ibm.com/security/community/app-exchange>) as the 'IBM Security Access Manager Extension for SAML Connectors' package.
- Web Reverse proxy flow data or PAM statistics are not supported.

- The embedded user registry can only be used to hold static data and should not be used to hold any user data. As a result the embedded user registry should only be used in conjunction with a federated registry, to store the user data, and basic users. The Security Verify Access integration component of the SCIM support will not be available if the embedded user registry is in use.
- Authentication using RSA SecurID tokens is not supported.
- The container cannot be executed from within a Docker user namespace.
- There are a few differences when managing junctions using the configuration container, namely:
 - Validation of junction server connectivity does not take place when creating a junction.
 - Fine grained authorization checks on junction management operations, and policy object space operations, does not take place. This means that any administrator who is able to authenticate to the policy server (using, for example, pdadmin) is able to manage junctions and the Web Reverse Proxy policy object space.

Shared configuration data

The shared configuration volume is a section of the file system that is reserved for the storage of data to be shared among multiple containers. The data on the shared configuration volume is persisted even if the containers are deleted.

The shared configuration volume is mounted in a Security Verify Access container at `/var/shared` . Snapshots, support files, and fix packs are stored in this volume. To manage these files, you can use the **System > Network Settings > Shared Volume** page of the configuration container LMI.

Snapshots

Snapshots are located in the `snapshots` directory of the configuration volume.

When a snapshot is published from the configuration container, it is stored on the shared volume. When a runtime container is started, it uses the snapshot to perform configuration and bootstrap successfully. Snapshots can only be created using the configuration container, though an administrator can also manually add or remove snapshots by directly accessing the Docker volume.

Support files

Support files are located in the `support` directory of the configuration volume.

Technically, you can create support files in containers of any type. However, support files are most commonly generated in one of the runtime containers. To generate and retrieve a support file in a runtime container, follow these steps:

1. Using the CLI or CLI web service, create a support file in the runtime container. This support file will be visible in the configuration container.
2. If the volume has not been directly mounted in the runtime container and a configuration service has been defined, use the **support -> publish** CLI command to send the snapshot to the configuration service.
3. Using the LMI or web service of the configuration container, retrieve the support file. Alternatively, you can also access the support folder on the Docker volume directly to retrieve the support file.

Fix packs

Fix packs are located in the `fixpacks` directory of the configuration volume.

When a container is started, fix packs that are specified in the **FIXPACKS** environment variable will be applied in the order that they are specified. If the **FIXPACKS** environment variable is not present, any fix packs present in the `fixpacks` directory of the configuration volume will be applied in alphanumeric order.

To manage fix packs, you can either access the Docker volume manually, or use the **System > Network Settings > Shared Volume** page of the configuration container LMI. On the **Shared Volume** page, you can view the contents of the `fixpacks` directory of the configuration volume, upload, delete, or rename fix packs.

The **System > Updates and Licensing > Fixpack** LMI page is read-only in a Docker environment. You can use that page to see which fix packs have been applied, but cannot use it to apply or roll back fix packs.

Log files

By default, Docker uses a layered file system to help reduce the disk space utilization of the Docker containers. However, this file system has slower write speeds than standard file systems. As such, a standard Docker practice is to place any files that are updated frequently (for example, log files) on a shared volume. All of the log files that are used by Security Verify Access are located in the `'/var/application.logs'` directory. Therefore the recommended approach is to create this directory as a shared volume when you create your container.

You can view the log files through the **Monitor > Application Log Files** panel of the LMI.

Multiple containers should not reference the same persistent volume for log storage, otherwise multiple containers will attempt to write to the same log file at the same time, causing data write and integrity issues. In a Kubernetes environment this problem can be overcome by deploying the containers in a StatefulSet (refer to the official Kubernetes documentation for information on StatefulSets). An alternative is to set the `USE_CONTAINER_LOG_DIR` environment variable in the container. When this variable is set the log files are written to a container specific log sub-directory. This environment variable is not available prior to version 10.0.0.

Note: In IBM Security Verify Access version 9.0.7.0 a container specific log sub-directory is always used.

The log file directory structure is shown in the following table.

<i>Table 13. Logs directory structure</i>	
Log file	Sub-directory (relative to the root log directory)
Local management interface log files	lmi
Runtime profile log files	rtprofile
Runtime audit logs	rtaudit
DSC log files	dsc
Security Verify Access policy server log and trace files	isam_runtime/policy
Embedded User Registry log files	isam_runtime/user_registry
Web Reverse Proxy log files	wrp/<instance>/log
Web Reverse Proxy statistic files	wrp/<instance>/stats
Web Reverse Proxy trace files	wrp/<instance>/trace
Web Reverse Proxy transaction files	wrp/<instance>/translog
System log files	system
Remote system log forwarder files	rsyslog_forwarder

The other option is to access the logs with the Web services on the configuration and the runtime containers. By invoking the corresponding "Application Logs" API on each container, the user can list and retrieve the log files on that container. See the Docker Web Services documentation for more information.

Note: The recommended approach is to configure Security Verify Access to send the log files to a remote syslog server wherever possible.

Docker image for PostgreSQL support

The **ibmcom/verify-access-postgresql** image extends the official **postgres** Docker image by adding SSL support and the Security Verify Access schema to the image. This image can be used to quickly deploy a database for use with the Federation and Advanced Access Control offerings of Security Verify Access.

Instructions on the use of the official postgres Docker image can be found at: [Docker Hub](#).

Additional environment variables

In addition to the standard **postgres** environment variables, the **ibmcom/verify-access-postgresql** Docker image defines the following environment variables:

Name	Description
POSTGRES_SSL_KEYDB	The name of the SSL file that contains both the SSL server certificate and key (the key should not be protected by a password). This key file must be made available to the Docker container at start-up. This is usually achieved by placing the key file in a Docker volume and making this volume available to the container.
POSTGRES_UNSECURE	By default unsecure communication with the database server is disabled. If set to the value of 'true', this environment variable will enable unsecure communications with the PostgreSQL server.
POSTGRES_SSL_CN	If a CN value is supplied, a self-signed certificate for the server will be automatically created when the container first starts. The public key will be available from the ' PGDATA/public.pem ' file of the running container.

Usage

Quick start

To start a container with the defaults, execute the command:

```
docker run --name isva-postgresql --detach ibmcom/verify-access-postgresql:latest
```

However, a more complete command, which would specify the volumes, ports and standard environment variables, could be:

```
docker run --hostname isva.postgresql --name isva.postgresql \  
--detach \  
--publish 5432:5432 \  
--volume /var/lib/postgresql/data \  
--env POSTGRES_USER=postgres \  
--env POSTGRES_PASSWORD=password \  
--env POSTGRES_DB=isva \  
--env POSTGRES_SSL_CN=isva.postgresql \  
ibmcom/verify-access-postgresql:latest
```

Security

By default the image will automatically generate a TLS certificate when the container is first started. The CN for the certificate is obtained from the **POSTGRES_SSL_CN** environment variable (if defined),

otherwise it will be obtained from the container hostname. The generated public key will be saved to the `'${PGDATA}/public.pem'` file within the container.

If you want to enable unsecure communication with the database server the `POSTGRES_UNSECURE` environment variable should be set to `'true'`.

If you want to provide your own certificate the public certificate and private key should be placed into a single file (without password protection) and made available to the container during initialization. The location of the key file within the container is defined by the `POSTGRES_SSL_KEYDB` environment variable.

If you want to create your own self-signed server certificate, you can do so using OpenSSL. For example:

```
openssl req -x509 -newkey rsa:4096 \  
-keyout postgres.key -out postgres.crt \  
-days 365 -nodes \  
-subj "/C=AU/ST=Queensland/L=Gold Coast/O=IBM/CN=isva-postgresql"  
cat postgres.key postgres.crt > container.pem
```

User ID

By default the container runs as the 'postgres' (uid: 70) user. If a volume is being used to persist the database, the 'postgres' user must be granted write permission to the volume. In a Kubernetes environment this can be achieved by setting the `fsGroup` field in the deployment yaml file.

License

The Dockerfile and associated scripts are licensed under the [Apache License 2.0](#) license.

Supported Docker versions

- This image is officially supported on Docker version v17 and later.
- Support for older versions is provided on a best-effort basis.

Community support

If you are a licensed IBM customer, you can request support through the official [IBM support channel](#). However, IBM does not provide support for the official **postgres** Docker image.

Community support is also available for this image via the DeveloperWorks communities. Both [DeveloperWorks Answers](#) and the [DeveloperWorks IBM Security Identity and Access Management Forum](#) are vibrant communities.

Supported tags

Tag	Purpose
latest	The latest stable version.
V.R.M.F	A particular release, of the format: <code>{version}.{release}.{modification}.{fixpack}</code> . For example, <code>10.0.0.0</code>

Related information

Docker image for OpenLDAP support

The **ibmcom/verify-access-openldap** image extends the **osixia/openldap** Docker image by adding the Security Verify Access "**secAuthority=Default**" schema and suffix to the registry. This image can be used to quickly build a user registry for use with Security Verify Access.

Instructions on the use of the **osixia/openldap** Docker image can be found at: <https://github.com/osixia/docker-openldap>.

Points to note

Some additional points to note about the extensions to the **osixia/openldap** Docker image include:

- The **secAuthority=Default** suffix is stored in the **"/var/lib/ldap.secAuthority"** directory and so this should be added to the list of volumes of the **osixia/openldap** container.
- Using the **osixia/openldap-backup** Docker container to back-up the user registry is not supported.
- The **secAuthority=Default** suffix will contain the **"cn=root,secAuthority=Default"** administrative user. The password for the user will be set to the same value as the admin user of the **osixia/openldap** container (controlled by the **LDAP_ADMIN_PASSWORD** variable).
- The user suffix is automatically determined from the **LDAP_DOMAIN** entry, where each element in the domain name is preceded by **"dc"**. For example, if the **LDAP_DOMAIN** is set to **ibm.com**, the corresponding suffix will be **"dc=ibm,dc=com"**.
- The default value of the **LDAP_TLS_VERIFY_CLIENT** environment variable has been changed from **'required'** to **'never'**.
- By default, the LDAP server will only listen on the LDAPS secure port (636) and will not listen on the LDAP unsecure port (389).

Additional environment variables

In addition to the standard **osixia/openldap** environment variables, the **ibmcom/verify-access-openldap** Docker image defines the following environment variables:

Name	Description
LDAP_ENABLE_PORT_389	By default, the OpenLDAP server will only listen on the secure 636 port. If you want the OpenLDAP server to also listen on the unsecure 389 port, this environment variable must be set to the value 'true' .

Usage

Quick start

To start a container with the defaults, execute the command:

```
docker run --name isva-openldap --detach ibmcom/verify-access-openldap:latest
```

However, a more complete command, which would specify the volumes, ports and standard environment variables, could be:

```
docker run --hostname isva.openldap --name isva.openldap \  
--detach \  
--publish 636:636 \  
--volume /var/lib/ldap \  
--volume /etc/ldap/slapd.d \  
--volume /var/lib/ldap.secAuthority \  
--env LDAP_DOMAIN=ibm.com \  
--env LDAP_ADMIN_PASSWORD=passwd \
```

```
--env LDAP_CONFIG_PASSWORD=passw0rd \  
ibmcom/verify-access-openldap:latest
```

TLS

By default, the image will automatically generate a TLS certificate when the container is first started. The CN for the certificate is obtained from the container hostname.

If you want to provide your own certificates, they should be made available to the container at initialization within the **/container/service/slapd/assets/certs** directory. The following files reside within this directory:

Filename	Description
ldap.crt	The server certificate to be used.
ldap.key	The private key for the server certificate.
ca.crt	The certificate for the trusted certificate authority, used to validate certificates that are presented to the LDAP server (aka mutual authentication).

If you want to create your own self-signed server certificate, you can do so using OpenSSL. For example:

```
openssl req -x509 -newkey rsa:4096 -keyout ldap.key -out ldap.crt \  
-days 365 -nodes \  
-subj "/C=AU/ST=Queensland/L=Gold Coast/O=IBM/CN=isva-openldap"
```

License

The Dockerfile and associated scripts are licensed under the [Apache License 2.0](#) license.

Supported Docker versions

- This image is officially supported on Docker version v17 and later.
- Support for older versions is provided on a best-effort basis.

OpenShift

To run this container with OpenShift, the service account must be allowed to run processes with elevated privileges. The 'anyuid' scc provides sufficient permission to run the `ibmcom/verify-access-openldap` image. An example service account is:

```
oc create sa openldap  
oc adm policy add-scc-to-user anyuid -z openldap
```

The `apps/v1.spec.template.spec.serviceAccountName` entry should also be set to this service account name. For example, `openldap`.

```
apiVersion: apps/v1  
kind: Deployment  
spec:  
  template:  
    spec:  
      serviceAccountName: openldap  
      containers:  
        . . .
```

Community support

If you are a licensed IBM customer, you can request support through the official [IBM support](#) channel. However, IBM does not provide support for the official **osixia/openldap** Docker image.

Community support is also available for this image via the DeveloperWorks communities. Both [DeveloperWorks Answers](#) and the [DeveloperWorks IBM Security Identity and Access Management Forum](#) are vibrant communities.

Supported tags

Tag	Purpose
latest	The latest stable version.
V.R.M.F	A particular release, of the format: {version}.{release}.{modification}.{fixpack}. For example, 10.0.0.0

Related information

Docker image for Verify Access Runtime

The Security Verify Access Runtime Docker image provides the Advanced Access Control and Federation capabilities of Security Verify Access. It is a light-weight and secure alternative to running the same capabilities using the main Security Verify Access Image.

The main differences between running the Security Verify Access runtime image and the main Security Verify Access image as the runtime persona include:

- The runtime image is more light weight, which means that it consumes less disk space and memory, runs fewer processes and is quicker to start.
- The runtime service listens on port 9443, by default, when using the Security Verify Access Runtime image, whereas it listens on port 443 when using the Security Verify Access image.
- The runtime image does not require any elevated container security capabilities/privileges (for example: the SETUID capability is not required).
- Logging and auditing records are sent to the console in JSON format so that the container logging infrastructure can manage the logging and auditing records. No support is provided for natively forwarding logging messages to a remote syslog server.
- The **isam_cli** program is not supported. The **runtime_reload** command can instead be used to reload the runtime configuration when the configuration snapshot changes.

Note: The definition for the Security Verify Access image and the Security Verify Access Runtime image are interchangeable. In order to migrate to the Security Verify Access Runtime image the image name will need to be changed, along with the port information. The HTTP_PORT and/or HTTPS_PORT environment variables can be set to indicate the ports on which the container will listen but note that if the supplied port number is less than 1024 the NET_BIND_SERVICE capability is required for the container. The rest of the definitions can remain the same. It is however recommended that the Security Verify Access Runtime image is used as it is a more secure and light-weight alternative, and the ability to execute the runtime capabilities from the Security Verify Access image will be deprecated in a future release. Please take note that the logging in the new container will by default be in JSON format, and the SETUID privilege is no longer required.

Consider the following points when you deploy a runtime container:

- The runtime container relies on a configuration snapshot which has been generated by the Security Verify Access configuration container.
- The runtime container should be started as the 'isam' user (UID: 6000). In a standard container environment this will happen automatically but in a Kubernetes environment the security context should be set to allow the container to start as this particular user.
- The following environment variables are used by the container:

CONFIG_SERVICE_URL

The URL that will be used to access the published configuration data. If using the configuration service of the Security Verify Access configuration container, the URL would be of the format: `https://<container-ip>:<mapped-port>/shared_volume`. A BA header will be supplied to handle authentication to the configuration service. If this environment variable is not specified the container will expect the configuration snapshot to be available in the shared configuration volume, at `/var/shared/snapshots`.

CONFIG_SERVICE_USER_NAME

The name of the user that will be used when accessing the configuration service.

CONFIG_SERVICE_USER_PWD

The password for the user that will be used when accessing the configuration service.

HTTP_PORT

The port on which the container will listen for HTTP request (default: 9080).

Note: If the port number is less than 1024 the `NET_BIND_SERVICE` capability is required for the container.

HTTPS_PORT

The port on which the container will listen for HTTPS request (default: 9443).

Note: If the port number is less than 1024 the `NET_BIND_SERVICE` capability is required for the container.

SNAPSHOT

The name of the configuration data snapshot file that is to be used when starting the container. No path information should be included in the filename. The standard snapshot locations will be searched for the specified snapshot file. This variable, if not specified, will default to the latest published configuration.

WLP_LOGGING_CONSOLE_FORMAT

The required format for the log messages. Valid values are `simple` or `json` format (default: `json`)

WLP_LOGGING_CONSOLE_SOURCE

The list of comma-separated sources that route to the console. This property applies only when the console format is JSON. Valid values are `message`, `trace`, `accessLog`, `ffdc`, and `audit` (default: `message`).

WLP_LOGGING_CONSOLE_LOGLEVEL

This filter controls the granularity of messages that go to the console. The valid values are `INFO`, `AUDIT`, `WARNING`, `ERROR`, and `OFF` (default: `AUDIT`).

LANG

The language in which messages which are sent to the console will be displayed. If no language is specified the messages will appear in English. The following table lists the supported languages:

Language	Environment Variable Value
Czech	<code>cs_CZ.utf8</code>
German	<code>de_DE.utf8</code>
Spanish	<code>es_ES.utf8</code>
French	<code>fr_FR.utf8</code>
Hungarian	<code>hu_HU.utf8</code>
Italian	<code>it_IT.utf8</code>
Japanese	<code>ja_JP.utf8</code>
Korean	<code>ko_KR.utf8</code>
Polish	<code>pl_PL.utf8</code>

Portuguese (Brazil)	pt_BR.utf8
Russian	ru_RU.utf8
Chinese (Simplified)	zh_CN.utf8
Chinese (Traditional)	zh_TW.utf8

Configuration

All configuration activities must be completed using the main Security Verify Access image, running as a configuration container. The configuration container supports a scaled-down version of the Security Verify Access appliance LMI. You can use this LMI to manage the configuration data.

Note: To make a configuration available to the runtime container, you must click **Publish configuration** in the LMI.

Service

The container, by default, will listen for incoming requests on port 9443, and optionally port 9080 (if HTTP access has been enabled in the configuration snapshot). Use the HTTPS_PORT and HTTP_PORT environment variables to change the ports on which the container will listen.

Shared configuration data

The shared configuration volume is a section of the file system that is reserved for the storage of data to be shared amongst multiple containers. The data on the shared configuration volume is persisted even if the containers are deleted.

The shared configuration volume is mounted in a Security Verify Access runtime container at '/var/shared'. Configuration snapshots are retrieved from the 'snapshots' sub-directory of this volume if the CONFIG_SERVICE_URL environment variable is not specified.

The configuration data which is used by the Security Verify Access Runtime image is fully compatible with the configuration data used by the legacy Security Verify Access image.

Logging

The logging and auditing of the runtime will, by default, be sent to the console of the container, in JSON format. This allows the logging infrastructure of the container environment itself to manage the auditing and message logs.

Some additional log files will however still be generated on the disk of the container. By default, Docker uses a layered file system to help reduce the disk space utilization of the containers. However, this file system has slower write speeds than standard file systems. As such, a standard Docker practice is to place any files that are updated frequently (for example, log files) on a shared volume. All of the log files that are written by the container are located in the '/var/application.logs.' directory. Therefore, the recommended approach is to create this directory as a shared volume when you create your container.

Note: Multiple containers should not reference the same persistent volume for log storage, otherwise multiple containers will attempt to write to the same log file at the same time, causing data write and integrity issues. In a Kubernetes environment this problem can be overcome by deploying the containers in a StatefulSet (refer to the official Kubernetes documentation for information on StatefulSets).

The log file directory structure is shown in the following table.

Log file	Sub-directory (relative to the root log directory)
Runtime server log files	rtprofile

Table 19. Logs directory structure (continued)	
Log file	Sub-directory (relative to the root log directory)
Database server log files	db

Docker image for Verify Access Web Reverse Proxy

The Security Verify Access Web Reverse Proxy (WRP) Docker image provides the Web Reverse Proxy capabilities of Security Verify Access.

It is a light-weight and secure alternative to running the same capabilities using the main Security Verify Access Image.

The main differences between running the Security Verify Access WRP image and the main Security Verify Access image as the WebSEAL persona include:

- The WRP image is more light weight, which means that it consumes less disk space and memory, runs fewer processes and is quicker to start.
- The WRP service listens on port 9443, by default, when using the Security Verify Access WRP image, whereas it listens on port 443 when using the Security Verify Access image.
- The WRP image does not require any elevated container security capabilities/privileges (for example: the SETUID capability is not required).
- Logging records are sent to the console in JSON format so that the container logging infrastructure can manage the logging records. No support is provided for natively forwarding logging messages to a remote syslog server.
- The 'reload' capability is not supported which means that the container must be restarted in order to pick up configuration changes. In a Kubernetes environment the 'rolling update' capability of Kubernetes should be used to ensure that there is zero downtime when applying a configuration update.

Note: The definition for the Security Verify Access image and the Security Verify Access WRP image are interchangeable. In order to migrate to the Security Verify Access WRP image the image name will need to be changed, along with the port information. The HTTP_PORT and/or HTTPS_PORT environment variables can be set to indicate the ports on which the container will listen but note that if the supplied port number is less than 1024 the NET_BIND_SERVICE capability is required for the container. The rest of the definitions can remain the same. It is however recommended that the Security Verify Access WRP image is used as it is a more secure and light-weight alternative, and the ability to execute the WRP capabilities from the Security Verify Access image will be deprecated in a future release. Please take note that the logging in the new container will by default be in JSON format, and the SETUID privilege is no longer required.

Consider the following points when you deploy a WRP container:

- The WRP container relies on a configuration snapshot which has been generated by the Security Verify Access configuration container.
- The WRP container should be started as the 'isam' user (UID: 6000). In a standard container environment this will happen automatically but in a Kubernetes environment the security context should be set to allow the container to start as this particular user.
- The following environment variables are used by the container:

CONFIG_SERVICE_URL

The URL that will be used to access the published configuration data. If using the configuration service of the Security Verify Access configuration container, the URL would be of the format: `https://<container-ip>:<mapped-port>/shared_volume`. A BA header will be supplied to handle authentication to the configuration service. If this environment variable is not specified the container will expect the configuration snapshot to be available in the shared configuration volume, at `/var/shared/snapshots`.

CONFIG_SERVICE_USER_NAME

The name of the user that will be used when accessing the configuration service.

CONFIG_SERVICE_USER_PWD

The password for the user that will be used when accessing the configuration service.

HTTP_PORT

The port on which the container will listen for HTTP request (default: 9080).

Note: If the port number is less than 1024 the NET_BIND_SERVICE capability is required for the container.

HTTPS_PORT

The port on which the container will listen for HTTPS request (default: 9443).

Note: If the port number is less than 1024 the NET_BIND_SERVICE capability is required for the container.

SNAPSHOT

The name of the configuration data snapshot file that is to be used when starting the container. No path information should be included in the filename. The standard snapshot locations will be searched for the specified snapshot file. This variable, if not specified, will default to the latest published configuration.

INSTANCE

The name of the WRP instance to be started. If no INSTANCE is specified, the instance with the name of 'default' will be used.

LOGGING_CONSOLE_FORMAT

The required format for the log messages. Valid values are 'basic' or 'json' (default: json).

LANG

The language in which messages which are sent to the console will be displayed. If no language is specified the messages will appear in English. The following table lists the supported languages:

Language	Environment Variable Value
Czech	cs_CZ.utf8
German	de_DE.utf8
Spanish	es_ES.utf8
French	fr_FR.utf8
Hungarian	hu_HU.utf8
Italian	it_IT.utf8
Japanese	ja_JP.utf8
Korean	ko_KR.utf8
Polish	pl_PL.utf8
Portuguese (Brazil)	pt_BR.utf8
Russian	ru_RU.utf8
Chinese (Simplified)	zh_CN.utf8
Chinese (Traditional)	zh_TW.utf8

Configuration

All configuration activities must be completed using the main Security Verify Access image, running as a configuration container. The configuration container supports a scaled-down version of the Security Verify Access appliance LMI. You can use this LMI to manage the configuration data.

Note: To make a configuration available to the WRP container, you must click **Publish configuration** in the LMI.

Service

The container, by default, will listen for incoming requests on port 9443, and optionally port 9080 (if HTTP access has been enabled in the configuration snapshot). Use the `HTTPS_PORT` and `HTTP_PORT` environment variables to change the ports on which the container will listen.

Shared configuration data

The shared configuration volume is a section of the file system that is reserved for the storage of data to be shared amongst multiple containers. The data on the shared configuration volume is persisted even if the containers are deleted.

The shared configuration volume is mounted in a Security Verify Access WRP container at `/var/shared`. Configuration snapshots are retrieved from the `snapshots` sub-directory of this volume if the `CONFIG_SERVICE_URL` environment variable is not specified.

The configuration data which is used by the Security Verify Access WRP image is fully compatible with the configuration data used by the legacy Security Verify Access image.

Logging

The logging of the WRP process will, by default, be sent to the console of the container, in JSON format. This allows the logging infrastructure of the container environment itself to manage the message logs.

The WRP configuration controls which auditing records are enabled, and where the auditing records are sent. It is recommended that all auditing records (including the WRP request log) are also written to the console of the container in JSON format. This can be achieved by:

1. Changing the logging agent which is used for the auditing, controlled by the `logcfg` configuration entry, to `stdout`.
2. Enabling JSON auditing, controlled by the `audit-json` configuration entry.
3. Modifying the request log destination so that it is set to `stdout`.
4. Modifying the request log format string so that the data is formatted as JSON.

For example:

```
[aznapi-configuration]
audit-json = yes
logcfg     = audit.azn:stdout

[logging]
requests-file      = stdout
request-log-format = {"host": "%h", "user": "%u", "time": "%t"}
```

Some additional log files will however still be generated on the disk of the container. By default, Docker uses a layered file system to help reduce the disk space utilization of the containers. However, this file system has slower write speeds than standard file systems. As such, a standard Docker practice is to place any files that are updated frequently (for example, log files) on a shared volume. All of the log files that are written by the container are located in the `/var/application.logs` directory. Therefore, the recommended approach is to create this directory as a shared volume when you create your container.

Note: Multiple containers should not reference the same persistent volume for log storage, otherwise multiple containers will attempt to write to the same log file at the same time, causing data write and integrity issues. In a Kubernetes environment this problem can be overcome by deploying the containers in a `StatefulSet` (refer to the official Kubernetes documentation for information on `StatefulSets`).

The log file directory structure is shown in the following table.

<i>Table 20. Logs Directory Structure</i>	
Log file	Sub-directory (relative to the root log directory)
WRP trace files	wrp/<instance>/trace
WRP statistic files	wrp/<instance>/stats
WRP crash files	wrp/<instance>/crash

Docker image for Verify Access Distributed Session Cache

The Security Verify Access Distributed Session Cache (DSC) Docker image provides the distributed session cache capabilities of Security Verify Access. It is a light-weight and secure alternative to running the same capabilities using the main Security Verify Access Image.

The main differences between running the Security Verify Access DSC image and the main Security Verify Access image as the DSC persona include:

- The DSC image is more light weight, which means that it consumes less disk space and memory, runs fewer processes and is quicker to start.
- The DSC image does not require any elevated container security capabilities/privileges (for example: the SETUID capability is not required).
- Logging records are sent to the console in JSON format so that the container logging infrastructure can manage the logging records. No support is provided for natively forwarding logging messages to a remote syslog server.
- The 'reload' capability is not supported which means that the container must be restarted in order to pick up configuration changes. In a Kubernetes environment the 'rolling update' capability of Kubernetes should be used to ensure that there is zero downtime when applying a configuration update.

Note: The definition for the Security Verify Access image and the Security Verify Access DSC image are interchangeable. In order to migrate to the Security Verify Access DSC image the image name simply needs to be changed. The rest of the definitions can remain the same. It is recommended that the Security Verify Access DSC image is used as it is a more secure and light-weight alternative, and the ability to execute the DSC capabilities from the Security Verify Access image will be deprecated in a future release. Please take note that the logging in the new container will by default be in JSON format, and the SETUID privilege is no longer required.

Consider the following points when you deploy a DSC container:

- The DSC container relies on a configuration snapshot which has been generated by the Security Verify Access configuration container.
- The DSC container should be started as the 'isam' user (UID: 6000). In a standard container environment this will happen automatically but in a Kubernetes environment the security context should be set to allow the container to start as this particular user.
- The following environment variables are used by the container:

CONFIG_SERVICE_URL

The URL that will be used to access the published configuration data. If using the configuration service of the Security Verify Access configuration container, the URL would be of the format: https://<container-ip>:<mapped-port>/shared_volume. A BA header will be supplied to handle authentication to the configuration service. If this environment variable is not specified the container will expect the configuration snapshot to be available in the shared configuration volume, at /var/shared/snapshots.

CONFIG_SERVICE_USER_NAME

The name of the user that will be used when accessing the configuration service.

CONFIG_SERVICE_USER_PWD

The password for the user that will be used when accessing the configuration service.

SNAPSHOT

The name of the configuration data snapshot file that is to be used when starting the container. No path information should be included in the filename. The standard snapshot locations will be searched for the specified snapshot file. This variable, if not specified, will default to the latest published configuration.

INSTANCE

The index of the DSC instance to be started (1-4). If no INSTANCE is specified, the instance with the index of '1' will be used.

LOGGING_CONSOLE_FORMAT

The required format for the log messages. Valid values are basic or json (default: json).

TRACE_LEVEL

The debugging trace level for the container. Valid values consist of the numbers 0 - 9 (default: 0).

LANG

The language in which messages which are sent to the console will be displayed. If no language is specified the messages will appear in English. The following table lists the supported languages:

Language	Environment Variable Value
Czech	cs_CZ.utf8
German	de_DE.utf8
Spanish	es_ES.utf8
French	fr_FR.utf8
Hungarian	hu_HU.utf8
Italian	it_IT.utf8
Japanese	ja_JP.utf8
Korean	ko_KR.utf8
Polish	pl_PL.utf8
Portuguese (Brazil)	pt_BR.utf8
Russian	ru_RU.utf8
Chinese (Simplified)	zh_CN.utf8
Chinese (Traditional)	zh_TW.utf8

Configuration

All configuration activities must be completed using the main Security Verify Access image, running as a configuration container. The configuration container supports a scaled-down version of the Security Verify Access appliance LMI. You can use this LMI to manage the configuration data.

Note: To make a configuration available to the DSC container, you must click **Publish configuration** in the LMI.

Service

The container will listen for incoming service and replication requests on the ports specified in the DSC configuration panel of the LMI.

Shared Configuration Data

The shared configuration volume is a section of the file system that is reserved for the storage of data to be shared amongst multiple containers. The data on the shared configuration volume is persisted even if the containers are deleted.

The shared configuration volume is mounted in a Security Verify Access DSC container at '/var/shared'. Configuration snapshots are retrieved from the 'snapshots' sub-directory of this volume if the CONFIG_SERVICE_URL environment variable is not specified.

The configuration data which is used by the Security Verify Access DSC image is fully compatible with the configuration data used by the legacy Security Verify Access image.

Logging

The logging of the DSC process will be sent to the console of the container, in JSON format. This allows the logging infrastructure of the container environment itself to manage the message logs.

Scenarios

These scenarios illustrate some of the typical situations an administrator encounters when using Security Verify Access in Docker environment and what actions the administrator can take in such situations.

Scenario - Initial configuration

The security administrator wants to construct a new Security Verify Access environment. This environment will be used to protect a single Web application.

The administrator completes the following steps.

1. Ensure that a user registry is available for the authentication of users. The administrator can use the **ibmcom/verify-access-openldap** image that was downloaded from Docker Hub to construct an OpenLDAP user registry if he wishes.
2. Pull the 'ibmcom/verify-access' image from [Docker Hub](#).
3. Start the configuration container.
4. Log onto the LMI of the configuration container, activate the base offering, and configure the Security Verify Access runtime environment and a Web Reverse Proxy instance. After the configuration changes have been completed, deploy the changes.
5. Start the Web Reverse Proxy "service" containers (multiple containers might be started for HA/load-balancing).

Scenario - Configuration update

The security administrator has a Security Verify Access environment already configured and running. Now the administrator is instructed to support a new application. This new application requires the addition of a new junction to an existing Web Reverse Proxy instance.

The administrator completes the following steps.

1. Start the configuration container (if not already started).
2. Log onto the LMI. Make the necessary configuration changes (for example, ACLs, junction creation, etc) and then deploy the changes.
3. If the service is using the main Verify Access image use the CLI to trigger a reload of the configuration. If the service is using the Verify Access Runtime image, the Verify Access Web Reverse Proxy image, or the Verify Access Distributed Session Cache image, you must restart the container.

Scenario - Replicated services

One of the Web Reverse Proxy instances is currently under load. So the security administrator wants to temporarily create a new Web Reverse Proxy instance on another docker host to help share the load.

The administrator completes the following steps.

1. Ensure that the configuration volume is available on the other docker host.
2. Start a new "Web Reverse Proxy instance" container on the other docker host.
3. Add the new docker container into the front-end load balancer.

Scenario - Upgrade

The security administrator currently has Security Verify Access running in a docker environment. A new version of Security Verify Access has just been released and so the administrator wants to upgrade the environment to this latest version.

The administrator completes the following steps.

1. Pull the latest Security Verify Access image from [Docker Hub](#).
2. Start a new configuration container using the latest Security Verify Access image.
 - When the image starts, it will automatically convert the data found in the configuration volume to the latest version.
 - The legacy data files will continue to exist so that Security Verify Access containers which are running the older version of Security Verify Access can continue to operate.
3. Start each service, one at a time, using the latest Security Verify Access images.
 - As each new service is started, stop the corresponding service that is running the older version of the image.
 - The services from the old version and services from the new version can co-exist in the environment. However, configuration changes to the services from the old version must be made using a configuration container also at the old version.

Note: In v10.0.2 a new light weight and more secure container has been provided for the Security Verify Access Runtime (AAC and Federation), the Web Reverse Proxy, and the Distributed Session Cache capabilities. In order to migrate to these new images from the legacy Verify Access image it should just be a matter of recreating the containers using all the existing definitions, with the exception of the image name, which should be changed to the new image name. The service port for the Runtime and Web Reverse Proxy containers should also be changed from 443 to 9443, or the port should be set to a custom value using the HTTPS_PORT environment variable.

Scenario - Support files

The security administrator has noticed that the Web Reverse Proxy is occasionally crashing. The administrator contacts the support team and they ask for the support file from the docker container that is experiencing the issue.

The administrator completes the following steps.

1. Execute the `/usr/sbin/isva_cli` command against the Docker container and then use the CLI to create a support file.
2. Start the configuration service (if not already started).
3. If the configuration data is being shared via the configuration service, the support file should be pushed to the configuration service using the CLI.
4. Log onto the LMI of the configuration container and go to **System > Network Settings > Shared Volume > Support files**. Locate the support file that was created in step 1. Download the support file and then send it to the support team.

Alternatively, the administrator can access the configuration volume directly to access the generated support file.

Scenario - AAC/Federation runtime configuration

The security administrator has a Security Verify Access environment that is already configured and running. Now the administrator wants to set up the Advanced Access Control (AAC)/Federation runtime container to use the AAC/Federation features, such as configuring the authentication service and OAuth authentication.

The administrator completes the following steps.

1. Ensure that the database server is running. One option would be to use the **ibmcom/verify-access-postgresql** image that is available for download from [Docker Hub](#).
2. Log on to the LMI of the configuration container.
3. Go to **System > Network Settings > Database Configuration**.
4. Configure the database settings.
5. Go to **System > Updates and Licensing > Overview**.
6. Activate the AAC/Federation module with the corresponding activation code.
7. Deploy the changes.
8. Start the AAC/Federation runtime container.

Note: When you run the AAC Auto Configuration Tool, use the configuration container's address and port for the Security Verify Access appliance LMI hostname and port, and the AAC LMI hostname and port arguments. Use the AAC runtime container's address and port (port 9080 or 9443 by default when using the Security Verify Access Runtime image) for the AAC runtime listening interface hostname and port arguments.

Orchestration

As each Docker container provides a single service, multiple containers with dependencies among them are usually required for a single environment. To simplify and automate the process, you can use Docker orchestration tools to deploy Security Verify Access to a Docker environment.

The orchestration tools that have been validated against Security Verify Access include Kubernetes and Docker Compose.

Kubernetes support

Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications.

It provides features such as:

- Self-healing
- Horizontal scaling
- Service discover and load balancing
- Secret and configuration management

Further information on Kubernetes can be found on the official Kubernetes Web site: <https://kubernetes.io/>.

Repository

The Security Verify Access image is available from the Docker Hub repository: 'ibmcom/verify-access', 'ibmcom/verify-access-runtime', 'ibmcom/verify-access-wrp', and 'ibmcom/verify-access-dsc'.

Secrets

Sensitive information, like passwords should never be stored directly in the yaml deployment descriptors. They should instead be stored within a Kubernetes secret and then the secret should be referenced in the yaml deployment descriptors. Instructions on how to use Kubernetes secrets can be found in the official Kubernetes documentation: <https://kubernetes.io/docs/concepts/configuration/secret/>

In the examples provided within this chapter, a 'secret' is used to store the Verify Access administration password. An example command to create the 'secret' is provided below (ensure that the **kubect1** context is set to the correct environment before running this command):

```
kubect1 create secret generic isva-passwords --type=string --from-literal=cfgsvc=Passw0rd
```

Service Accounts

Service accounts can be used to provide an identity for processes that run in a Pod. Information on the usage of service accounts can be found in the official Kubernetes documentation: <https://kubernetes.io/docs/tasks/configure-pod-container/configure-service-account/>.

In the examples that are provided within this chapter, the deployment descriptors use the 'isva' service account. The **kubect1** utility can be used to create the 'isva' service account (ensure that the **kubect1** context is set to the correct environment before running this command):

```
kubect1 create serviceaccount isva
```

Readiness, Liveness and Startup Probes

Kubernetes uses *liveness* probes to help determine whether a container has become unresponsive. If a container does become unresponsive Kubernetes automatically attempts to restart the container to help rectify the problem.

Kubernetes uses *readiness* probes to determine whether a container is ready to serve traffic. A pod with containers reporting that they are not ready will not receive traffic through Kubernetes Services.

Kubernetes uses *startup* probes to deal with applications that might require an additional startup time on their first initialization. The startup probe allows you to wait for a container to signal that it has fully started before the readiness and liveness probes are commenced.

Note: The startup probe was introduced in Kubernetes v1.16 and is not available prior to v1.16.

The Verify Access images provide a shell script which can be used to respond to liveness, readiness, and startup probes: `/sbin/health_check.sh`. If the `livenessProbe` command line option is provided to the script it will report on the liveness of the container, if the `startupProbe` command line option is provided to the script it will report on the startup status of the container, otherwise it will report on the readiness of the container. For a liveness probe the container will first check to see if it is still in the process of starting or reloading. If it is in the process of starting or reloading it will return a `healthy` result. Once the container is ready both the liveness' and 'readiness probes will return the network connectivity state of the service which is hosted by the container.

For more information on liveness, readiness, and startup probes, refer to the official Kubernetes documentation.

Deployment

The following section illustrates how to deploy Security Verify Access containers into a Kubernetes environment.

Configuration Container

Instructions on how to create the Security Verify Access configuration container are provided in the following steps:

1. Ensure that the **kubect1** context is set to the correct environment. The mechanism to do this differs based on the Kubernetes environment in use.

2. Create a configuration file that is named `config-container.yaml`. This configuration file defines a configuration container that can be used to configure your environment.

```
#
# The deployment description of the Verify Access configuration container. This
# container is used to manage the configuration of the Verify Access
# environment.
#

apiVersion: apps/v1
kind: Deployment

metadata:
  name: isva-config
  labels:
    app: isva-config

spec:
  selector:
    matchLabels:
      app: isva-config

  template:
    metadata:
      labels:
        app: isva-config

    spec:
      # The name of the service account which has the required
      # capabilities enabled for the isva container.
      serviceAccountName: isva

      # We want to run the container as the isam (uid: 6000) user.
      securityContext:
        runAsNonRoot: true
        runAsUser: 6000

      # We use a volume to store the configuration snapshot for the
      # environment.
      volumes:
        - name: isva-config
          emptyDir: {}

      containers:
        - name: isva-config

          # The fully qualified name of the verify-access image.
          image: ibmcom/verify-access:10.0.0.0

          # The port on which the container will be listening.
          ports:
            - containerPort: 9443

          # Environment definition. The administrator password is
          # contained within a Kubernetes secret.
          env:
            - name: SERVICE
              value: config
            - name: ADMIN_PWD
              valueFrom:
                secretKeyRef:
                  name: isva-passwords
                  key: cfigsvc

          # The liveness, readiness and startup probes are used by
          # Kubernetes to monitor the health of the container. Our
          # health is governed by the health_check.sh script which is
          # provided by the container.
          livenessProbe:
            exec:
              command:
                - /sbin/health_check.sh
                - livenessProbe
            initialDelaySeconds: 5
            periodSeconds: 10

          readinessProbe:
            exec:
              command:
                - /sbin/health_check.sh
```

```

    initialDelaySeconds: 5
    periodSeconds: 10

  startupProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - startupProbe
    initialDelaySeconds: 30
    periodSeconds: 10
    failureThreshold: 30

  # The '/var/shared' directory contains the configuration
  # snapshots and should be persistent. We use a volume for
  # this directory.
  volumeMounts:
    - mountPath: /var/shared
      name: isva-config

---

#
# The service description of the Verify Access configuration service. The
# service is only accessible from within the Kubernetes cluster.
#

apiVersion: v1
kind: Service

metadata:
  name: isva-config

spec:
  ports:
    - port: 9443
      name: isva-config

  selector:
    app: isva-config

  type: ClusterIP

```

3. Create the container:

```
kubectl create -f config-container.yaml
```

4. You can monitor the bootstrapping of the container using the **'logs'** command:

```
kubectl logs -f `kubectl get -o json pods -l app=isva-config | jq -r .items[0].metadata.name`
```

5. Start the Kubernetes proxy so that you are able to access the Web management console of the configuration container. An alternative approach is to create a Kubernetes service that directly exposes the LMI port of the configuration container.

```
kubectl port-forward `kubectl get -o json pods -l app=isva-config | jq -r .items[0].metadata.name` 9443
```

6. Access the proxied Web administration console (<https://127.0.0.1:9443>) authenticating as the **'admin'** user, with a password of **'PasswOrd'** (as defined in the isva-passwords secret). Proceed through the first-steps and then configure your environment.

7. Using the Web administration console, publish the configuration of the environment.

Web Reverse Proxy Container

The following steps illustrate how to create a Web Reverse Proxy container for the **'default'** instance:

1. Ensure that the **kubect1** context is set to the correct environment. The mechanism to do this differs, based on the Kubernetes environment being used.

2. Create a configuration file that is named **wrp-container.yaml**. This configuration file defines a WebSEAL container that can be used to secure access to your Web applications:

```
#
# The deployment description of the Verify Access Web Reverse Proxy
# container.
#

apiVersion: apps/v1
kind: Deployment

metadata:
  name: isva-wrp
  labels:
    app: isva-wrp

spec:
  selector:
    matchLabels:
      app: isva-wrp

  replicas: 1

  template:
    metadata:
      labels:
        app: isva-wrp

    spec:
      # The name of the service account which has the required
      # capabilities enabled for the verify-access container.
      serviceAccountName: isva

      # We want to run the container as the isam (uid: 6000) user.
      securityContext:
        runAsNonRoot: true
        runAsUser: 6000

      containers:
        - name: isva-wrp

          # The fully qualified name of the image.
          image: ibmcom/verify-access-wrp:10.0.2.0

          # The port on which the container will be listening.
          ports:
            - containerPort: 9443

          # Environment definition for the 'default' Web reverse
          # proxy instance. The administrator password is contained
          # within a Kubernetes secret.
          env:
            - name: INSTANCE
              value: default
            - name: CONFIG_SERVICE_URL
              value: https://isva-config:9443/shared_volume
            - name: CONFIG_SERVICE_USER_NAME
              value: admin
            - name: CONFIG_SERVICE_USER_PWD
              valueFrom:
                secretKeyRef:
                  name: isva-passwords
                  key: cfigsvc

          # The liveness, readiness and startup probes are used by
          # Kubernetes to monitor the health of the container. Our
          # health is governed by the health_check.sh script which
          # is provided by the container.
          livenessProbe:
            exec:
              command:
                - /sbin/health_check.sh
                - livenessProbe
            initialDelaySeconds: 5
            periodSeconds: 10

          readinessProbe:
            exec:
              command:
                - /sbin/health_check.sh
```

```

    initialDelaySeconds: 5
    periodSeconds: 10

  startupProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - startupProbe
    initialDelaySeconds: 30
    periodSeconds: 10
    failureThreshold: 30

```

3. Create the container:

```
kubectl create -f wrp-container.yaml
```

4. The **'isva_cli'** command can be used to directly administer a WebSEAL container:

```
kubectl exec -it `kubectl get -o json pods -l app=isva-wrp | jq -r .items[0].metadata.name` isva_cli
```

5. You can monitor the bootstrapping of the container using the **'logs'** command:

```
kubectl logs -f `kubectl get -o json pods -l app=isva-wrp | jq -r .items[0].metadata.name`
```

6. Create a configuration file that is named **wrp-service.yaml**. This configuration file defines a WebSEAL service that can be used to access WebSEAL. The type of service defined is different based on whether the 'load balancer' service type is supported in the environment.

The following definition can be used if the 'load balancer' service type is not supported in your environment:

```

#
# The service description of the Verify Access Web Reverse Proxy
# service. This is the entry point into the environment and can be
# accessed over port 30443 from outside of the Kubernetes cluster.
#

apiVersion: v1
kind: Service

metadata:
  name: isva-wrp

spec:
  ports:
    - port: 9443
      name: isva-wrp
      protocol: TCP
      nodePort: 30443

  selector:
    app: isva-wrp

  type: NodePort

```

The following definition can be used if the 'load balancer' service type is supported in your environment:

```

# LoadBalancer service definition....
apiVersion: v1
kind: Service
metadata:
  name: isva-wrp
spec:
  type: LoadBalancer
  ports:
    - port: 443
      targetPort: 9443
  selector:
    app: isva-wrp

```

7. Create the service:

```
kubectl create -f wrp-service.yaml
```

8. a. If a **'LoadBalancer'** service was defined, determine the external IP address of the service and then use your browser to access WebSEAL (port 443):

```
kubectl get service isva-wrp --watch
```

- b. If a **'NodePort'** service was defined, determine the IP address of the Kubernetes cluster and then use your browser to access the Web Reverse Proxy (port 30443). In a **'minikube'** environment the IP address of the cluster can be obtained with the following command:

```
minikube ip
```

In an IBM cloud environment, the IP address of the cluster can be obtained with the following command:

```
bluemix cs workers mycluster --json | jq -r .[0].publicIP
```

Runtime Container

The Verify Access Runtime Container (called `verify-access-runtime` or Verify Access Liberty Runtime) provides the advanced authentication, context-based access and federation services. The `verify-access-runtime` container retrieves a snapshot from the configuration container in the same manner as the Web Reverse Proxy Container. Due to the fact that the Web Reverse Proxy container should always act as a point of contact for the runtime service there should be no need to listen externally on a NodePort. Instead it only exposes its HTTPS interface on the cluster network with the `isva-runtime` service.

Note: The `verify-access-runtime` container can be used interchangeably with the `verify-access` main container for establishing a Runtime Container. It is recommended to use the more efficient and secure `verify-access-runtime` image

The following steps illustrate how to create a runtime container:

1. Ensure that the **kubectl** context is set to the correct environment. The mechanism to do this differs, based on the Kubernetes environment that is being used.
2. Create a configuration file that is named **runtime-container.yaml**. This configuration file defines a runtime container that can be used to secure access to your Web applications:

```
#
# The deployment description of the Verify Access runtime profile container.
# This container provides the Federation and Advanced Access Control
# capabilities of Verify Access.
#

apiVersion: apps/v1
kind: Deployment

metadata:
  name: isva-runtime
  labels:
    app: isva-runtime

spec:
  selector:
    matchLabels:
      app: isva-runtime

  replicas: 1

  template:
    metadata:
      labels:
        app: isva-runtime

    spec:
      # The name of the service account which has the required
      # capabilities enabled for the isva container.
      serviceAccountName: isva

      # We want to run the container as the isam (uid: 6000) user.
```

```

securityContext:
  runAsNonRoot: true
  runAsUser: 6000

containers:
- name: isva-runtime

  # The fully qualified name of the verify-access image.
  image: ibmcom/verify-access-runtime:10.0.2.0

  # The port on which the container will be listening.
  ports:
    - containerPort: 9443

  # Environment definition. The administrator password is
  # contained within a Kubernetes secret.
  env:
    - name: SERVICE
      value: runtime
    - name: CONFIG_SERVICE_URL
      value: https://isva-config:9443/shared_volume
    - name: CONFIG_SERVICE_USER_NAME
      value: admin
    - name: CONFIG_SERVICE_USER_PWD
      valueFrom:
        secretKeyRef:
          name: isva-passwords
          key: cfgsvc

  # The liveness, readiness and startup probes are used by
  # Kubernetes to monitor the health of the container. Our
  # health is governed by the health_check.sh script which is
  # provided by the container.
  livenessProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - livenessProbe
    initialDelaySeconds: 5
    periodSeconds: 10

  readinessProbe:
    exec:
      command:
        - /sbin/health_check.sh
    initialDelaySeconds: 5
    periodSeconds: 10

  startupProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - startupProbe
    initialDelaySeconds: 30
    periodSeconds: 10
    failureThreshold: 30

---

#
# The service description of the isva runtime profile service. The
# service is only accessible from within the Kubernetes cluster.
#

apiVersion: v1
kind: Service

metadata:
  name: isva-runtime

spec:
  ports:
    - port: 443
      targetPort: 9443
      name: isva-runtime
  selector:
    app: isva-runtime

  type: ClusterIP

```

3. Create the container:

```
kubectl create -f runtime-container.yaml
```

4. You can monitor the bootstrapping of the container using the **'logs'** command:

```
kubectl logs -f `kubectl get -o json pods -l app=isva-runtime | jq -r .items[0].metadata.name`
```

Distributed Session Cache

The Verify Access Distributed Session Cache Container (called isva-dsc) can be used by the Web Reverse Proxy and Runtime to share sessions between multiple containers. The isva-dsc container also retrieves a snapshot from the configuration container in the same manner as the Web Reverse Proxy Container. Besides the technical function of the container, the difference is that this container has no need to listen externally on a NodePort. Instead, it only exposes its HTTPS and replication interfaces on the cluster network with the isva-dsc service.

The following steps illustrate how to create a DSC container:

1. Ensure that the **kubectl** context is set to the correct environment. The mechanism to do this differs, based on the Kubernetes environment being used.
2. Create a configuration file that is named **dsc-container.yaml**. This configuration file defines a DSC container that can be used to share sessions:

```
#
# The deployment description of the Verify Access distributed session
# cache container.
#

apiVersion: apps/v1
kind: Deployment

metadata:
  name: isva-dsc
  labels:
    app: isva-dsc

spec:
  selector:
    matchLabels:
      app: isva-dsc

  template:
    metadata:
      labels:
        app: isva-dsc

    spec:
      # The name of the service account which has the required
      # capabilities enabled for the isva container.
      serviceAccountName: isva

      # We want to run the container as the isam (uid: 6000) user.
      securityContext:
        runAsNonRoot: true
        runAsUser: 6000

      containers:
        - name: isva-dsc

          # The fully qualified name of the verify-access image.
          image: ibmcom/verify-access-dsc:10.0.2.0

          # The ports on which the container will be listening. Port
          # 443 provides the main DSC service, and port 444 provides
          # the replication service which is used when replicating
          # session data between DSC instances.
          ports:
            - containerPort: 443
            - containerPort: 444

          # Environment definition. The administrator password is
          # contained within a Kubernetes secret.
          env:
            - name: INSTANCE
```



```

    value: '1'
  - name: CONFIG_SERVICE_URL
    value: https://isva-config:9443/shared_volume
  - name: CONFIG_SERVICE_USER_NAME
    value: admin
  - name: CONFIG_SERVICE_USER_PWD
    valueFrom:
      secretKeyRef:
        name: isva-passwords
        key: cfigsvc

  # The liveness, readiness and startup probes are used by
  # Kubernetes to monitor the health of the container. Our
  # health is governed by the health_check.sh script which is
  # provided by the container.
  livenessProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - livenessProbe
    initialDelaySeconds: 5
    periodSeconds: 10

  readinessProbe:
    exec:
      command:
        - /sbin/health_check.sh
    initialDelaySeconds: 5
    periodSeconds: 10

  startupProbe:
    exec:
      command:
        - /sbin/health_check.sh
        - startupProbe
    initialDelaySeconds: 30
    periodSeconds: 10
    failureThreshold: 30

---

#
# The service description of the verify-access distributed session
# cache service. The service is only accessible from within the
# Kubernetes cluster.
#

apiVersion: v1
kind: Service

metadata:
  name: isva-dsc

spec:
  ports:
    - port: 443
      name: isva-dsc
    - port: 444
      name: isva-dsc-replica
  selector:
    app: isva-dsc

  type: ClusterIP

```

3. Create the container:

```
kubectl create -f dsc-container.yaml
```

4. The '**dscadmin**' command can be used to directly administer the distributed session cache:

```
kubectl exec -it `kubectl get -o json pods -l app=isva-dsc | jq -r .items[0].metadata.name`
dscadmin
```

5. You can monitor the bootstrapping of the container using the '**logs**' command:

```
kubectl logs -f `kubectl get -o json pods -l app=isva-dsc | jq -r .items[0].metadata.name`
```

Kubernetes Environments

The following Kubernetes environments is validated by using the Security Verify Access image:

Minikube

Minikube is a tool that makes it easy to run Kubernetes locally. Minikube runs a single-node Kubernetes cluster inside a VM on your laptop for users looking to try out Kubernetes or develop with it day-to-day. Further information can be obtained from the Minikube Web site: <https://kubernetes.io/docs/getting-started-guides/minikube/>

To set the context for the **kubect1** utility use the following command:

```
kubect1 config use-context minikube
```

IBM Cloud

The IBM cloud container service provides advanced capabilities for building cloud-native apps, adding DevOps to existing apps, and relieving the pain around security, scale, and infrastructure management. Further information can be obtained from the IBM Cloud Web site: <https://www.ibm.com/cloud/container-service>

To set the context for the **kubect1** utility use the IBM Cloud CLI to obtain the **kubect1** configuration file:

```
bx cs cluster-config <cluster-name>
```

Microsoft Azure Container Registry

Azure Container Service (AKS) manages your hosted Kubernetes environment, making it quick and easy to deploy and manage containerized applications without container orchestration expertise. It also eliminates the burden of ongoing operations and maintenance by provisioning, upgrading, and scaling resources on demand, without taking your applications offline. Further information can be obtained from the Microsoft Azure AKS Web Site: <https://docs.microsoft.com/en-us/azure/aks/>

To set the context for the **kubect1** utility use the Microsoft Azure CLI:

```
az aks get-credentials --resource-group <group-name> --name <cluster-name>
```

Google Cloud Platform

Google Cloud Platform lets you build and host applications and websites, store data, and analyze data on Google's scalable infrastructure. Further information can be obtained from the Google Cloud Web Site: <https://cloud.google.com/kubernetes-engine/>

To set the context for the **kubect1** utility use the Google Cloud CLI:

```
gcloud container clusters get-credentials <cluster-name>
```

Redhat OpenShift

RedHat OpenShift is an open source container application platform based on the Kubernetes container orchestrator for enterprise application development and deployment. For more information, see: <https://www.openshift.com/>.

To set the context for the **kubect1** utility use the OpenShift CLI:

```
oc login
```

The **oc** binary is the preferred mechanism for accessing the OpenShift CLI and can be used interchangeably with the **kubect1** utility.

The default security context which is enabled by RedHat OpenShift is too restrictive for the main Verify Access container. As a result of this a less restrictive security context should be enabled for the service account which will run the main Verify Access containers (in the examples provided in this chapter we use the 'Verify Access' service account).

Note: The default security context is adequate for the verify-access-runtime, verify-access-wrp, and verify-access-dsc containers. When using the default security context the 'isam' user (uid: 6000), defined by the runAsUser entry, should not be specified in the deployment descriptor.

The pre-defined 'anyuid' security context can be used, but this does provide additional capabilities that are not required by the main Verify Access containers. To create a security context with the minimum set of capabilities required for the main Verify Access containers:

1. Ensure that the **oc** binary is available in the environment and that a login has already been performed.
2. Create a configuration file that is named `-scc.yaml`. This configuration file defines a new security context which can be used by the Verify Access containers:

```
#
# The minimum security context constraints which are required to run
# the Verify Access container. We cannot use the 'restricted' security
# constraint as we need additional capabilities which would otherwise
# be denied to the container. The 'anyuid' security constraint may
# be used, but it allows additional capabilities which are not
# required by the container.
#

kind: SecurityContextConstraints
apiVersion: v1

# The name and description for the security context constraint to be
# created.
metadata:
  name: isva-scc
  annotations:
    kubernetes.io/description: The Verify Access SCC allows the container to run
    as any non-root user.

# The following capabilities are not required.
allowHostDirVolumePlugin: false
allowHostIPC:             false
allowHostNetwork:        false
allowHostPID:            false
allowHostPorts:          false
allowPrivilegedContainer: false
readOnlyRootFilesystem:  false

# The priority is set to '10', otherwise the security constraint does
# not take affect when applied to a service account.
priority: 10

# The Verify Access container needs to be run as a 'custom' user, but does
# not need to run as the root user.
runAsUser:
  type: MustRunAsNonRoot
seLinuxContext:
  type: MustRunAs
fsGroup:
  type: RunAsAny
supplementalGroups:
  type: RunAsAny

# The following volumes are required by the Verify Access container.
volumes:
- configMap
- emptyDir
- projected
- secret
- downwardAPI
- persistentVolumeClaim

# By default we drop all capabilities and then only add back in the
# capabilities which are required by the Verify Access container.
requiredDropCapabilities:
- ALL

# The capabilities which are required by the Verify Access container.
allowedCapabilities:
- CHOWN
- DAC_OVERRIDE
- FOWNER
- KILL
```

```
- NET_BIND_SERVICE
- SETFCAP
- SETGID
- SETUID

defaultAddCapabilities:
- CHOWN
- DAC_OVERRIDE
- FOWNER
- KILL
- NET_BIND_SERVICE
- SETFCAP
- SETGID
- SETUID
```

3. Create the container:

```
oc create -f -isva-scc.yaml
```

4. Associate the new security context with the 'isva' user:

```
oc adm policy add-scc-to-user isva-scc -z isva
```

Helm Charts

Helm Charts help you define, install, and upgrade even the most complex Kubernetes application.

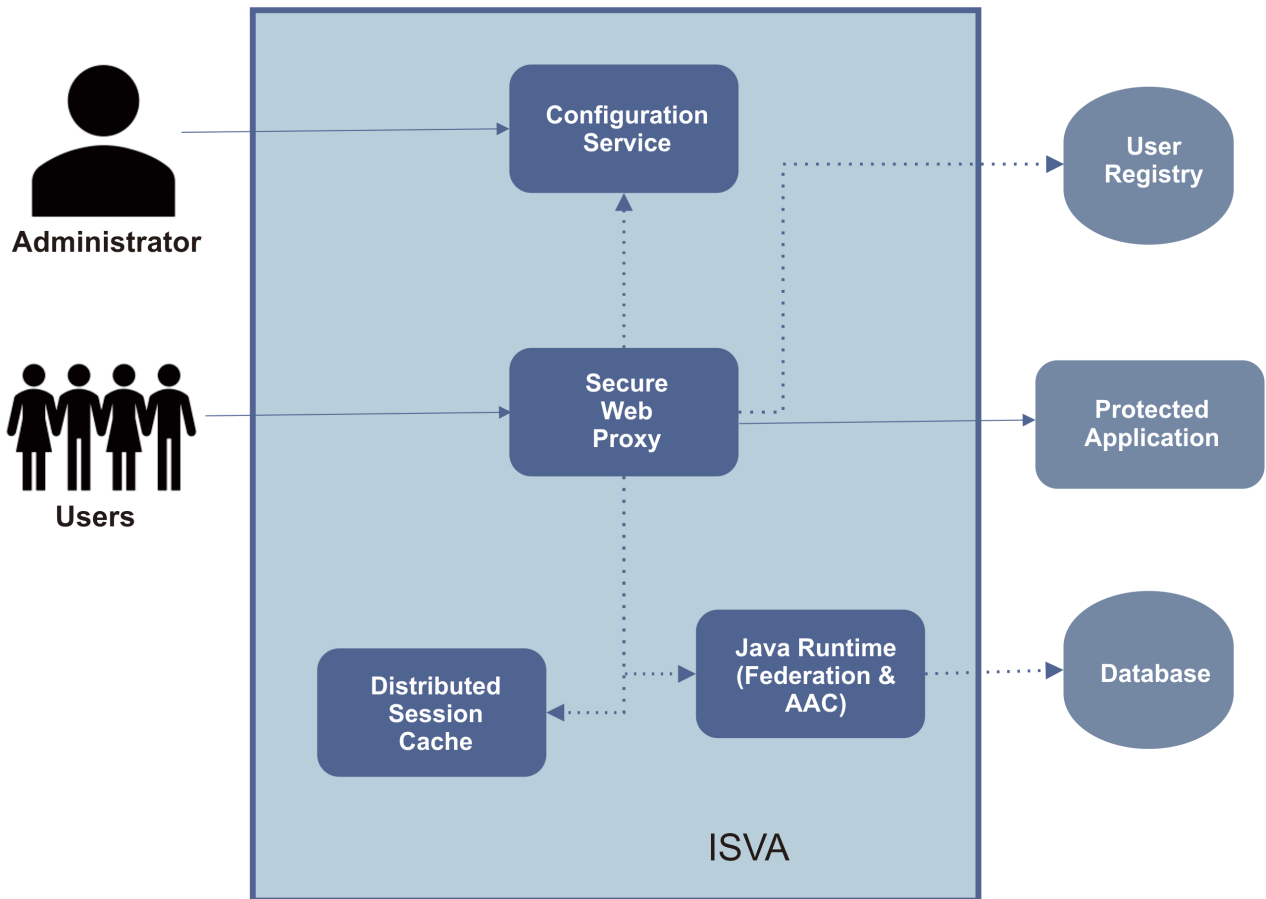
Helm helps you manage Kubernetes applications. Charts are easy to create, version, share, and publish. For more information on Helm, see <https://helm.sh/>.

IBM now provides a helm chart that can be used to deploy an IBM Security Verify Access environment. The helm chart is available from the IBM Security charts repository in GitHub: <https://github.com/IBM-Security/helm-charts/tree/master/stable/ibm-sam>.

This chart allows you to deploy a complete IBM Security Verify Access environment, including the following containers:

- Configuration
- WebSEAL
- Runtime
- Distributed Session Cache
- PostgreSQL Database (although this should only be used in test environments - an external database should always be used in production environments)

The following diagram shows the components of the `ibm-sam` Helm chart:



Note:

- Multiple Web Reverse Proxy instances can be created;
- The Distributed Session Cache will support a primary/secondary instance (active-passive);
- Each of the containers will retrieve configuration information from the configuration service of the configuration container.

Additional information can be found in the README .md file provided with the Helm chart.

Docker Compose support

Docker Compose provides a simple mechanism for defining multi-container environments.

Developers who want to familiarize themselves with the anatomy of a Security Verify Access Docker environment can use the following sample .yaml and .env file to easily build an environment on their workstation for development purposes. This practical example is used to illustrate the composition of an example Security Verify Access Docker environment.

docker-compose.yaml

```

version: '3'
services:

#
# Security Verify Access Containers
#

  isva-config:
    image: ibmcom/verify-access:${ISVA_VERSION}
    hostname: isva-conf
    environment:
      - SERVICE=config
#
# - SNAPSHOT=${SNAPSHOT?}
#
# - FIXPACKS=${FIXPACKS?}
#
# - ADMIN_PWD=${ADMIN_PWD?}

```

```

- CONTAINER_TIMEZONE=${TIMEZONE}
volumes:
- ./isva-volume:/var/shared
- ./isva-logs/conf:/var/application.logs
ports:
- ${CONFIG_HTTPS_PORT}:9443
depends_on:
- isva-ldap
- isva-db
cap_add:
- SYS_PTRACE
- SYS_RESOURCE

isva-webseal:
image: ibmcom/verify-access-wrp:${ISVA_VERSION}
hostname: isva-webseal
environment:
- INSTANCE=${WEBSEAL_INSTANCE_NAME}
# - SNAPSHOT=${SNAPSHOT}
volumes:
- ./isva-volume:/var/shared
- ./isva-logs/webseal:/var/application.logs
ports:
- "${WEBSEAL_HTTPS_PORT}:9443"
- "${WEBSEAL_HTTP_PORT}:9080"
depends_on:
- isva-ldap
- isva-dsc

isva-aac:
image: ibmcom/verify-access-runtime:${ISVA_VERSION}
hostname: isva-aac
environment:
- SERVICE=runtime
volumes:
- ./isva-volume:/var/shared
- ./isva-logs/aac:/var/application.logs
ports:
- "${AAC_HTTPS_PORT}:9443"
- "${AAC_HTTP_PORT}:9080"
depends_on:
- isva-ldap
- isva-db
- isva-webseal
- isva-dsc

isva-dsc:
image: ibmcom/verify-access-dsc:${ISVA_VERSION}
hostname: isva-dsc
environment:
- INSTANCE=1
# - SNAPSHOT=${SNAPSHOT}
# - FIXPACKS=${FIXPACKS}
volumes:
- ./isva-volume:/var/shared
ports:
- "${DSC_SERVICE_PORT}:443"
- "${DSC_REPLICA_PORT}:444"

#
# Service Containers
#

isva-ldap:
image: ibmcom/verify-access-ldap:${ISVA_VERSION}
hostname: isva-ldap
environment:
- LDAP_ADMIN_PASSWORD=${LDAP_PASSWORD}
# - LDAP_CONFIG_PASSWORD=${LDAP_CONFIG_PASSWORD}
# - LDAP_BASE_DN=${LDAP_BASE_DN}
# - LDAP_TLS_VERIFY_CLIENT=${LDAP_TLS_VERIFY_CLIENT}
# - LDAP_DOMAIN=${LDAP_DOMAIN}
# - LDAP_ORGANISATION=${LDAP_ORGANISATION}
# - LDAP_ENABLE_PORT_389=${LDAP_SSL_DISABLED}
volumes:
- libldap:/var/lib/ldap
- ldapslapd:/etc/ldap/slapd.d
- libsecauthority:/var/lib/ldap.secAuthority
ports:
# - ${LDAP_PORT}:389
- ${LDAPS_PORT}:636

```

```

isva-db:
  image: ibmcom/verify-access-postgresql:${ISVA_VERSION}
  hostname: isva-db
  environment:
    - POSTGRES_DB=${DB_NAME}
    - POSTGRES_USER=${DB_USER}
    - POSTGRES_PASSWORD=${DB_PASSWORD}
    - POSTGRES_SSL_CN=${DB_CN}
  #
  # POSTGRES_UNSECURE=${DB_SSL_DISABLED}
  volumes:
    - pgdata:/var/lib/postgresql/data
  ports:
    - "${DB_PORT}:5432"

```

Environment

The environment is defined in the following .env file.

```

ISVA_VERSION=10.0.0.0
TIMEZONE=Australia/Brisbane

#
# Security Verify Access CONTAINERS
#

# The name of the snapshot which is to be used when starting the container.
# The snapshot must reside in <shared-volume>/snapshots
# SNAPSHOT=

# A list of fixpacks to apply when starting the container.
# The fixpacks must reside in <shared-volume>/snapshots
# FIXPACKS=

# The password to be set for the default 'admin' user account.
# ADMIN_PWD=

# Config Container
CONFIG_HTTPS_PORT=10443

# AAC Container
AAC_HTTP_PORT=11080
AAC_HTTPS_PORT=11443

# WebSEAL default Container
WEBSEAL_INSTANCE_NAME=default
WEBSEAL_HTTP_PORT=12080
WEBSEAL_HTTPS_PORT=12443

# DSC Container
DSC_SERVICE_PORT=13443
DSC_REPLICA_PORT=13444

#
# SERVICE CONTAINERS
#

# LDAP Container
LDAP_PORT=14389
LDAPS_PORT=14636
LDAP_DOMAIN=ldap.ibm.com
LDAP_PASSWORD=passwd0rd
LDAP_ORGANISATION=isva
LDAP_BASE_DN=cn=isva
LDAP_CONFIG_PASSWORD=passwd0rd
LDAP_TLS_VERIFY_CLIENT=false
LDAP_SSL_DISABLED=true

# Database Container
DB_PORT=15432
DB_CN=isva
DB_SSL_DISABLED=false
DB_USER=postgres
DB_PASSWORD=passwd0rd
DB_NAME=isva

```

Overview

This Docker Compose configuration defines an environment with the following containers:

- Security Verify Access containers (**ibmcom/verify-access**, **ibmcom-verify-access-wrp**, **ibmcom/verify-access-runtime**, and **ibmcom/verify-access-dsc**)
 - Configuration container
 - WebSEAL instance container
 - AAC runtime container
 - DSC container
- Services
 - PostgreSQL server container (**ibmcom/verify-access-postgresql**)
 - OpenLDAP server container (**ibmcom/verify-access-openldap**)

This environment has been created for simplicity to demonstrate:

- The concept of the shared configuration volume.

The shared configuration volume is created in a folder named '**isva-volume**'. All Security Verify Access containers share this volume.
- Log file storage

The log file directories are created in a folder name '**isva-logs**'. Each Security Verify Access container has its own log directory within this folder.
- Port mappings that are used by each container

All environment variables and port mappings are externalized to the file '**.env**' for convenience.
- How to persist data within the OpenLDAP and PostgreSQL containers.

The PostgreSQL and OpenLDAP containers will store their data in folders named '**db**' and '**ldap**' respectively.

Note:

- If you are not using the Advanced Access Control capability, you do not need the **isva-postgres** and **isva-aac** containers. However, if you are using the Federation capabilities in your environment, you will need similar containers created.
- The name of the WebSEAL instance that is run in the **isva-webseal** container must be defined when the container is created. Customize the value of **WEBSEAL_INSTANCE_NAME** in **.env** or create your WebSEAL instance with the default name '**default**'.

Quick start

Place the '**docker-compose.yaml**' and '**.env**' files into a new directory. From that directory, execute the following command to start the test environment:

```
docker-compose up -d
```

This command will create and start all of the containers in the environment.

To access the LMI, open your web browser and visit:

```
https://{docker-host}:10443
or
https://{docker-host}:CONFIG_HTTPS_PORT if .env has been customized
```

To access the Security Verify Access CLI, execute:

```
docker exec -it <container-name> isva_cli
```

To destroy the environment, execute the following command:

```
docker-compose down
```


Note that the data stored on the shared configuration volume and log file directories will not be removed when the environment is destroyed.

Additional commands

Some example commands for some common Docker Compose tasks are listed in the following table:

Task	Command
Run just the configuration service container and its dependencies	<code>docker-compose run isva-config</code>
Stop the LDAP service container	<code>docker-compose stop isva-ldap</code>
Remove the stopped LDAP service container	<code>docker-compose rm isva-ldap</code>
Recreate the Database service container	<code>docker-compose up --force-recreate -d isva-db</code>

For more information about Docker Compose, see the Docker Compose website. (<https://docs.docker.com/compose/>)

CLI in a Docker environment

In a Docker environment, a subset of the appliance CLI commands are available for you to manage the runtime aspects of the appliance.

Note: The CLI commands are only available in the main verify-access image and are not available in the verify-access-runtime, verify-access-wrp images, and verify-access-dsc images.

The CLI can be accessed by invoking the "**isva_cli**" command in the container. For example, the command to access the CLI in a container with the name "isva_config" would be: "`docker exec -it isva_config isva_cli`".

The reload command

The **reload** global command is a new command that is used to reload the configuration for Docker containers. After making configuration changes, use this command to reload the latest configuration data and apply the changes to the running services.

The supported options include:

```
reload [all|check|policy|runtime] [force]
```

all

Reload the entire configuration. This will involve some minimal service interruption while the services are restarted.

check

Check whether the container is running with the latest snapshot.

policy

Reload the Security Verify Access policy database. No service interruption will occur as a result of this operation. The **policy** option is only available in Web Reverse Proxy containers.

runtime

Reload the federation and advanced access control runtime information. No service interruption will occur as a result of this operation. The **runtime** option is only available in runtime containers.

force

Use this option to force the use of the locally cached data in the event that the configuration service is unavailable.

Distributed Session Cache in Docker environment

The Distributed Session Cache (DSC) is an independent service that acts as a centralized session repository for a Web Reverse Proxy server environment. Servers in the environment can use the DSC to provide failover for user sessions.

When Security Verify Access is running in a Docker environment, you can use the DSC Configuration page of the LMI to configure the DSC. See [“Managing Distributed Session Cache in Docker”](#) on page 85.

To configure a Web Reverse Proxy instance to use the DSC, go to **Web > Manage > Reverse Proxy** and select to edit the instance. On the **Session** tab, select the **Enable Distributed Session Cache** option. If you enable the DSC within a Web Reverse Proxy instance but do not want the configuration to be automatically updated if the DSC configuration changes, set the value of the **dsess-auto-update** entry in the **[session]** stanza in the WebSEAL configuration file to no.

The SSL certificates that are used by the DSC are stored in the **dsc_key_store** key store. This key store is initially populated with a self-signed certificate that is used when connecting to the DSC servers. The self-signed certificate can be replaced with a CA-signed certificate using the **SSL Certificates** management page of the LMI.

To start the DSC container within a Docker environment, specify the Docker environment variable **INSTANCE = '1|2|3|4'** at container start time. The instance number corresponds to the role that the DSC container will play in the environment (1 corresponds to primary, 2 corresponds to secondary, 3 corresponds to tertiary, 4 corresponds to quaternary). You can configure up to four DSC servers in your environment for high availability of the DSC. See [Failover for the distributed session cache](#).

License usage with IBM Security Verify Access deployed on Kubernetes

The IBM License Metric Tool (iLMT) is an application that is provided by IBM to audit licensed products to ensure that licensed software is being used appropriately.

Any IBM Security Verify Access deployment used in a production environment must be licensed. To assist customers, this guide demonstrates how licensing information can be collected for containerized Verify Access deployments by using the Kubernetes infrastructure.

To ensure IBM Security Verify Access license compliance using Kubernetes, customers are required to:

1. [Deploy the iLMT operator.](#)
2. [Deploy IBM Security Verify Access with appropriate annotations.](#)
3. [Verify that license metrics are being collected.](#)

The following sections describe in detail what is required for each of these steps.

Note: The IBM Licence Metric Tool (iLMT) is only applicable when a processor based licensing model (PVU) is in use. It cannot be used to monitor user based licensing (UVU).

Deploy iLMT operator

The iLMT container can be deployed using a Kubernetes Operator maintained by IBM. This operator is available at <https://github.com/IBM/ibm-licensing-operator>. The iLMT Team provides a guide to deploy the operator using the Kubernetes cli tool (https://github.com/IBM/ibm-licensing-operator/blob/master/docs/Content/Install_from_scratch.md).

Once the license service container is running you can deploy IBM Security Verify Access containers and verify that license audit information is being recorded.

Deploy IBM Security Verify Access

The license service relies on fixed annotations being added to deployed pods in order to report on license usage. For IBM Security Verify Access the annotations which should be added to each deployment descriptor include:

Annotation	Value	Description
Product name	IBM Security Verify Access Virtual Edition	Name of application which is being licensed
Product Id	e2ba21cf5df245bb8524be1957857d9f	Internal identifier of the application being licensed
Product metric	PROCESSOR_VALUE_UNIT	Metric used for license usage calculation
Product charged containers	ibmcom/verify-access	Names of containers which will be charged

These annotations should be added to the deployment metadata for each IBM Security Verify Access container. The following code snippet demonstrates how the required Kubernetes annotations can be added to a deployments.

Note: Details about the container template selection or configuration have been omitted for conciseness.

```
### ISVA Config ###
apiVersion: apps/v1
kind: Deployment
metadata:
  name: isamconfig
  labels:
    app: isamconfig
  annotations:
    productName: "IBM Security Verify Access Virtual Edition"
    productId: "e2ba21cf5df245bb8524be1957857d9f"
    productMetric: "PROCESSOR_VALUE_UNIT"
    productChargedContainers: "ibmcom/verify-access"
spec:
  selector:
    matchLabels:
      app: isamconfig
  replicas: 1
  template:
    . . .
```

Verify the license server metrics

The Rest API of the license service container should be queried in order to verify that the container is able to correct collect the license audit data. Detailed information on the API can be found at https://github.com/IBM/ibm-licensing-operator/blob/master/docs/Content/Retrieving_data.md.

Chapter 9. Deployment methodologies

Read this section to understand the deployment architectures.

Cluster support

The Security Verify Access appliance includes cluster support, which allows multiple appliances to share configuration information and runtime information to work together in a clustered environment.

For information about how to configure and administer a cluster in the LMI, see [“Managing cluster configuration ”](#) on page 69.

Cluster support overview

To share configuration information between appliances and provide failover for services, you can configure your Security Verify Access appliances into clusters.

Every cluster has a *primary* master and up to three back-up masters, known as the *secondary*, *tertiary* and *quaternary* masters for high availability of cluster services.

By default, an individual appliance is configured as the primary master of a stand-alone cluster. You can configure other appliances to join the cluster as *nodes*. When an appliance is configured as a node, it can access and share the configuration information of the primary master.

Roles and services in a cluster

The nodes in a cluster share the cluster services, which include the distributed session cache, configuration database, geolocation database, and runtime database.

The IBM Security Verify Access appliance provides services that can be shared across the cluster.

You can configure more than one master appliance to provide failover for some of these services as described in [“Failover in a cluster”](#) on page 165.

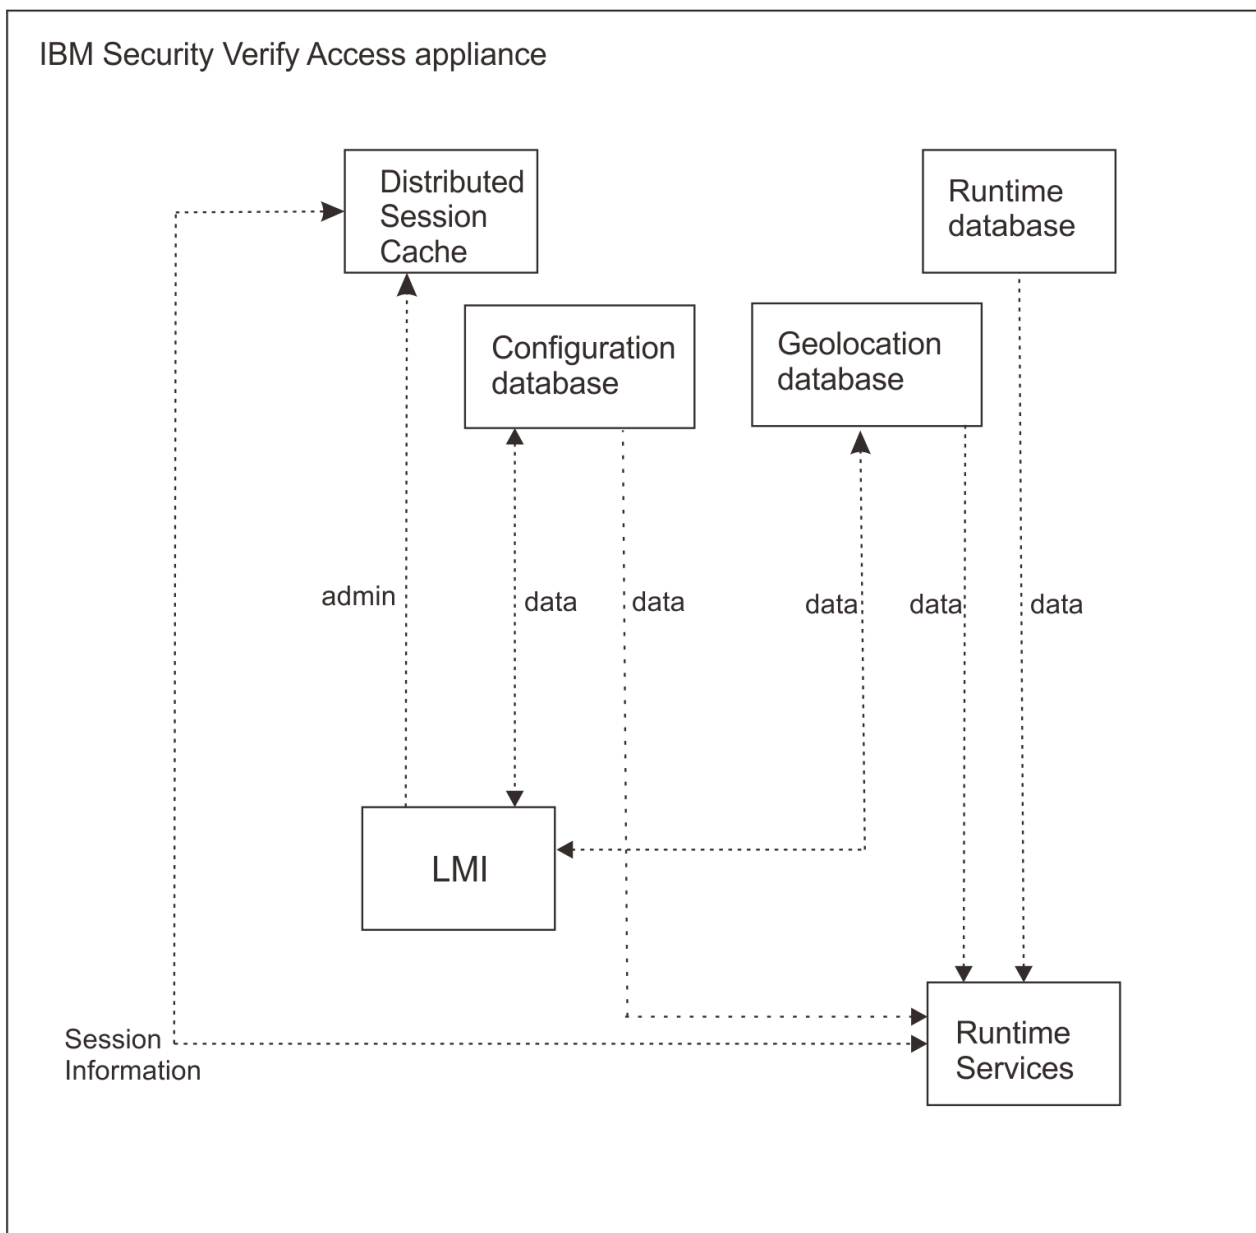


Figure 2. Services architecture

Distributed Session Cache

The distributed session cache is a central cache to hold user session information.

Configuration database

The configuration database stores configuration data that includes policy information, which is shared between the appliances in the cluster.

Note: You can update configuration data on the primary master only.

Geolocation database

The geolocation database provides geographic location information.

Runtime database

The context-based access component populates the high-volume database with runtime data. You can configure this database as an embedded database or an external database.

The embedded database is suitable for small environments only. For large-scale, production environments, configure an external database.

Data replication in a cluster

Cluster members share data that is relevant to the Security Verify Access configuration. You can update the configuration data on the primary master only. The other nodes in the cluster maintain local read-only replicas of the data from the primary master.

Any change to the cluster configuration or runtime parameters policy is automatically synchronized and applied to every node in the cluster. The **Cluster Configuration** management page in the LMI lists the nodes in the cluster. This list includes a **Status** column to indicate the status of the synchronization of system settings across the cluster.

If the changes to the system settings are not synchronized correctly on a particular node, the cluster administrator must investigate the problem. The administrator can examine the various log files on the node to determine why the change did not deploy successfully. When the problem is fixed, the administrator can either reboot the node, or rejoin the node to the cluster so that it applies the changes again.

Note: The **Status** column indicates whether the system settings on each node are up-to-date. This column does not indicate the status of any other synchronizations.

The data that is replicated across the cluster includes security settings, geolocation data, and system settings.

You can optionally configure the cluster to replicate the Security Verify Access runtime settings and the certificate database settings. Replicating the runtime settings can provide high availability for the Policy Server. For more information, see [“High availability for the policy server” on page 166](#).

Security Settings

In an IBM Security Verify Access appliance cluster, the nodes share configuration data and runtime data that is related to the security settings.

Configuration data

- One-time password (OTP) mapping rules.
- Policy information such as risk profiles, attributes, and obligations.
- Configuration information such as user registry data.
- All of the advanced configuration data.

Geolocation data

- Data that maps ranges of IP addresses to geographic locations.

Runtime data

- Session data.
- Non-session data that is relevant to the cluster, such as one-time passwords.
- Template files.

System settings

In an IBM Security Verify Access appliance cluster, the nodes share some system settings.

Cluster configuration

The cluster configuration information is replicated across the nodes of the cluster.

Runtime tuning parameters

The advanced tuning parameters are replicated across the nodes of the cluster.

Runtime settings

By default, the policy server configuration and policy database is not replicated across the cluster. However, you can choose to replicate this data. For more information about this configuration, see the "Replicate settings across the cluster" details in [“Managing cluster configuration ” on page 69](#).

SSL certificates

By default, the key file that is used by external clients to communicate with the DSC is not automatically distributed to nodes in the cluster. However, you can choose to replicate this data by selecting the 'Replicate with Cluster' check box on the **SSL certificates** management page.

High availability of cluster services

When you plan the architecture of your cluster, consider the services that you use in your environment along with your failover requirements for high availability. Include an External Reference Entity (ERE) for the primary and secondary masters in your architecture to assist in the failover process.

Cluster service considerations

A cluster requires at least one master, called the primary master, which provides the cluster services. For failover purposes in a cluster with multiple nodes, you can configure up to three more masters in the environment. The required number of masters depends on which services you use and your failover requirements.

The following table depicts the valid master configurations.

Number of masters	Combination of masters	Considerations
1	Primary master only.	No failover for cluster services.
2	Primary master and secondary master.	This configuration includes a secondary master to provide failover for the cluster services, which include the distributed session cache (DSC), configuration database, geolocation database, and runtime database.
3	Primary master, secondary master, and tertiary master.	You can optionally designate a tertiary master to provide extra failover for the distributed session cache. Only the distributed session cache recognizes the tertiary master node. The configuration, geolocation, and runtime databases consider the tertiary node as a non-master node.
4	Primary master, secondary master, tertiary master, and quaternary master.	You can optionally designate tertiary and quaternary masters to provide extra failover for the distributed session cache. Only the distributed session cache recognizes the tertiary and quaternary master nodes. The configuration, geolocation, and runtime databases consider these nodes as non-master nodes.

For high availability in a cross data center environment, you can consider separating the master appliances between the data centers as depicted in [Figure 3](#) on page 165.

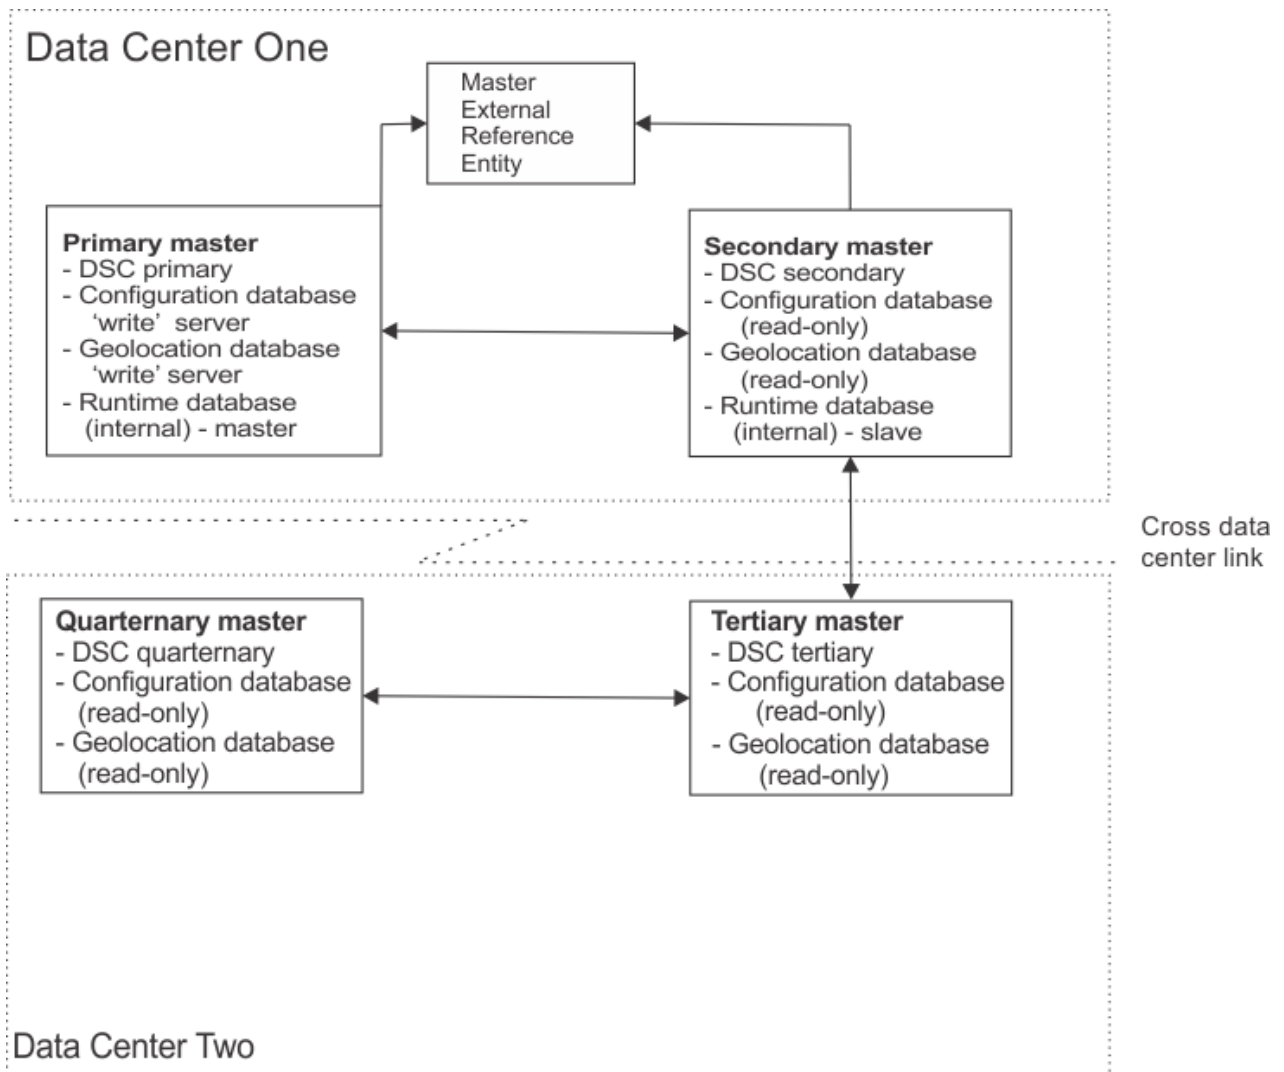


Figure 3. Example cluster architecture

This figure shows the data replication and service availability across the master nodes.

Distributed session cache

The primary master maintains the master copy of the distributed session cache and the other master nodes keep slave copies for failover purposes.

Runtime database

If you are using the internal runtime database, the primary master maintains the master copy of this data, while the secondary master keeps a slave copy for failover purposes.

If you are using an external runtime database, the cluster does not provide failover. In this case, the external database server is responsible for ensuring high availability.

Configuration and geolocation databases

The primary master is the only master on which you can update the configuration and geolocation databases. The other nodes in the cluster, including secondary, tertiary, and quaternary masters, maintain a read-only copy of the information from these databases.

Failover in a cluster

The distributed session cache, internal runtime database, geolocation database, and configuration database have varying failover capabilities in a clustered environment.

If you configure a secondary master and the primary master fails, the distributed session cache and the internal runtime database failover to the secondary master. When the primary master is restored,

reconciliation occurs and the primary master resumes control of these services. For the distributed session cache, a full copy of all sessions is restored.

You can also configure tertiary and quaternary masters for distributed session cache failover. If the primary and secondary masters are both unavailable, the distributed session cache fails over to the tertiary master. If the tertiary master is also unavailable, the distributed session cache fails over to the quaternary master.

The distributed session cache forms a chain of session replication from the primary to the quaternary master. If any node in the chain fails, it can request a full copy of all sessions from either partner in the chain when it recovers. There is no disk caching of sessions, so a full copy is required.

There is no failover between the master servers for the configuration and geolocation databases. If the primary master fails, the other nodes have a local read-only copy of the information that they can use in the interim. However, no configuration or geolocation updates are possible until the primary master is back online or a new primary master is designated.

Note: For WebSEAL to successfully fail over, set the following parameter in the **[junction]** stanza of the WebSEAL configuration file:

```
[junction] use-new-stateful-on-error = yes
```

See [use-new-stateful-on-error](#) for more information.

External Reference Entity

To prepare for failover situations, you must configure an External Reference Entity (ERE) for the primary and secondary master nodes.

When the communication link between the primary and secondary master nodes fails, both database servers might mistakenly assume that the other one is down. As a result, a dual primary situation can arise and you might lose transactions when databases are later synchronized. To avoid this situation, you can use a network reference device, such as a network router, as an ERE to check the health of the network.

If you configure a secondary master, you must also configure an ERE for the primary and secondary masters. If the primary master loses its connection to the secondary master, it can contact the ERE to determine whether there is a network fault or the secondary master is down.

In a distributed configuration, you can separate the primary and secondary masters into one data center and the tertiary and quaternary masters into another data center. If the data center link fails, the primary and tertiary masters operate in parallel and service requests in their local networks. When the data center link is restored, the tertiary master becomes inactive and reconciles its updates with the primary master.

Note: The above mentioned data reconciliation applies only to the distributed session cache. It does not apply to other databases such as the configuration database and geolocation database.

High availability for the policy server

You can enable the replication of the Security Verify Access runtime settings and the certificate database settings to achieve high availability for the policy server.

In a clustered environment, the Security Verify Access policy server can run on any node in the cluster. However, you must configure the policy server on the primary master if you want high availability.

To achieve high availability, you must adhere to the following requirements:

- The policy server must run on the primary master.
- You must configure replication for the runtime settings.
- If you are using SSL communication with an external directory server, you must configure replication for the certificate database settings.

You can configure the cluster to replicate the runtime settings and certificate database settings on the **Replication** tab of the **Cluster Configuration** page. For more information about these settings, see the "Replicate settings across the cluster" details in ["Managing cluster configuration" on page 69](#).

When you enable replication of the runtime settings, the policy server configuration and policy database information is copied from the primary master to every node in the cluster. The keys that are used for SSL communication between the Security Verify Access servers are also distributed across the cluster. If these settings are changed, the primary master sends the updates to the other nodes in the cluster.

The following process occurs when you enable replication of the runtime settings from the local management interface of the primary master:

- Any policy servers on other nodes in the cluster are stopped.
- The policy server configuration and policy database information is copied from the primary master to all other nodes in the cluster. Any existing policy server configuration on these nodes is overwritten by the configuration from the primary master.
- SSL keys for communication between the Security Verify Access servers are copied to every node.
- LDAP servers on other nodes in the cluster are stopped.
- If the Policy Server is configured to use a local LDAP, the LDAP data is copied to every node in the cluster and LDAP is started on each node.

Note: If there are WebSEAL instances or authorization servers, which are configured against a different policy server, you must reconfigure them to use the policy server on the primary master.

If you are using an external directory server with SSL enabled, you must configure the cluster to replicate the certificate database settings. If you enable this replication setting, the key files for SSL communication with the external directory server are distributed across the cluster.

If the primary master fails, you can promote any other node in the cluster to be the new primary master. The policy server starts automatically on the new primary master. All of the Security Verify Access servers on the other nodes are automatically reconfigured to use the policy server on the new primary master. The Security Verify Access servers can connect to the new policy server without requiring a restart. For more information about promoting a node to primary master, see ["Promoting a node to primary master when the original primary master is unavailable" on page 169](#).

When a node is promoted to primary master and replication for the runtime settings is enabled, the following process occurs:

- The replicated policy server configuration is modified to listen on one of the Management Interfaces.
- The policy server is started on the promoted node.
- If the Policy Server is configured to use a local LDAP, the local LDAP is started on the promoted node.
- Any configured WebSEAL and authorization servers on nodes in the cluster are modified to point to the policy server on the new primary master.

When you disable replication of the runtime settings, the policy server configuration and policy database information is removed from the other nodes in the cluster. If you are using the local LDAP on the primary master, the replicated copies of the LDAP files are removed from the other nodes. The WebSEAL instances and authorization servers in the cluster continue to use the policy server on the primary master.

Note: After you disable the replication, restart the Security Verify Access server on each node in the cluster.

If the policy server is configured with a local LDAP server as the user registry, high availability is provided. Each node of the cluster contains a read-only replica of the LDAP server that is used automatically in failover scenarios.

If the LDAP server provided by the primary master becomes unavailable to a node, any authorization servers that run on that node will failover to their local replicas. During this time, only read operations are possible. When the primary master LDAP server becomes available again, the node will automatically revert to normal operation.

Cluster failure management

If a cluster member fails, you must take different administrative actions, depending on the role of the node in the cluster.

Failure of the primary master

1. Promote a different node to the primary master. For detailed steps that describe how to promote a different node, see [“Promoting a node to master” on page 168](#).

You can promote a non-master node to the primary master so that other master nodes in the environment remain for failover purposes.

If there is a secondary master in the environment, you can optionally promote it to primary master. The process for this promotion depends on whether there are tertiary and quaternary masters in the environment:

- If there are tertiary and quaternary masters, you must take either of the following actions at the same time as you promote the secondary master to primary:
 - Promote a non-master node to secondary master, or
 - Demote the tertiary and quaternary nodes to non-master nodes.

You cannot have a tertiary and quaternary master without a secondary master.

- If you do not have tertiary and quaternary masters, you can promote the secondary master to primary master and the cluster can operate with a single master. However, for high availability purposes, you might also want to promote a non-master node to secondary master.
2. Remove the failed node from the cluster. For detailed steps, see [“Removing an unreachable master node from the cluster” on page 170](#).
 3. Export the signature file from the new master. You must use this signature file when you are adding new nodes to the cluster.

Failure of a secondary, tertiary, or quaternary master

1. Demote the failed node on the primary master.
2. Promote a non-master node to replace the failed master.

Note: You might need to complete steps 1 and 2 simultaneously to ensure that you maintain a valid combination of master nodes. For more information about valid architectures, see [“Cluster architecture rules” on page 172](#).

3. Remove the failed node from the cluster.

Failure of a node

1. Unregister the node on the primary master.
2. Optionally, you can add a node to the cluster to replace the failed node.

Promoting a node to master

If a master node fails, you might want to promote a different node to master while you resolve the failure.

About this task

When you are promoting a node to master, ensure that you adhere to the cluster architecture rules. For example, you must specify the supplementary masters in order. You cannot specify tertiary and quaternary masters if there is no secondary master. For a complete list of the cluster configuration rules, see [“Cluster configuration rules” on page 172](#).

Promoting a node to a master falls into two main categories:

- Promoting a node to a supplementary master - secondary master, tertiary master, or quaternary master.
- Promoting a node to primary master.

Promoting a node to a supplementary master

Procedure

You can use the local management interface of the primary master to update the cluster configuration and select the supplementary masters. To promote a node to secondary, tertiary, or quaternary master, complete these steps:

1. Open the **Cluster Configuration** page from the primary master local management interface.
2. Go to the **General** tab.
3. Change the values in the master fields. That is, **Secondary master**, **Tertiary master**, **Quaternary master**.
4. Save and deploy the updates.

Promoting a node to primary master when the original primary master is unavailable

About this task

- Nodes are automatically updated with information for the new primary master. If a node is not reachable by the primary master at the time of promotion, there is a delay of up to 15 minutes from the time that connectivity is restored before the node is notified of the new primary master.
- If the original primary master is reconnected to the primary master, it is automatically demoted to the role of a normal node.
- If the network is segregated and two different nodes are promoted to primary master in the different networks, automatic recovery is not possible when the network connectivity is re-established. In this situation, a manual merge of the segregated cluster is required. This step is achieved by removing all nodes from one of the clusters and joining these nodes back into the other cluster. This situation occurs only when both of the following conditions are met:
 - Connectivity in the cluster is lost.
 - The administrator promotes two different nodes to the primary master role while network connectivity is lost.

Procedure

Use the local management interface of the appliance that you are promoting to primary master to update the configuration. You can promote a non-master node or one of the supplementary masters if available. To promote the selected node to primary master, complete these steps:

1. Access the local management interface of the node that you want to promote to primary master.
2. Select **System > Network Settings > Cluster Configuration**.
3. Select the **General** tab.
4. Select **Set this appliance as a Primary Master**.
5. Use the available menu to set the Primary master IP address. Select the first management interface of the appliance.
6. Save and deploy the changes.

Promoting a node to primary master when the original primary master is available

About this task

- Nodes are automatically updated with information for the new primary master. If a node is not reachable by the primary master at the time of promotion, there is a delay of up to 15 minutes from the time that connectivity is restored before the node is notified of the new primary master.
- You can promote another node to primary master only if it is currently contactable by the current primary master.

Procedure

Use the local management interface of the current primary master to update the configuration.

1. Access the local management interface of the current primary master.
2. Select **System > Network Settings > Cluster Configuration**.
3. Select the **General** tab.
4. Select a new primary master from the list of nodes in the drop-down list.
5. If applicable, update the rest of the configuration to ensure that you do not break any of the clustering rules.
6. Save and deploy the changes.

Removing an unreachable master node from the cluster

If a master node is unreachable, you can demote it from master and then remove it from the cluster to resolve the failure. When the node is restored, you can register it with the cluster again as a non-master node.

Procedure

To remove the failed node from the cluster, complete the following steps in the local management interface of the new primary master:

1. Go to the **Overview** tab on the **Cluster Configuration** page.
2. Under the Nodes section, select the node to remove.
3. Click **Delete**.
4. Select the **Force** check box to force the removal of the node even if the node cannot be reached.
5. Click **Yes** to confirm the operation.
6. Deploy the changes.

After you remove the failed node from the cluster, you might want to restart it and ultimately restore it as a cluster member. In this case, you must complete some additional steps. While the node is disconnected from the network, change it to a stand-alone cluster with only a single node, as described in the following steps.

7. Restore the node and use its local management interface to access the **Cluster Configuration** page.
8. Go to **General** tab.
9. From the overview page, remove all other nodes.
10. Change the **Primary master** IP address to 127.0.0.1.
11. Save and deploy the change.
12. Troubleshoot the original failure and resolve any problems.

You can now join the restored appliance back in to the original cluster. This process joins the restored node to the cluster as a non-master node:

13. In the local management interface of the restored appliance, go to the **Overview** tab on the **Cluster Configuration** page.

14. Click **Import**.
15. In the **Join Cluster** window, click **Browse** to select the cluster signature file of the new primary master.

Note: You can generate the cluster signature file by using the local management interface of the new primary master and selecting the **Export** option in the **Overview** tab.

16. Click **Join** to add the current appliance to the cluster.
17. Deploy the changes.

Managing restricted nodes in a cluster

You can restrict nodes that are in the DMZ so that your network is secure. You can specify which nodes are restricted in the local management interface.

About this task

The following restrictions apply to restricted nodes:

- Restricted nodes cannot be promoted to any of the master roles.
- Restricted nodes cannot use the Policy Administration tool to modify the security policy.
- Restricted nodes do not contain a replica of the data that is stored by the embedded user registry.

You can restrict a node when you register a node in a cluster or at any time from the master local management interface. You can also restrict several nodes in a cluster.

Procedure

Select the steps for the task you want to complete:

- Configuring a restricted node during registration

Configure a restricted node when you register the node by using the local management interface.

- a. Register a node to a cluster.

For more information, see [Managing cluster configuration](#).

- b. Check **Join as restricted node** in the **Join Cluster** window.
- c. Click **Join** to add the appliance to a cluster as a restricted node.

- Configure a restricted node in a cluster

Use the local management interface to specify a restricted node in a cluster.

- a. Log on to the master appliance.
- b. From the top menu of the local management interface, select **System > Cluster Configuration**.
- c. Select the **Overview** tab.
- d. Select the node to be set as restricted in the **Nodes** grid.
- e. Click **Restricted Node**.
- f. Click **Submit**.

Cluster maintenance

Firmware updates in a cluster

To apply firmware updates in a cluster configuration, you must change the cluster configuration temporarily before the update so that changes can be written to the database.

For detailed instructions, see the [Use the local management interface for a cluster of appliances](#) section in [Upgrading to the current version](#).

Back up procedures

In a clustered environment, you cannot use VMWare snapshots to back up your virtual machines. For reliable backups, use appliance snapshots to back up the cluster.

You can complete an appliance snapshot on each cluster member to effectively back up the cluster. An appliance snapshot of the primary master includes all of the cluster configuration and runtime data. When the primary master is restored from an appliance snapshot, it updates every cluster member with the restored configuration.

An appliance snapshot of a node other than the primary master excludes the runtime database information. When a cluster member is restored from a snapshot, it contacts the primary master to obtain up-to-date configuration and runtime information.

To effectively back up the cluster, complete an appliance snapshot of the primary master after any change to the cluster configuration. For example, take a snapshot after you add or remove a node to ensure that the correct nodes are included in the cluster after a restore.

Cluster configuration rules

When you are configuring a cluster of Security Verify Access appliances, consider the following rules that govern cluster configuration.

General notes:

- Try to limit the number of changes that are made to the cluster configuration in a single policy update.
- After you save the policy changes, you must deploy the updates for the changes to take effect.

Cluster architecture rules

The architecture of a cluster, including the appointment of masters, is governed by numerous rules.

- A node must be a registered member of the cluster before it can be promoted to a master. The only exception is the primary master when there are no other nodes in the cluster.
- At a minimum, you must specify a primary master for the cluster.
- You must activate the product on the primary master of the cluster before any other node. If you use the internal runtime database in an Advanced Access Control-activated cluster, activate the Advanced Access Control-activated appliance on the secondary master before the other nodes.

Ensure that the product is activated on the masters before it is activated on any of the individual nodes in the cluster.

- The primary and secondary masters of the cluster must be activated at the highest level of all the nodes in the cluster. If any node in the cluster is activated with Security Verify Access base, the primary and secondary masters must also be activated with Security Verify Access base. Similarly, if any node in the cluster is activated with Advanced Access Control, the primary and secondary masters must also be activated with Advanced Access Control. Activation levels are validated when:
 - A node joins the cluster. Such validation is to ensure that the primary and secondary masters are activated to at least the same level.
 - A new primary or secondary master is set. Such validation is to ensure that the activation level of the new master is at least at the same level as the current primary master.
- You cannot specify a master without first specifying each of the prior masters. For example, you must specify the secondary master before you can specify a tertiary master.
- If you specify a secondary master, you must also specify the master external reference entity (ERE).
- You can modify the cluster policy on the primary master only, unless you are promoting a local node to primary master in a disaster recovery situation.

Cluster node availability

If a node is unavailable when you update the cluster configuration, it contacts the primary master to get the updated configuration information when it comes back online. If the primary master is offline at the same time as the secondary master, the primary master comes back online with read-only databases until the secondary master is available.

A node can become unavailable for a number of reasons, including a shutdown request, system failure, or networking failure. If a cluster node is not available during a cluster configuration change, it contacts the primary master for up-to-date information when it restarts. There might be a slight delay where the restored node tries to use the old policy and configuration information before it retrieves the missed updates.

The relationship between the primary and secondary nodes can be temporarily affected if both nodes are shut down simultaneously, and only one is powered back up. Until the other node is up, the databases on the newly powered up node are in read-only mode. When you power up the other node, the databases on the primary node become writable.

You can then shut down the secondary node without affecting the write capability on the primary server. It is only if both master nodes are offline at the same time that the restored primary master becomes read-only until the secondary master is back online.

This situation can be serious if the secondary node fails and the primary node stops for any reason. In this case, the primary node is not writable when it restarts until a secondary node is either started, or removed from the cluster. If the secondary node is removed, the primary master can operate as a single master in the cluster. You must address a failed primary or secondary master as soon as possible to avoid this situation.

Note: The above discussion about cluster node availability applies only to the configuration database and an embedded runtime database.

First management interface

In a clustered environment, the IP address of the first management interface is used as the node identifier. For this reason, a static IP address must be assigned to the first management interface of the appliance.

When you change the first management interface of a non-master node, the cluster is updated automatically.

You cannot change the IP address of the first management interface on a master. If you want to change the first management interface on a master node, you must first demote the node from master. You can then promote the node to master again and update any external client references in the distributed session cache.

Cluster registration

Before you register or unregister a node in a cluster, consider these registration rules.

- You must activate your products on the primary master before you activate the product on any other nodes.

If you are using the internal runtime database in a cluster, you must also activate the product on the secondary master before the other nodes.

For more information about the activation process, see [Activating the product and buying support](#).

- A node cannot be registered with a cluster if it is already a member of another cluster. In this situation, the node must first be unregistered from its current cluster.
- Node registration must occur directly through the local management interface of the appliance that you want to join the cluster. The appliance that you are registering must be able to communicate with the primary master.
- Node unregistration must occur on the primary master.

- A node cannot be unregistered if it is configured as a master. You must first demote the node from master and promote another node as the master.

Cluster ports

When you configure an appliance cluster, you are required to specify the starting port number for a range of ports to be dedicated to the services that are provided by the cluster.

It is important to note that these ports are for internal use only and are not used by the cluster for communication between nodes. All of the communication that takes place between nodes in the cluster occurs over port 22. This means that if your nodes are separated by a firewall, you only need to open up traffic on port 22 to allow the cluster services of the various nodes to communicate with each other.

The following diagram illustrates the communication requirements of the various roles in the cluster.




Figure 4. Communication in a cluster using port 22

If you want to manually configure a DSC client on a node within the cluster, use the following addresses and ports:

Note: The following examples assume that the first port is 2020.

Node	Port	Example
Primary	first_port + 15	127.0.0.1 15 127.0.0.1:2035
Secondary	first_port + 16	127.0.0.1 16 127.0.0.1:2036
Tertiary	first_port + 17	127.0.0.1 17 127.0.0.1:2037
Quarternary	first_port + 18	127.0.0.1 18 127.0.0.1:2038

Some additional settings are required to configure a DSC client on a node within the cluster. Set the priority for each distributed session cache server to 9 within the **server** stanza entry in the **[dsess-cluster]** stanza. Also set the **load-balance** stanza entry in the **[dsess-cluster]** stanza to no. The DSC does not support load balancing. Setting the **load-balance** configuration entry to no prevents connection attempts to servers for which the connection attempts will certainly fail.

Related reference

[server](#)

[load-balance](#)

Data loss considerations

The cluster services might lose data under certain circumstances.

Distributed session cache

- The policy data, which is used to indicate the first port that is available for use by the cluster, is changed.
- The policy data that defines the masters is changed.

Configuration database

An appliance that is operating as a single node cluster fails. In this situation, you must rely on snapshot information to restore the configuration database.

Internal runtime database

- An appliance that is operating as a single node cluster fails. In this situation, there is no recovery possible.
- The primary master fails, and no secondary master is configured.
- The maximum size of the internal runtime database is adjusted such that the new maximum size is smaller than the existing database.

Deployment pattern

Read this section to understand the components of a typical cluster environment and how to set up such environments. In this typical deployment scenario, the cluster incorporates both a Security Verify Access base appliance and an appliance with Advanced Access Control activated.

The following diagram illustrates a sample cluster environment.




Figure 5. Sample cluster environment

This environment consists of the following components:

- An external user registry, which can be a federated registry.
- Numerous appliances, namely:
 - A policy server
 - One or more appliances that provide the Advanced Access Control runtime service
 - Potentially an internal web reverse proxy to handle corporate traffic
 - One or more web reverse proxies in the DMZ to handle public traffic

In this scenario, all of the appliances reside in the same appliance cluster, with the policy server running on the primary master. Any of the other appliances that are running in the trusted zone can be enrolled as the secondary master, or you can have a dedicated secondary master appliance. The tertiary and quaternary masters are only required if you are using the distributed session cache across multiple data centers.

It is advisable to enroll the appliances that reside in the DMZ as restricted nodes. A restricted node imposes extra security constraints on the appliance, namely you cannot modify the security policy on these appliances or promote any of these appliances to a master.

In this environment, it is preferable to enable the replication of the Security Verify Access runtime environment and SSL certificate key files. For instructions on how to enable such settings, see [“Managing cluster configuration”](#) on page 69. The replication of the Security Verify Access runtime environment has the following advantages:

- You no longer need to configure the runtime environment manually on any node in the environment. The configuration information is automatically obtained from the primary master.
- If the primary master becomes unavailable (for example, due to hardware failure), you can promote one of the other unrestricted nodes to become a primary master and you do not lose the policy database. Nodes within the cluster are also automatically notified of the new policy server.

The following steps describe the recommended way in which to set up the environment:

1. Install each of the appliances. You should also:
 - Configure the networking.
 - Activate the required offerings.

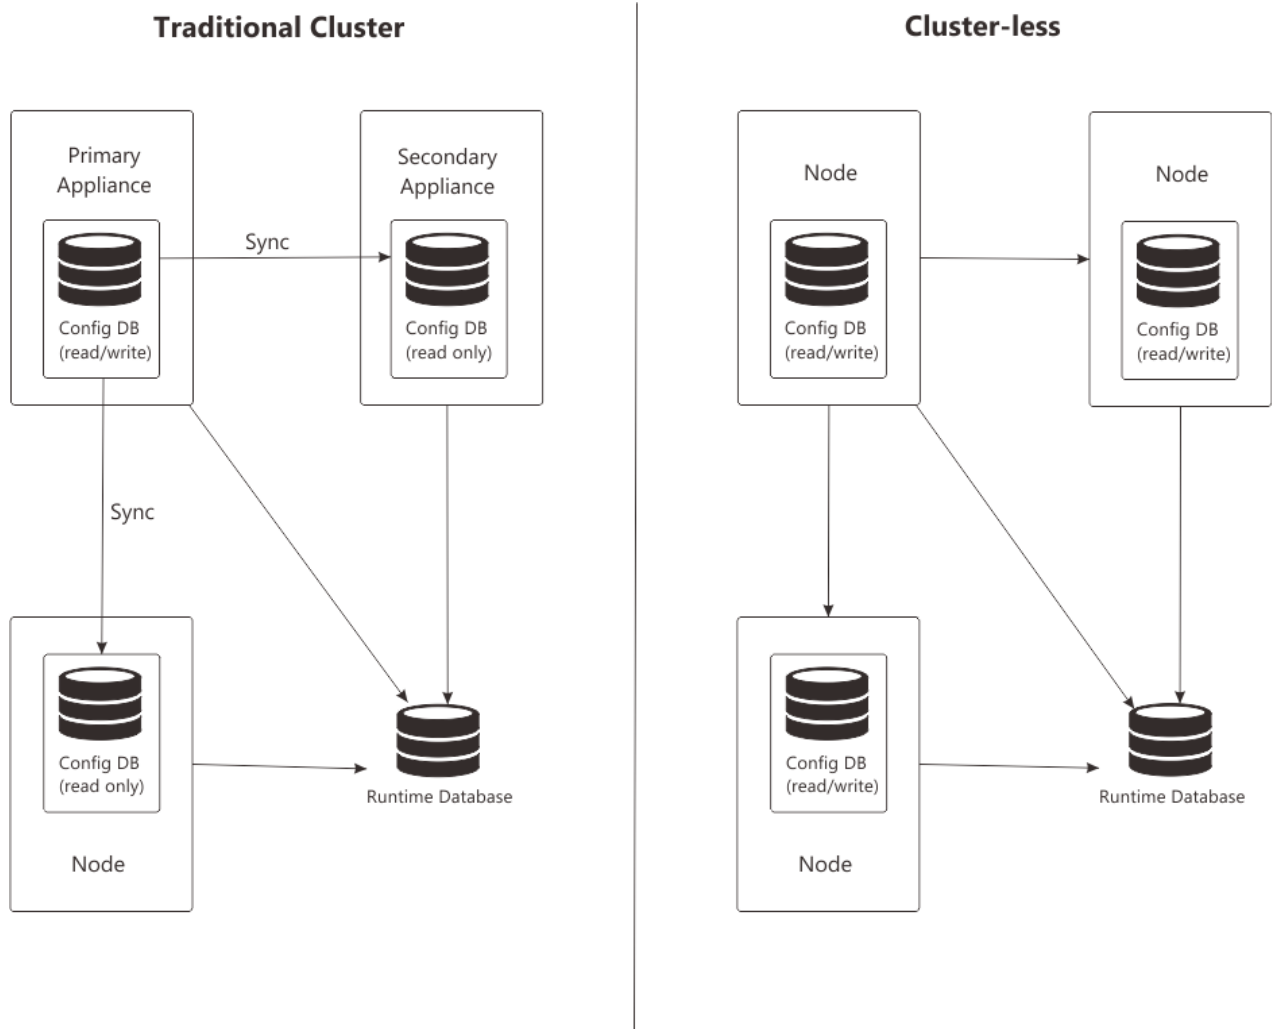
Note: The primary master must be activated with each offering that you will be using in your environment (for example, in this environment the primary master would be activated with both Security Verify Access base and Advanced Access Control).

2. Change the cluster configuration on the policy server to make it the primary master of a multi-node cluster.
3. On the primary master, configure the Security Verify Access runtime environment, including the policy server.
4. Enable the cluster replication of the runtime environment and certificate database.
5. Join each appliance to the cluster, one at a time. Join any appliances that reside in the DMZ as a restricted node.
6. Change the cluster configuration on the primary master to promote one of the unrestricted nodes to the role of secondary master. The node that is being promoted to secondary master must also be activated with each of the offerings that are used in the environment.
7. Configure the Security Verify Access base and Advanced Access Control security policies.
8. Configure the web reverse proxy instances on each of your Security Verify Access nodes.

Cluster-less AAC Deployment

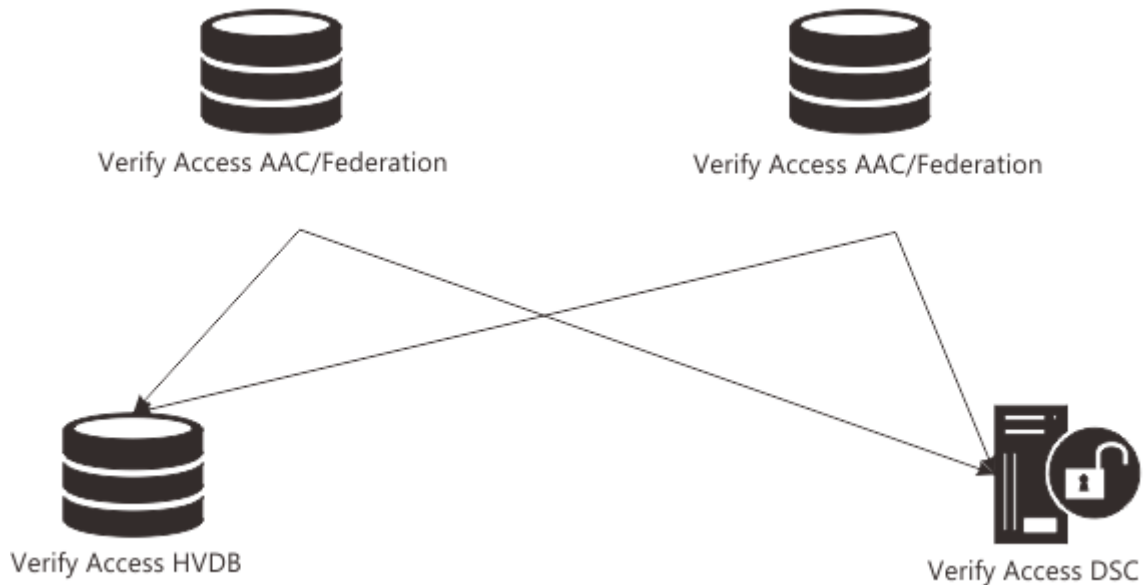
Feature enhancements from IBM Security Access Manager v9.0.6 onwards have enabled the AAC component to be deployed in a cluster-less architecture rather than the traditional clustered configuration.

A cluster-less deployment may be preferred because of technical limitations (Docker does not support clustering) or operational reasons, for instance a node can be upgraded by rebuilding the node rather than updating in situ.



In a traditional clustered deployment, the configuration is automatically synchronized by Verify Access between multiple appliances and can only be modified on the primary node of the cluster. Whereas in a cluster-less deployment it is the responsibility of the administrator of the environment to ensure that each appliance has an identical configuration using alternative mechanisms. In a containerized deployments (See Docker Support), this is easy and configuration snapshots can be used to instantiate a replicated instance (See Scenario - Replicated Services). For appliance deployments automation should be used instead to configure each appliance identically.

In order to achieve the cluster-less configuration of AAC with a seamless failover, the AAC nodes must be configured to exercise an external runtime database and distributed session cache.



AAC Specific Configuration

To enable consistent state sharing between nodes during authentication service flows (context based access or authentication policies), the authentication service must be configured to operate in cookie-less mode.

See [Configuring the authentication and access module for cookieless operation](#).

This enables the required session and state data to be stored within the DSC or runtime database, which is then available to other nodes in failover scenarios.

FIDO2 Specific Configuration

Prior to IBM Security Verify Access v10.0.1, FIDO2 registration and authentication ceremonies can not work in a cluster-less environment.

This is because both the FIDO2 server endpoints and the authentication mechanism that retrieving the Relying Party configuration by the randomly generated, read-only configuration ID. From v10.0.1 the configuration ID can be specified as a custom value when you are creating a new relying party, either through the UI or the REST API. Therefore, when you are creating relying parties in a cluster-less deployment, the configuration ID must be identical for an individual relying party across all nodes.

Federation Specific Configuration

For protocols such as SAML20, SAML11, and WS-Federation no additional configuration is required for cluster-less deployments to work, besides synchronized configuration and shared external runtime database across nodes.

The following features are tested and works in a cluster-less deployment:

Protocol	Scenario
SAML20	Single Sign-On – Different Binding (Post, Artifact, Redirect)
SAML20	Single Sign-On – Different NameIdFormats Email, Persistent (HVDB and LDAP data sources) and Transient
SAML20	NameIdManagement – Update and Terminate
SAML20	Single Logout – Different Bindings (Post, Artifact, SOAP)
SAML20	Single Sign-On With AccessPolicy

Protocol	Scenario
SAML11	Single Sign-On – Different Binding (Post, Artifact)
SAML11	Invoking an STS chain from a mapping rule during an SSO flow.
WS-Federation	Single Sign-On
WS-Federation	Single Sign-On with one-time assertion use enforcement set to true.

OpenID Connect Specific Configuration

OpenID Connect Provider dynamic clients must be migrated for the cluster-less deployment to work.

Migration can be performed for a specific API Protection Definition or for a specific dynamic client. Details about running the migration script can be found [here](#).

Once the migration successful, the cluster-less deployment for dynamic client works as expected.

Protocol	Scenario
OAuth 2.0/OIDC	AuthorizationCode, Implicit, and Hybrid flows with static client
OAuth 2.0/OIDC	AuthorizationCode, Implicit, and Hybrid flows with dynamic client
OAuth 2.0/OIDC	AuthorizationCode, Implicit, and Hybrid flows with different response types and response modes
OAuth 2.0/OIDC	AuthorizationCode, Implicit, and Hybrid flows with different token endpoint authentication mechanisms such as Post, Basic, JWT and Client Certificate.

Limitations

Each Verify Access instance must have the following items configured identically to ensure consistent behaviour.

- Reverse proxy authentication mechanisms
- Junctions
- Verify Access runtime configuration (including managed groups, ACLs, and POPs)

Chapter 10. Supported Web Reverse Proxy functionality

The IBM Security Verify Access appliance Web Reverse Proxy functionality is based on the technology included with the IBM Security Verify Access WebSEAL product. The appliance supports the majority of features that are offered by WebSEAL, with the exception of the items contained in the following table:

Feature	Description
Custom libraries, including CDAS and EAS	<p>The appliance does not support custom CDAS modules. As a result, the appliance does not support the following authentication mechanisms:</p> <ul style="list-style-type: none">• IP address• HTTP header• Post password change <p>WebSEAL does not provide CDAS modules for these mechanisms.</p> <p>Note: The appliance does support the IBM Security Identity Manager Password Synchronization Plug-in. For more information, see the [itim] stanza in the Stanza Reference topics in the Knowledge Center.</p>
Local junctions	<p>The following limitations apply to local junction support on the appliance:</p> <ul style="list-style-type: none">• The appliance can support a single fixed file system path for the local junction of a WebSEAL instance.• Local junctions on the appliance cannot execute any CGI scripts.
Application Response Measurement (ARM)	<p>WebSEAL software includes support for ARM to monitor transactions throughout the request and response processing stream. The appliance does not include ARM support.</p>
Tivoli Common Directory Logging	<p>The Tivoli Common Directory Logging feature stores all log files for IBM Security software applications in a common file system directory. The appliance does not support this common logging. Logging for the appliance is managed through the LMI.</p>
Auditing to a pipe or CARS	<p>The appliance cannot send audit records directly to a pipe or a CARS server. It can however, use an intermediate Verify Access authorization server to indirectly send audit records to the destinations.</p>

Table 23. WebSEAL features that the appliance does not support (continued)

Feature	Description
ARS (web service)	The IBM Security Verify Access for Web ARS web service can send request information to an external ARS server for authorization. ARS is not available on the appliance.

Chapter 11. Migration

Migrating an existing WebSEAL instance to the appliance

You can migrate an existing WebSEAL instance to the appliance.

Before you begin

1. Custom CDAS or EAS libraries are not supported. Make sure that there is no dependency on custom CDAS or EAS libraries before you start to migrate the system. For example, any custom CDAS processing must be converted to an EAI.
2. Local junctions are supported, but a fixed location is used as the document root. A local junction is also not permitted to run any CGI scripts. It can serve only static page content. Any CGI scripts must be migrated to a remote server. The appliance supports only a single local junction. The content for all other local junctions (if any) must also be migrated to a remote server.
3. As part of the migration process, you must collect the files that are necessary for the migration. You can use either of the following methods to collect the necessary files:

- Run the provided Perl script (`wga_migrate.pl`) to automatically collect the necessary files.

A Perl utility is provided to help facilitate the collection of files that are required by the WebSEAL instance. This utility can process the configuration for the specified WebSEAL instance. It can also copy the necessary files into the directory structure that is required by the import facility of the appliance.

To set up and run this utility, follow these steps:

- a. In the appliance top menu, go to **System > File Downloads**.
- b. Under **common > migrate**, select the `wga_migrate.pl` file to download it.
- c. Copy the script to the WebSEAL server.
- d. Ensure that Perl is installed and available on the WebSEAL server.
- e. Locate the name of the configuration file for the WebSEAL instance that is to be migrated.
- f. Run the `wga_migrate.pl` script, specifying the name of the WebSEAL configuration file and the destination directory. Use the following format for the script:

```
perl wga_migrate.pl [-c config-file] [-d dst-dir] {-v}
```

-c *config-file*

The name of the WebSEAL configuration file.

d *dst-dir*

The name of the destination directory. This directory must not exist on the file system.

-v

Display more status messages during the execution of the script.

For example, use the following script:

```
perl wga_migrate.pl -c /var/pdweb/etc/webseald-default.conf  
-d /tmp/migrate_out
```

- g. Review the files that are contained within the destination directory to ensure that all of the necessary files are located.
- Manually create the directory structure and copy the files to those directories.

On the source WebSEAL server, create the directory structure of configuration files, as defined in the following table. Only those directories for which files are to be migrated must be created. Create these directories as subdirectories under a single source directory.

<i>Table 24. Directory structure</i>	
Directory	Description
dynurl	Dynamic URL configuration files.
fsso	Forms-Based Single Sign-on configuration files.
jmt	Junction Mapping Table configuration files.
keytab	The key database (kdb/sth) files that are used by the WebSEAL instance. The files do not include the keyfile that is used to communicate with the policy server.
ltpa-keys	LTPA key files.
tam-keys	Key files that are generated with the cdsso-key-gen utility. They are used for things such as encrypting the failover cookie.
xslt/user-map-cdas	XSLT configuration file that is used by the client certificate user mapping CDAS.
xslt/http-transformation	XSLT configuration file that is used by the HTTP transformation rules function.
doc-root/docs	The files that are served by the WebSEAL local junction. These files are typically located under the /opt/pdweb/www- <i><instance></i> /lib/docs directory.
doc-root/errors	The error pages that are served by the WebSEAL instance. These files are typically located under the /opt/pdweb/www- <i><instance></i> /lib/errors directory.
doc-root/html	The management HTML pages (for example, login.html) which are served by the WebSEAL instance. These files are typically located under the /opt/pdweb/www- <i><instance></i> /lib/html directory.
doc-root/oauth	The OAuth response files, as defined within the [oauth-eas] stanza of the WebSEAL configuration file.
junctions	The XML files that contain the junction definitions for the WebSEAL instance. These files are typically located under the /opt/pdweb/www- <i><instance></i> /jct directory.
etc	The configuration files that are used by the WebSEAL instance. In particular, the routing file, the webseald- <i><instance></i> .conf, and the webseald- <i><instance></i> .conf.obf files.

Note: When you create the directory structure, additional subdirectories are not supported for any directory other than the doc-root ones (doc-root/docs, doc-root/errors, doc-root/html, doc-root/oauth). For example, you can create a directory structure such as /doc-root/error/<folder>/<file>, but a structure such as xslt/http-transformation/<folder>/<file> is not valid. For directories other than the doc-root ones, files can be placed only in the default root directories that are listed in [Table 24 on page 184](#). For example, xslt/http-transformation/<file>.

Note: All files to be copied must have unique file names. If two files have the same name, the migration tool copies only the first file that matches the name. For example, you might have the following structure:

```
[http-transformation]
request_pop1 = <path1>/pop1.xml
response_pop1 = <path2>/pop1.xml
```

Only <path1>/pop1.xml are created in the directory structure. All references to <path1>/pop1.xml and <path2>/pop1.xml in the configuration file are reduced to pop1.xml, which now points to the same file.

4. The WebSEAL configuration file must be included in the set of configuration files to be migrated. The obfuscated configuration file, as defined by the [configuration-database] stanza and **file** configuration entry, must also be included.
5. Modify the copied WebSEAL configuration file so that any configuration entries that are not applicable to the new WebSEAL instance are removed. Examples of entries that you might potentially not want to migrate would include network settings. The following configuration entries are ignored when the configuration file is imported into the appliance:
 - **token-card** configuration entry from the [authentication-levels] stanza
 - **server-name** configuration entry from the [server] stanza
 - **network-interface** configuration entry from the [server] stanza
 - [interfaces] configuration stanza
6. Create a compressed file, with the contents relative to the location that contains the copied files. For example, on a UNIX system, if the directory structure was created in /tmp/migrate, the command would be:

```
cd /tmp/migrate; zip -r /tmp/migrate.zip *
```

About this task

Migration is supported for the following versions:

- IBM Tivoli Access Manager Version 6.1 and later
- IBM Security Access Manager Version 7.0 and later

Procedure

1. Create a WebSEAL instance on the appliance with the local management interface.
2. Import the migration compressed file.

Note:

- If you are warned that files might be overwritten as a part of the import operation, you must validate the overwrite operation before you can continue. Make sure that the overwrite operation does not affect any other WebSEAL instances that might be running on the appliance. For detailed steps when you import with the local management interface, see [Import the contents of a compressed file into the administration pages root](#).
- If the appliance is a non-primary node in a clustered environment, and you enabled replication of SSL certificates in the cluster, first manually import the required SSL key files into the primary node and wait for these certificates to be replicated to the non-primary node. After the replication is complete, you can then import the WebSEAL configuration bundle into the non-primary node. If you do not follow this procedure in this type of environment and instead import the WebSEAL configuration directly on the non-primary node, the certificates from the WebSEAL configuration compressed file might be replaced during the next replication event by the certificates from the primary master and this will cause deployment issues.
- The import function audits the configuration file changes and logs the auditing details in the file migrate_YYYYMMDDHHMM.log. To access this file, go to **Monitor > Manage > Reverse Proxy Log Files**, select the instance from **Reverse Proxy Instances**, this log file is accessible under **Log Files for Selected Instance**.

3. Deploy the changes.
4. Restart the WebSEAL instance.
5. Examine the WebSEAL log file for any potential migration issues.

Migrating an existing Access Manager software environment to the appliance

You can migrate an existing Access Manager software environment to the appliance with the provided mechanism.

Before you begin

To achieve the migration, ensure that Perl is installed and available on the policy server to be migrated.

To migrate from an environment that is using Active Directory as the user registry, ensure that:

- IBM Directory Server client is installed on the policy server.
- The AD DS Snap-Ins and Command-Line Tools component is available on the policy server.

The appliance provides a Perl script to help with the collection of files that are necessary for the migration. These files include the IBM Security Verify Access configuration files, key files, and the authorization database.

Note: Such migration is supported for the following versions:

- IBM Tivoli Access Manager Version 6.1 and later
- IBM Security Access Manager Version 7.0 and later

Procedure

1. In the appliance top menu, go to **System > File Downloads**.
2. Under **common > migrate**, select the `isam_migrate.pl` file to download it.
This file is a Perl utility to help facilitate the collection of files that are required by the migration.
3. Copy the `isam_migrate.pl` file to the existing Access Manager software environment.
4. Run the `isam_migrate.pl` script, specifying the location of the runtime environment and policy server configuration path.

```
perl isam_migrate.pl [-c <config-path>] [-d <working-dir>] [-o <zip-file>] {-v}
```

-c <config-path>

The path of the IBM Security Access Manager configuration files.

-d <working-dir>

The name of the working directory. This directory must not exist on the file system.

-o <zip-file>

The name of the configuration bundle .zip file to produce. This file must not exist on the file system.

-v

Display more status messages during the execution of the script.

The following script is an example:

```
perl isam_migrate.pl -c /opt/PolicyDirector/etc/ -d /tmp/isam -o /tmp/isam.zip -v
```

Note: In most situations, the existing user registry is used by the migrated policy server. An exception to this situation is the environment where Active Directory is used as the user registry. In this situation, the Security Verify Access metadata must be migrated from the existing user registry to a new user registry. The `isam_migrate.pl` utility also provides this capability.

To migrate from a Windows computer that runs the Security Access Manager policy server, that uses Active Directory as the user registry, you can use the following commands:

-U

Unconfigure the old Active Directory policy server. This parameter is used to clean up the Security Verify Access user data from the Active Directory server after the data is migrated.

-i

The user registry that is embedded in the appliance is used by the policy server. If this parameter is not present, then the LDAP server is external to the destination appliance.

-h <ldap-host>

The host name of the user registry against which the policy server is configured. This option is not required if the "-i" option is used.

-p <ldap-port>

The port of the user registry against which the policy server is configured. This option is not required if the "-i" option is used.

-s

If this parameter is present, then SSL is used by the policy server when it is communicating with the external user registry. This option is not required if the "-i" option is used.

-D <ldap-admin-dn>

The distinguished name of the administrator of the external user registry that is used. This option is not required if the "-i" option is used.

-a <authority-suffix>

The LDAP suffix that is used to hold the Security Verify Access secAuthority data. This option is not required if the "-i" option is used.

-w <ldap-pwd>

The password for the administrator of the external or internal user registry.

-b

Migrate the users as Security Verify Access basic users.

-k <keyfile>

A GSKit CMS keyfile that contains the Active Directory CA certificate. If the option "-i" was not supplied and "-s" was supplied, then it must also contain the external LDAP server SSL CA certificate.

-W <keyfile-pwd>

The password for the specified keyfile.

-f <usergroup-ldif-file>

The file that stores all non-system user and group metadata in LDIF format. This file must be added after the policy server is migrated.

Note: This file is not used for the migration on the appliance. Do not include this file in the <zip-file>.

- Generate a migration .zip file that can be used to configure a policy server on the appliance with the embedded LDAP server.

```
perl isam_migrate.pl -i -c <config-path> [-v] -d <working-dir>
-o <zip-file> -w <ldap-pwd> [-b ] -f <usergroup-ldif-file> -k <keyfile> -W <keyfile-pwd>
```

As an example, use the following set of assumptions:

- The user is logged in to the Active Directory machine that is running the policy server and has administrative access to Active Directory.
- Perl is installed into the directory C:\perl.
- The isam_migrate.pl file is in C:\.
- The current working directory is C:\.
- A temporary directory is created: C:\tmp.
- The appliance has the default LDAP administrator password of "password".

- The Active Directory signer certificate is placed in the GSKit CMS file C:\adkeyfile.kdb with the password "passw0rd".
- The destination uses full Security Verify Access users, not basic users. The -b option is not provided.

The following command is based on the list of assumptions:

```
C:\perl\bin\perl.exe isam_migrate.pl -i -c "C:\Program Files\Tivoli
\Policy Director\etc" -d "C:\tmp\mig" -o "C:\tmp\migrate.zip"
-w passw0rd -k "C:\adkeyfile.kdb" -W passw0rd -f "C:\tmp\usergroup.ldif"
```

- Generate a migration .zip file that can be used to configure a policy server on the appliance with an external LDAP server.

```
perl isam_migrate.pl -c <config-path> [-v] -d <working-dir>
-o <zip-file> -w <ldap-pwd> [-b ] -f <usergroup-ldif-file> -k <keyfile> -W <keyfile-pwd>
-h <ldap-host> -p <ldap-port> [-s] -D <ldap-admin-dn> [-a <authority-suffix>]
```

As an example, use the following set of assumptions:

- The user is logged in to the Active Directory machine that is running the policy server and has administrative access to Active Directory.
- Perl is installed into the directory C:\perl.
- The isam_migrate.pl file is in C:\.
- The current working directory is C:\.
- A temporary directory is created: C:\tmp.
- The external LDAP server administrator is "cn=root" with password of "passw0rd".
- The Active Directory signer certificate is placed in the GSKit CMS file C:\adextkeyfile.kdb with the password "passw0rd".
- The external LDAP server, host name of extldap.ibm.com, requires SSL access on port 636 and its signer certificate is placed in C:\adextkeyfile.kdb.
- The external LDAP server has a suffix "secAuthority=Default" at which the Security Verify Access metadata is placed.
- The destination uses full Security Verify Access users, not basic users. The -b option is not provided.

The following command is based on the list of assumptions:

```
C:\perl\bin\perl.exe isam_migrate.pl -c "C:\Program Files\Tivoli
\Policy Director\etc" -d "C:\tmp\mig" -o "C:\tmp\migrate.zip"
-D "cn=root" -w passw0rd -h extldap.ibm.com -p 636 -s
-k "C:\adextkeyfile.kdb" -W passw0rd -f "C:\tmp\usergroup.ldif"
```

- Unconfigure the Active Directory server. This command is used to clean up the Security Verify Access user data from the Active Directory server after the data is migrated.

```
perl isam_migrate.pl -U -c <config-path> [-v]
```

Note: Use this unconfigure command only after you finish generating the migration .zip file.

As an example, use the following set of assumptions:

- The user is logged in to the Active Directory machine that is running the policy server and has administrative access to Active Directory and the local machine.
- Perl is installed into the directory C:\perl.
- The isam_migrate.pl file is in C:\.
- The current working directory is C:\.

The following command is based on the list of assumptions:

```
C:\perl\bin\perl.exe isam_migrate.pl -U -c
"C:\Program Files\Tivoli\Policy Director\etc"
```

5. If a compressed file is not automatically created on your platform, create a compressed file where the contents are relative to the location that contains the copied files.

For example, on a UNIX system, if the directory structure was created in `/tmp/isam`, the command would be:

```
cd /tmp/isam; zip -r /tmp/isam.zip *
```

6. In the destination appliance's local management console, import the compressed file created in the previous step.

- a) Go to **Web > Manage > Runtime Component**.
- b) Click **Configure**.
- c) Click **Import**.
- d) In the pop-up window, click **Browse**.
- e) Select the compressed file that contains the necessary migration files.
- f) Click **Import**.
- g) Deploy the changes.

Note:

- If you are migrating from an environment that uses a local LDAP server, you might need to manually change the host values (`localhost`) in the `pd.conf` and `ldap.conf` files to IP addresses that suit your new environment.
- The behavior of "[ssl] ssl-v3-enable" in `pd.conf` changed after version 6.1.1. It now provides the default for all other Security Verify Access servers on the same machine, unless their `.conf` file explicitly sets its value. Previously this option only affected the `pdadmin` command. So if "[ssl] ssl-v3-enable = yes" is set in the migrated `pd.conf`, and is not explicitly set in the migrated `ivmgrd.conf` file, then the policy server starts with SSLv3 enabled. To obtain the behavior before migration, add "[ssl] ssl-v3-enable = no" into the `ivmgrd.conf` file. It would be better to not use SSLv3 at all and set "[ssl] ssl-v3-enable = no" in the migrated `pd.conf` file.

What to do next

If you want to add the `<usergroup-ldif-file>` after migration, you must apply this file to the LDAP server that is used by the new policy server by using an LDIF tool.

For example, use the following **ldapadd** command:

```
/opt/ibm/ldap/V6.3/bin/ldapadd -h <ldap-host> -p <ldap-port> -D <ldap-admin-dn>
-w <ldap-pwd> -K <keyfile> -P <keyfile-pwd> -Z -i <usergroup-ldif-file>
```

Chapter 12. Configuration changes commit process

The LMI uses a two-stage commit process when you make changes to the appliance.

Stage 1

Changes are made by using the LMI and saved to a staging area.

Stage 2

The user explicitly deploys the changes into production.

Multiple changes can exist in a pending state at the same time. They are committed or rolled back together when a user deploys or rolls back these changes.

Pending changes are managed on a per user identity basis. This means that changes made by one user identity will not be visible to another user identity until the changes are deployed.

Note: As there is no validation or merging of changes that are made by different user identities to the same component, changes that are made by one user can potentially overwrite changes that are made by another user.

Any changes that affect running reverse proxy instances require a restart of the effected instances before the changes can take effect.

Certain appliance updates require either the appliance or the web server to be restarted before the changes can take effect. When one or more of these updates are made alongside other reverse proxy updates, an additional step is required to deploy the reverse proxy updates. You must:

1. Deploy all updates.
2. Restart the appliance or the web server.
3. Deploy all remaining updates.

If there are conflicts between the pending changes and the production files, then all pending changes are automatically rolled back and the production files remain unchanged.

Web service

Deploy the pending configuration changes

URL

```
https://{appliance_hostname}/isam/pending_changes/deploy
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON error response where applicable.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/deploy
```

Response:

```
200 ok
```

Roll back the pending configuration changes

URL

```
https://{appliance_hostname}/isam/pending_changes/forget
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON error response where applicable.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/forget
```

Response:

```
200 ok
```

Retrieve the number of outstanding changes

URL

```
https://{appliance_hostname}/isam/pending_changes/count
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON data that represents the number of pending changes.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes/count
```

Response:

```
{ "count": 3 }
```

Retrieve the list of outstanding changes

URL

```
https://{appliance_hostname}/isam/pending_changes
```

Method

```
GET
```

Parameters

N/A

Response

HTTP response code and JSON data that represents the list of pending changes.

Example

Request:

```
GET https://{appliance_hostname}/isam/pending_changes
```

Response:

```
200 ok

[{"id": 0,
 "policy": "SSL Certificates",
 "user": "admin",
 "date": "2012-11-05T11:22:20+10:00"
}]
```

Local management interface

When there are pending changes, a warning message is displayed at the top of the main pane. To deploy or roll back the pending changes:

1. Click the **Click here to review the changes or apply them to the system** link within the warning message.
2. In the **Deploy Pending Changes** page:
 - To view the details of changes that are made to a particular module, click the link to that module.
 - To deploy the changes, click **Deploy**.
 - To abandon the changes, click **Roll Back**.
 - To close the pop-up page without any actions against the changes, click **Cancel**.

Chapter 13. Runtime environment

In the local management interface, go to **Web > Manage > Runtime Component**.

Stopping, starting, or restarting the runtime environment

After you change the runtime configuration, you must restart the runtime environment to apply the changes.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.

Information about the status and the mode of the runtime environment is displayed.

Note: If the runtime environment is configured as either local stand-alone or remote stand-alone mode, you can stop, start, or restart it with this management page. Otherwise, the **Stop**, **Start**, and **Restart** buttons are disabled.

2. Depending on your needs, choose to stop, start, or restart the runtime environment.

- a) To stop the runtime environment, click **Stop**.
- b) To start the runtime environment, click **Start**.
- c) To restart the runtime environment, click **Restart**.

The records of these operations are logged to the policy server log files and user registry log files.

3. Optional: To manage the policy server and user registry log files, click the **Go to Application Log Files to view the Policy Server and User Registry log files** link. You can also access these log files by selecting **Monitor > Application Log Files** from the top menu.

Relevant entries can be found under `isam_runtime/policy_server` and `isam_runtime/user_registry`.

Configuring the runtime environment

To configure the runtime environment with the local management interface, use the Runtime Component management page.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Click **Configure**.

You can configure your policy server to be local or remote.

- **Local policy server with a remote LDAP user registry**
 - a. Under **Policy Server**, select **Local**.
 - b. Under **User Registry**, select **LDAP Remote**.
 - c. Under **Common**, check the check-box **Restrict Management Interfaces** if the local policy server should listen only on the local interface. If the box is not checked, the policy server will listen on all the available management interfaces.
 - d. Click **Next**.
 - e. On the Policy Server tab, provide settings for the fields displayed. Fields with an asterisk are required and must be completed.
 - **Management Suffix:** The LDAP suffix that is used to hold the IBM Security Verify Access secAuthority data.

Note: To create the domain at the secAuthority=Default tree, you must leave this field blank.

- **Management Domain:** The IBM Security Verify Access domain name.

Note: Make sure that the domain name you specify is unique among all domains on the LDAP server. The existence of a domain with the same name in a different suffix also causes an error. As this field is the name of the management domain, do not specify an LDAP DN.

Here are some example settings and the corresponding result data:

Setting	Result
Management Suffix: <blank> Management Domain: Default	secAuthority=Default
Management Suffix: OU=TAMDATA Management Domain: Default	secAuthority=Default,OU=TAMDATA

- **Administrator Password:** The security administrator's password.
- **Confirm Administrator Password:** The security administrator's password.
- **SSL Server Certificate Lifetime (days):** The lifetime in days for the SSL server certificate.
- **SSL Compliance:** Specifies any additional SSL compliance.

Note: If FIPS is enabled on the appliance, the **SSL Compliance** field cannot be set to No additional compliance.

f. Click **Next**.

g. On the LDAP tab, provide settings for the fields displayed.

- **Host name:** The name of the LDAP server.
- **Port:** The port to be used the system communicates with the LDAP server.
- **DN:** The distinguished name that is used when the system contacts the user registry.
- **Password:** The password for the DN.
- **Enable SSL:** Whether SSL is enabled.
- **Certificate Database:** The KDB file that contains the certificate that is used to communicate with the user registry. This field is required if "Enable SSL" is selected.
- **Certificate Label:** The label of the SSL certificate that is presented to the user registry upon request. This field is optional and is only required if SSL is enabled, and the user registry is configured to require a client certificate.

h. Click **Finish** to save the settings.

- **Local policy server with a local user registry**

Note: Users and groups within the local user registry are managed through the Security Verify Access administration framework; for example, pdadmin. All these users and groups are housed under the suffix "dc=iswga".

a. Under **Policy Server**, select **Local**.

b. Under **User Registry**, select **LDAP Local**.

c. Under **Common**, check the check-box **Restrict Management Interfaces** if the local policy server and user registry should listen only on the local interface. If the box is not checked, the policy server and user registry will listen on all the available management interfaces.

d. Click **Next**.

e. On the Policy Server tab, provide settings for the fields displayed. Fields with an asterisk are required and must be completed.

- **Administrator Password:** The security administrator's password.
 - **Confirm Administrator Password:** The security administrator's password.
 - **SSL Server Certificate Lifetime (days):** The lifetime in days for the SSL server certificate.
 - **SSL Compliance:** Specifies any additional SSL compliance.
- f. Click **Next**.
- g. On the **LDAP** tab, provide settings for the fields displayed. Fields with an asterisk are required and must be completed.
- Clean existing data**
Select this check box to delete any existing data in the embedded LDAP server before the configuration.
- h. Click **Finish** to save the settings.
- **Remote policy server**
 - a. Under **Policy Server**, select **Remote**.
 - b. Under **User Registry**, select whether to use **LDAP**.
 - c. Click **Next**.
 - d. On the Policy Server tab, provide settings for the fields displayed.
 - **Host name:** The name of the host that hosts the IBM Security Verify Access policy server.
 - **Port:** The port over which communication with the IBM Security Verify Access policy server takes place.
 - **Management Domain:** The IBM Security Verify Access domain name.
 - e. Click **Next** and complete settings on the **LDAP** tab.
 - **Host name:** The name of the LDAP server.
 - **Port:** The port to be used when the system communicates with the LDAP server.
 - f. Click **Finish** to save the settings.

Unconfiguring the runtime environment

To unconfigure the runtime environment component of the appliance with the local management interface, use the Runtime Component management page.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Click **Unconfigure**.
3. Take one of the following sets of actions.
 - **Unconfigure a local policy server with a remote LDAP user registry**
 - a. Enter the LDAP DN and LDAP password.
 - b. Select the **Clear user registry entries** check box if you want the unconfigure operation to remove all Security Verify Access domain, user, and group information. By default, this check box is not selected.
 - c. Click the **Force** check box if you want the unconfigure operation to forcefully remove all of the configuration data. By default, this check box is not selected.

Note: Select the **Force** check box only if the unconfiguration fails repeatedly. Use this option only as a last resort.
 - d. Click **Submit** to confirm the operation.
 - **Unconfigure a local policy server with a local user registry**

- a. Select the **Clear user registry entries** check box if you want the unconfigure operation to remove all Security Verify Access domain, user, and group information. By default, this check box is not selected.
- b. Select the **Force** check box if you want the unconfigure operation to forcefully remove all of the configuration data. By default, this check box is not selected.
Note: Select the **Force** check box only if the unconfiguration fails repeatedly. Use this option only as a last resort.
- c. Click **Submit** to confirm the operation.
- **Unconfigure a remote policy server**
 - a. Select the **Force** check box if you want the unconfigure operation to forcefully remove all of the configuration data. By default, this check box is not selected.
Note: Select the **Force** check box only if the unconfiguration fails repeatedly. Use this option only as a last resort.
 - b. Click **Submit** to confirm operation.

Managing runtime configuration files

To manage configuration files with the local management interface, use the Runtime Component management page.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Click **Manage > Configuration Files**.
3. Select one of the following runtime configuration files.

```
pd.conf
ivmgrd.conf
ldap.conf
activedir_ldap.conf
Routing File
```

Note: The **ivmgrd.conf** and **Routing File** options are only available when a policy server is configured on the appliance.

4. Edit the configuration file and then click **Save** to save the changes. If you do not want to save the changes, click **Cancel**. If you want to revert to the previous version of the configuration file, click **Revert**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Configuring JVM debugging for the runtime profile

Enable JVM debugging for the runtime profile so that you can debug new Java™ extension points.

Procedure

1. From the top menu, select **System > System Settings > Advanced Tuning Parameters**.
2. Click **New**.
3. In the **Key** field, enter `runtime_profile.jvm_option`.
4. In the **Value** field, enter the JVM debug options that suits your environment. For example, `-Xdebug -Xrunjdpw:transport=dt_socket,server=y,suspend=n,address=1044`.
5. Click **Save Configuration**.
6. Deploy your changes.

Chapter 14. Users and user registries

Configuring the runtime to authenticate basic users

Basic users are users in the registry that are not imported in to Security Verify Access. Edit the `ldap.conf` file so that basic users can authenticate in Security Verify Access.

Before you begin

The following limitations apply to basic users:

- Basic users work in minimal registry mode only.
- Basic users cannot use global sign-on.
- You cannot set access control lists for individual basic users. However, basic users can be members of a Security Verify Access group with access control lists.
- Registry direct Java API does not support basic users.
- Account and password valid settings are set to yes. You cannot modify them for basic users.



Warning: Basic users are not subject to any Security Verify Access account and password policies. They always have their `account-valid` and `password-valid` values set to yes. Basic users do not record the last login or last password change even if `[ldap] enable-last-login` is set. You must use the underlying registry equivalents for these capabilities.

About this task

Configure the run time so that basic users can authenticate to Security Verify Access. Basic users have limitations.

When `basic-user-support` is enabled, basic and full users are located by using the `basic-user-principal-attribute` suffix in the LDAP native user entry. If the located native user entry has full Security Verify Access user metadata then it is treated as a full user. The value of the `basic-user-principal-attribute` is used for the user ID even if the Security Verify Access full user metadata has a different `principalName`.

Basic users are managed in the corporate user registry by using LDAP management tools. These users are not managed through Security Verify Access, except when you change and reset passwords for basic users.

When searching for basic or full users, Security Verify Access:

- Uses the configured `basic-user-principal-attribute` and the `user-search-filter` values to locate users in the registry.
- Searches all suffixes that are defined by `basic-user-search-suffix` entries and in the order that they are defined, unless `basic-user-suffix-optimizer` is enabled. If no `basic-user-search-suffix` entries are specified, all suffixes are searched in an unspecified order.
- If `basic-user-suffix-optimizer` is enabled, a hit count is kept for each suffix that is used to search for users. The suffix search order is based on a dynamic most-used suffix order. This dynamic search order is not used if `basic-user-no-duplicates` is enabled since in that situation, all suffixes must be searched to ensure that there are no duplicates, thus the order is irrelevant.

Procedure

1. Log in the local management interface.
2. From the top menu, select **Web > Manage > Runtime Component**.
3. Click **Manage > Configuration Files**.

4. Select **ldap.conf**.
5. Add the following lines under the [ldap] stanza.

basic-user-support = yes

Set this option to *yes* to support basic users.

basic-user-principal-attribute = <uid>

This attribute is the principalName of the basic and full users.

basic-user-search-suffix = <DN>

Set this option for each suffix to search for full and basic users. This must include suffixes to search on the primary LDAP server and all federated registries.

If *basic-user-support* is enabled and one or more *basic-user-search-suffix* values are configured, the *ignore-suffix* entries are disregarded. The *basic-user-search-suffix* configuration entries determine the suffixes that are searched.

Note: When there are no *basic-user-search-suffix* entries, the system searches all available suffixes, except for those specified by the *ignore-suffix* entries. If you do not specify any *basic-user-search-suffix* values, you can use *ignore-suffix* entries to specify one or more suffixes to exclude from the search.

If *basic-user-search-suffix* is not set, then all suffixes are chosen in an unspecified order.

If you choose to specify one or more *basic-user-search-suffix* entries, ensure that you include an entry for every suffix that must be searched. Ensure that you include the primary suffix for Security Verify Access accounts. For example, *secAuthority=Default*. If you specify one or more *basic-user-search-suffix* entries, but you do not include this suffix, the search does not return the full Security Verify Access accounts. In this case, you are not able to authenticate to *pdadmin* with the *sec_master* account or any other Security Verify Access accounts.

basic-user-no-duplicates = {yes | no}

If set to *yes*, the search for basic users covers all suffixes to ensure that no users with the same name are found. If set to *no*, the search for basic users stops immediately and ignores possible duplicates.

Avoid configuring your environment to include suffixes that contain duplicates. Ensure that the *basic-user-principal-attribute* is unique for all accounts across the specified suffixes. If there are no duplicates in the environment, you can set *basic-user-no-duplicates* to *no* to improve search efficiency. However, if duplicates exist in your environment, set *basic-user-no-duplicates* to *yes* so that the system can return an error if it encounters more than one account with the same principal attribute value.

basic-user-suffix-optimizer = {yes | no}

If set to *yes* and *basic-user-no-duplicates* is set to *no*, the search order of suffixes is sorted, with the most hit of the basic user suffix at the head of the search suffix list. If set to *no*, the search order is provided by the *basic-user-search-suffix* order.

Note: If *basic-user-no-duplicates* is set to *yes*, the *basic-user-suffix-optimizer* entry is disregarded. In this case, all suffixes are searched to check for duplicates.

6. Add the following line under the [server: <fedreg>] stanza.

```
basic-user-principal-attribute = <uid>
```

7. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [Configuration changes commit process](#).

Embedded LDAP server management

When you configure the Security Verify Access runtime environment, you can choose to use an external user registry for storing the Security Verify Access metadata, or use the embedded user registry.

This same registry can optionally be used to also store the associated user data for the users. For more information, see [“Managing federated directories” on page 203](#).

By default, the contents of the embedded user registry are not included in snapshot files. To include the user registry data from the embedded user registry in snapshot files, set the **wga_rte.embedded.ldap.include.in.snapshot** advanced tuning parameter to **true**.

SSL support

The embedded LDAP server provides an SSL interface for management of the data contained in the user registry.

The embedded LDAP server listens on port 636 of the management interface of the appliance by default. The administrator can choose a port other than the default by modifying the advanced tuning parameter **wga_rte.embedded.ldap.ssl.port**. The advanced tuning parameters are accessed through **System > Advanced Tuning Parameters**. After you modify this advanced tuning parameter, you must restart the Security Verify Access runtime environment for the change to take effect.

The SSL certificates that are used by the LDAP server can be managed through the **SSL Certificates** panels of the LMI. For further details, see [“Managing SSL certificates” on page 109](#). The certificates are contained in the **embedded_ldap_keys** database file.

Two certificates are used by the LDAP server:

1. The certificate with the **server** label is used as the server certificate by the LDAP server. By default, the server certificate is a self-signed certificate. But this should be replaced in a production environment.
2. The certificate with the **ca** label is used as the CA certificate by the LDAP server. If no **ca** certificate is found in the key database, the server then uses the **server** certificate as the CA certificate. That is, it expects the server certificate to be a self-signed certificate.

In addition to this, the LDAP server can support mutual authentication by client certificates, providing that:

1. The client certificate has been signed by the CA that is known to the LDAP server. That is, the CA certificate is stored in the keyfile with a label of **ca**.
2. The distinguished name (DN) contained in the client certificate precisely matches a known LDAP user.

The FIPS setting of the appliance controls the ciphers that are supported by the OpenLDAP server.

Managing passwords

Administration of the data contained in the embedded LDAP server can be performed as the **cn=root,secAuthority=Default** user.

About this task

The default password for this user is **passw0rd**. The password should be modified in a production environment.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Select **Manage > Embedded LDAP > Password**.
3. Enter the new password in the **Password** field.
4. Enter the new password again in the **Confirm Password** field.
5. Click **OK** to change the password.

Managing suffixes

A *suffix* (also known as a naming context) is a DN that identifies the top entry in a locally held directory hierarchy. Because of the relative naming scheme used in LDAP, this DN is also the suffix of every other entry in that directory hierarchy. The embedded LDAP server can have multiple suffixes, each identifying a locally held directory hierarchy, for example, `o=ibm, c=us`.

About this task

The embedded LDAP server is pre-configured with a default suffix, `dc=iswga`, to make it easier to get started with the server. There is no requirement that you use this suffix. You can add your own suffixes and delete the pre-configured suffix.

There are two commonly used naming conventions for suffixes. One is based on the TCP/IP domain for your organization. The other is based on the organization's name and location. For example:

- Given a TCP/IP domain of `mycompany.com`, you might choose a suffix like `dc=mycompany, dc=com`, where the `dc` attribute refers to the domain component.
- If your company name is `My Company` and it is located in the United States, you might choose a suffix like one of the following examples:

```
o=My Company
o=My Company, c=US
ou=Widget Division, o=My Company, c=US
```

Where `ou` is the name for the **organizationalUnit** object class, `o` is the organization name for the **organization** object class, and `c` is a standard two letter country abbreviation used to name the **country** object class.

The following table lists the supported suffix elements and the corresponding object classes that are used when creating the top level entry for the suffix:

Element	Object class
dc	domain
c	country
o	organization
ou	organizationalUnit
l	locality

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Select **Manage > Embedded LDAP > Suffixes**.
All current suffixes are listed. You can then add or delete suffixes as needed.
3. Follow the prompts to complete the action you want to take.

Setting debug log level

Customize the log levels of the embedded LDAP server to suit your debugging needs.

Procedure

1. Select **Web > Manage > Runtime Component**.
2. On the Runtime Component page, select **Manage > Embedded LDAP > Change Debug Level**.

3. Select or clear the check boxes to indicate the wanted debug level. You can select zero to multiple debug level options.

Tip: Use the check box at the top to select or clear all debug level options.

Table 26. . Debug level option		
Debug level option	Keyword	Description
trace	trace	Trace function calls
connection	conns	Connection management
search.filter	filter	Search filter processing
config.file	config	Configuration processing
acl.processing	ACL	Access control list processing
statistics	stats	Statistics log connections, operations, or results
statistics.entries	stats2	Statistics log entries sent
shell.backend	shell	Print communication with shell backends
entry.parsing	parse	Print entry parsing debugging
sync.replication	sync	Sync replication consumer processing
uncategorized	none	Log messages that are not categorized including critical messages

4. Click **Submit**.

Managing federated directories

Keep your federated directories up-to-date so that Security Verify Access can access the most recent user information that is stored in external user registries. You can add a new directory, remove an existing one, or modify its settings.

About this task

Federated directories store the data that is associated with different users in different user registries. With federated directories, the appliance can access user information that is stored in a user registry external to Security Verify Access.

The DN of the user controls the user registry that is used when you search for user information. The Security Verify Access data that is associated with each user record is still stored in the Security Verify Access user registry. The Security Verify Access user registry is defined when you configure the runtime environment.

The **Federated Directories** menu item is enabled only if the runtime component is already configured.

Note: If the federated directories configuration is changed on the appliance that is running the policy server, the policy server is automatically restarted.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Select **Manage > Federated Directories**.

Note: All configured directories are displayed. By default, only the number of configured suffixes is shown. To view the suffixes in a particular directory, expand the relevant row.

3. Follow the prompts to complete the action you want to take.

Note: After you make any of the following changes, you must restart the Security Verify Access runtime environment for the changes to take effect.

- Add a directory
 - Click **New** and provide values for the displayed fields.
 - Multiple suffixes can be added on separate lines in the **Suffix** field.
 - If the **Enable SSL** option is selected, an extra field **Client Certificate** is displayed. Use the **Client Certificate** field to define the client personal certificate to present to the federated user directory server. This field is not required when one of the certificates in the keyfile was identified as the default certificate. The decision of whether to identify a certificate as the default depends on the configuration of the target user directory server.
 - You can click **Save** only if all of the fields are valid.
- Modify the settings for a configured directory
 - Select the directory to update and click **Edit**.
- Remove a directory or suffix
 - If you select a directory row and click **Delete**, the selected directory is removed. If you select a suffix row and click **Delete**, the selected suffix is removed.
 - Note:** Before you delete a federated directory, delete all federated users in this directory from Security Verify Access first.
 - The confirmation message indicates whether a directory or a suffix is being removed.
 - You cannot delete a suffix if it is the only suffix left in a directory, as such operation would leave the configuration in an invalid state. A directory must have at least one suffix to be valid.
- Update the LDAP SSL settings
 - Click **SSL Settings**.
 - This function updates the values in the `ldap.conf` configuration file. These values are only used if SSL settings do not exist in the configuration file of the hosting server. For example, if the settings exist in the WebSEAL configuration file, they take precedence over the settings that are contained in the `ldap.conf` configuration file.

Chapter 15. Reverse proxy instance management

In the local management interface, go to **Web > Manage > Reverse Proxy**. A list of all instances and their current states is displayed.

Stopping, starting, or restarting an instance

To stop, start or restart an instance with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. You can select single or multiple instances.

Stop an instance

- a. Select the instance(s) of interest.
- b. Click **Stop**.
- c. A message is displayed indicating that the instance has been stopped successfully.

Start an instance

- a. Select the instance(s) of interest.
- b. Click **Start**.
- c. A message is displayed indicating that the instance has been started successfully.

Restart an instance

- a. Select the instance(s) of interest.
- b. Click **Restart**.
- c. A message is displayed indicating that the instance has been restarted successfully.

3. To operate on all instances do not select any instances:

Stop all instances

Ensure that no instances are selected.

- a. Click **Stop All**.
- b. A message is displayed indicating that the instances are stopped successfully.

Start all instances

Ensure that no instances are selected.

- a. Click **Start All**.
- b. A message is displayed indicating that the instances are started successfully.

Restart all instances

Ensure that no instances are selected.

- a. Click **Restart All**.
- b. A message is displayed indicating that the instances are restarted successfully.

Configuring an instance

To configure an instance with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Click **New**.
3. Provide settings for the fields that are displayed on the **Instance**, **IBM Security Verify Access**, **Transport**, and **User Registry** tabs.

- On the **Instance** tab:

Field	Description
Instance Name	This is the new instance name, which is a unique name that identifies the instance. Multiple instances can be installed on the same computer system. Each instance must have a unique name.
Host Name	The host name that is used by the IBM Security Verify Access policy server to contact the appliance. The address that corresponds to this host name must match a management interface address of the appliance. The addresses that are associated with the application interface of the appliance cannot be used for communication with the IBM Security Verify Access policy server. Valid values include any valid host name or IP address. For example: libra.dallas.ibm.com
Listening Port	This is the listening port through which the instance communicates with the Security Verify Access policy server.
IP Address for the Primary Interface	The IP address for the logical interface.

- On the **IBM Security Verify Access** tab:

Field	Description
Administrator Name	The Security Verify Access administrator name.
Administrator Password	The Security Verify Access administrator password.
Domain	The Security Verify Access domain.

- On the **Transport** tab:

Field	Description
Enable HTTP	Specifies whether to accept user requests across the HTTP protocol.
HTTP Port	The port to listen for HTTP requests. This field is only valid if the Enable HTTP check box is selected.
Enable HTTPS	Specifies whether to accept user requests across the HTTPS protocol.
HTTPS Port	The port to listen for HTTPS requests. This field is only valid if the Enable HTTPS check box is selected.

- On the **User Registry** tab:

Field	Description
Enable SSL	Specifies whether to enable SSL communication between the instance and the LDAP server.
Key File Name	The file that contains the LDAP SSL certificate. This field is only valid if the Enable SSL check box is selected.
Certificate Label	The LDAP client certificate label. This field is only valid if the Enable SSL check box is selected.
Port	The port number through which to communicate with the LDAP server. This field is only valid if the Enable SSL check box is selected.

- Click **Finish**.

A message is displayed indicating that the instance has been configured successfully.

Unconfiguring an instance

To unconfigure an instance with the local management interface, use the Reverse Proxy management page.

Procedure

- From the top menu, select **Web > Manage > Reverse Proxy**.
- Select the instance to unconfigure.
- Click **Delete**.
- Enter the administrator name and password.
- Click **Delete**

Note: Select the **Force** check box if unconfiguration fails multiple times. Use this option only as a last resort.

Managing web reverse proxy configuration entries

To manage the web reverse proxy basic configuration, use the Reverse Proxy management page.

Procedure

- From the top menu, select **Web > Manage > Reverse Proxy**.
- Select the instance of interest.
- Select **Edit**.
- Make your changes to the settings on the **Server, SSL, Junction, Authentication, SSO, Session, Response, Logging, and Interfaces** tabs.

Server

The Server tab contains entries that are related to the general server configuration.

Field	Description
HTTPS	Select this check box to enable the HTTPS port within Reverse Proxy.
HTTPS Port	The port over which Reverse Proxy listens for HTTPS requests.
Enable HTTPS Proxy Protocol	Select this check box to enable proxy protocol support for incoming connections on the primary interface over HTTPS.
HTTP	Select this check box to enable the HTTP port within Reverse Proxy.
HTTP Port	The port over which Reverse Proxy listens for HTTP requests.

Field	Description
Enable HTTP Proxy Protocol	Select this check box to enable proxy protocol support for incoming connections on the primary interface over HTTP.
Interface Address	The network interface on which the Reverse Proxy server listens for requests.
Enable HTTP/2	Select this check box to enable HTTP/2 incoming connections on the primary interface from clients (browsers).
Persistent Connection Timeout	The maximum number of seconds that a persistent connection with a client can remain inactive before it is closed by the server.
Worker Threads	The number of threads that are allocated to service requests.
Cluster is Master	If the Reverse Proxy clustering function is used, this check box controls whether this Reverse Proxy server acts as the cluster master.
Master Instance Name	The server name for the Reverse Proxy instance which is acting as the master within the cluster. This option is only enabled if the Cluster is Master check box is not selected.
Message Locale	The locale in which the Reverse Proxy runs.

SSL

The SSL tab contains entries that are related to the general SSL configuration of the server.

Field	Description
SSL Certificate Key File	The key database that is used to store the certificates which are presented by Reverse Proxy to the client.
Network HSM Key File	The key database that stores the certificates to be used by the network Hardware Security Module (HSM) device.
SSL Server Certificate	The name of the SSL certificate, within the key database, which is presented to the client. The drop-down list includes certificates from both the local and network key files. The certificates from the network key file are prefixed with the token label for the network HSM device.
JCT Certificate Key File	The key database that is used to store the certificates which are presented by Reverse Proxy to the junctioned Web servers.

Junction

The Junction tab contains entries that are related to the general junction configuration.

Field	Description
HTTP Timeout	Timeout in seconds for sending to and reading from a TCP junction.
HTTPS Timeout	Timeout in seconds for sending to and reading from an SSL junction.
Ping Interval	The interval in seconds between requests which are sent by Reverse Proxy to junctioned Web servers to determine the state of the junctioned Web server.
Ping Method	The HTTP method that Reverse Proxy uses when it sends health check requests to the junctioned Web server.

Field	Description
Ping URI	The URI that Reverse Proxy uses when it sends health check requests to the junctioned Web server.
Maximum Cached Persistent Connections	The maximum number of connections between Reverse Proxy and a junctioned Web server that will be cached for future use.
Persistent Connection Timeout	The maximum length of time, in seconds, that a cached connection with a junctioned Web server can remain idle before it is closed by Reverse Proxy.
Managed Cookie List	A pattern-matched and comma-separated list of cookie names for those cookies which are stored in the Reverse Proxy cookie jar. Other cookies are passed by Reverse Proxy back to the client.

Authentication

The Authentication tab contains entries that are related to the configuration of the authentication mechanisms which are used by the server.

Basic Authentication

Field	Description
Transport	The transport over which basic authentication is supported.
Realm Name	Realm name for basic authentication.

Forms Authentication

Field	Description
Forms Authentication	The transport over which forms authentication is supported.

Client Certificate Authentication

Field	Description
Accept Client Certificates	Defines the condition under which client certificates are required by Reverse Proxy.
Certificate EAI URI	The resource identifier of the application that is invoked to perform external client certificate authentication.
Certificate Data	The client certificate data that are passed to the EAI application.

Kerberos Authentication

Field	Description
Transport	The transport over which Kerberos authentication is supported.
Keytab File	Name of the Kerberos keytab file. The keytab file must contain each of the service principal names used for SPNEGO authentication.

Field	Description
Use Domain Qualified Name	Kerberos authentication provides a principal name in the form of "shortname@domain.com". By default, only the shortname is used as the Security Verify Access user ID. If this checkbox is selected, then the domain is also included as part of the Security Verify Access user ID.
Kerberos Service Names	The list of Kerberos service principal names used for the server. The first service name in the list is the default service name. To make a service name the default, select the service name and then click Default .

EAI Authentication

Field	Description
Transport	The transport over which EAI authentication is supported.
Trigger URL	A URL pattern that is used by Reverse Proxy to determine whether a response is examined for EAI authentication headers.
Authentication Levels	The designated authentication level for each of the configuration authentication mechanisms.

Token Authentication

Field	Description
Transport	The transport over which RSA authentication is supported.

You can also click **Go to RSA Configuration** to access the RSA Configuration page.

OIDC Authentication

Field	Description
Transport	Specifies the transport for which authentication using the OIDC authentication mechanism is enabled.
Redirect URI	The redirect URI which has been registered with the OIDC OP. The redirect URI should correspond to the /pkmsoidc resource of the WebSEAL server (for example: https://isva.ibm.com/pkmsoidc). If no redirect URI is configured it will be automatically constructed from the host header of the request.
Discovery Endpoint	The discovery end-point for the OP. The CA certificate for the discovery-endpoint and corresponding authorization and token endpoints must be added to the WebSEAL key database.

Field	Description
Proxy URL	The URL of the proxy which will be used when communicating with the OP.
Client Id	The Security Verify Access client identity, as registered with the OP.
Client Secret	The Security Verify Access client secret, as registered with the OP.
Response Type	<p>The required response type for authentication responses. The possible values are:</p> <p>code The authorization code flow will be used to retrieve both an access token and identity token.</p> <p>id_token The implicit flow will be used to retrieve the identity token.</p> <p>id_token token The implicit flow will be used to retrieve both an access token and identity token.</p>
Mapped Identity	A formatted string which is used to construct the Security Verify Access principal name from elements of the ID token. Claims can be added to the identity string, surrounded by '{ }'. For example: {iss}/{sub} - would construct a principal name like the following: https://server.example.com/248289761001.
External User	Whether the mapped identity should correspond to a known Security Verify Access identity.
Bearer Token Attributes	The list of JSON data elements from the bearer token response which should be included in the credential as an extended attribute. The JSON name can contain pattern matching characters: '*', '?'. The JSON data name will be evaluated against each rule in sequence until a match is found. The corresponding code (+/-) will then be used to determine whether the JSON data will be added to the credential or not. If the JSON data name does not match a configured rule it will by default be added to the credential.

Field	Description
Id Token Attributes	The list of claims from the ID token which should be included in the credential as an extended attribute. The claim name can contain pattern matching characters: '*', '?'. The claims will be evaluated against each rule in sequence until a match is found. The corresponding code (+/-) will then be used to determine whether the claim will be added to the credential or not. If the claim does not match a configured rule it will by default be added to the credential.

Click the **Load Key** button to load the SSL key for the discovery URI into the WebSEAL key file. This will be achieved by retrieving the root certificate from the server. If the CA certificate is not provided by the server it should be loaded manually into the WebSEAL SSL key file. This operation is not supported when a proxy is configured. In this environment the key should be loaded manually into the SSL key file.

Click the **Test Endpoint** button to see whether the endpoint can be successfully accessed by WebSEAL and that it returns the expected OIDC meta-data.

Session

The Session tab contains entries that are related to the general session configuration.

Field	Description
Re-authentication for Inactive	Whether to prompt users to re-authenticate if their entry in the server credential cache has timed out because of inactivity.
Max Cache Entries	The maximum number of concurrent entries in the session cache.
Lifetime Timeout	Maximum lifetime in seconds for an entry in the session cache.
Inactivity Timeout	The maximum time, in seconds, that a session can remain idle before it is removed from the session cache.
TCP Session Cookie Name	The name of the cookie to be used to hold the HTTP session identifier.
SSL Session Cookie Name	The name of the cookie to be used to hold HTTPS session identifier.
Use Same Session	Select the check box to use the same session for both HTTP and HTTPS requests.

Session Cache

Field	Description
Enable Distributed Sessions	Select the check box to enable distributed sessions on this reverse proxy instance.

Field	Description
Session cache type	Select the type of session cache to be used, either Redis session cache or Distributed session cache. Note: The appliance must be a part of an appliance cluster to enable the distributed session cache. Also, if the cluster configuration changes and a new master is specified, this option must be disabled and then re-enabled. The instance can then pick up the details of the new cluster configuration.
Redis Collections	Specify which of the pre-defined Redis collections (see “Managing the Redis configuration” on page 253) will be used by this Reverse Proxy. The first collection in the list will be set as the default collection.

Response

The Response tab contains entries that are related to response generation.

Field	Description
Enable HTML Redirect	Select the check box to enable the HTML redirect function.
Enable Local Response Redirect	Select the check box to enable the local response redirect function.
Local Response Redirect URI	When local response redirect is enabled, this field contains the URI to which the client is redirected for Reverse Proxy responses.
Local Response Redirect Macros	The macro information which is included in the local response redirect.

SSO

The SSO tab contains entries that are related to the configuration of the different single-sign-on mechanisms that are used by the server.

Failover

Field	Description
Transport	The transport over which failover authentication is supported.
Cookies Lifetime	Maximum lifetime in seconds for failover cookies.
Cookies Key File	The key file which is used to encrypt the failover cookie.

LTPA

Field	Description
Transport	The transport over which LTPA authentication is supported.
Cookie Name	The name of the cookie which is used to transport the LTPA token.
Key File	The key file that is used when accessing LTPA cookies.
Key File Password	The password that is used to access the LTPA key file.

Logging

The Logging tab contains entries that are related to the logging and auditing configuration.

Field	Description
Enable Agent Logging	Select the check box to enable the agent log.
Enable Referer Logging	Select the check box to enable the referrer log.
Enable Request Logging	Select the check box to enable the request log.
Request Log Format	The format of the entries that are contained within the request log.
Maximum Log Size	The maximum size of the log file before it is rolled over.
Flush Time	The period, in seconds, that Reverse Proxy caches the log entries before the system writes the entries to the log file.
Enable Audit Log	Select the check box to enable the generation of audit events.
Audit Log Type	Select the events to be audited.
Audit Log Size	The maximum size of the audit log file before it is rolled over.
Audit Log Flush	The period, in seconds, that Reverse Proxy caches the audit log entries before the system writes the entries to the log file.

Interfaces

The Interfaces tab contains settings that are related to WebSEAL secondary interfaces.

- To add a new secondary interface, click **New**. Then, define your settings in the pop-up window that contains the following fields:

Field	Description
Application Interface IP Address	The IP address on which the WebSEAL instance listens for requests.
HTTP Port	This field contains the port on which the WebSEAL instance listens for HTTP requests.
Enable HTTP Proxy Protocol	Select this check box to enable proxy protocol support for incoming connections to the interface over HTTP.
HTTPS Port	This field contains the port on which the WebSEAL instance listens for HTTPS requests.
Enable HTTPS Proxy Protocol	Select this check box to enable proxy protocol support for incoming connections to the interface over HTTPS.
Web HTTP Port	This is the port that the client perceives WebSEAL to be using.
Web HTTP Protocol	This is the protocol that the client perceives WebSEAL to be using.
Certificate Label	The label of the SSL server certificate that is presented to the client by the WebSEAL instance.
Accept Client Certificates	Defines the condition under which client certificates are required by WebSEAL.

Field	Description
Worker Threads	The number of threads that is allocated to service requests.
HTTP/2	Enables HTTP/2 connection.
HTTP/2 Maximum Connections	The maximum number of HTTP/2 connections allowed per specified port.
HTTP/2 Header Table Size	The size of HTTP/2 header table.
HTTP/2 Maximum Concurrent Streams	The maximum concurrent HTTP/2 streams allowed.
HTTP/2 Initial Window Size	The initial window size of HTTP/2 connections.
HTTP/2 Maximum Frame Size	The maximum frame size of HTTP/2 connections.
HTTP/2 Maximum Header List Size	The maximum header list size of HTTP/2 connections.

Click **Save** to save the settings.

- To delete a secondary interface, select the interface and then click **Delete**.
- To edit a secondary interface, select the interface and click **Edit**. Then, update your settings in the pop-up window that contains the fields that described previously.

5. Click **Save** to apply the changes.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Managing web reverse proxy configuration files

To manage reverse proxy configurations with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Manage > Configuration > Edit Configuration File**.
4. Edit the configuration file that is displayed and then click **Save** to save the changes. If you do not want to save the changes, click **Cancel**. If you want to revert to the previous version of the configuration file, click **Revert**.

Tip: When you are editing the configuration file, you can use the search function of the browser to locate a string. For example, press Ctrl+F.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Exporting WebSEAL configuration

Export the configuration bundle of WebSEAL from the appliance so that you can migrate the WebSEAL instances between different appliances.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.

2. Select the instance of interest.
3. Select **Manage > Configuration > Export Configuration**.
4. Confirm the save operation when your browser displays a confirmation window.

Related tasks

[“Migrating an existing WebSEAL instance to the appliance” on page 183](#)

You can migrate an existing WebSEAL instance to the appliance.

[“Migrating an existing Access Manager software environment to the appliance” on page 186](#)

You can migrate an existing Access Manager software environment to the appliance with the provided mechanism.

Configuring Web Application Firewall

To configure the Web Application Firewall with the local management interface, use the Reverse Proxy management page.

About this task

Important: The Web Application Firewall capability will reach end of service on 31st December, 2022. After this date, no further updates will be made available. Customers can continue to use the capability on an as-is basis, and support will be available for general information and existing functionality only. There will be no defect support available.

The Security Verify Access Web Application Firewall (WAF) can be seen as three modules that flow one after the other, namely:

- Resource filtering
- Issue detection
- Issue response/action

The resource filtering is based on the items that are listed in the **Registered Resources** table on the **Operating Configuration** tab within the **Web Content Protection Configuration** pane. It is a list of URIs, which can include wild cards. If there is a match to any of these, the request then goes to the detection engine.

The issues that the detection engine will check for depends on what items are enabled on the **Issues** tab. You can enable or disable these individually, or click **Trust X-Force** to automatically disable all issues for which there is not a default response. Events that are detected go to the action module.

Lastly, the response/action module specifies what happens when there has been a detection. This is configured in the **Resource Actions** section back on the **Issues** tab. If you do not specify an action in this part, then the specified default action (or 'default response') for this issue will be performed.

The Web Application Firewall logging data is stored in the **pam.log** file. To access this log file:

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the reverse proxy instance.
3. Click **Manage > Logging**. The **Manage Reverse Proxy Log Files** window will be displayed.
4. The log file **pam.log** contains the Web Application Firewall logging data.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the Reverse Proxy instance to configure web application firewall for.
3. Click **Manage > Configuration > Web Content Protection**.
4. On the Operating Configuration tab, you can configure general Web Content Protection settings.
 - a) Select the **Enable Web Content Protection** check box to turn on the web application firewall.

b) To run the firewall in a simulation mode without actually affecting the client traffic, select the **Enable Simulation Mode** check box. When the simulation mode is enabled, any detected issues are audited and then ignored. You can preview the issues that are detected and adjust the settings if necessary before any real actions are taken against the offending requests.

c) Select the **Use Proxy HTTP Header** check box as needed.

This is used to control whether the audit log contains the IP address of the client as obtained from the network connection, or the IP address that is obtained from the x-forwarded-for HTTP header. This setting is useful when a network terminating firewall sits between the reverse proxy and the client.

d) Provide a value in bytes for the **Maximum Memory Size** field. This defines the maximum memory that can be used by the PAM engine.

Note: PAM has a pre-defined minimum memory size. If the configured value is set to less than the minimum, the allocated memory is automatically increased to this minimum size.

e) Under **Resource Actions**:

Note: Use this table to customize the actions that are taken when issues are encountered for a particular resource. This is a pattern-matched list that is searched in order. The resource name can contain the "*" and "?" pattern-matching characters. If no matching resource is found, the default actions, as recommended by the x-force team, are taken.

- To add a resource:

- i) Click **New**.

- ii) On the Add Custom Resource page, provide the resource name. All issues available to the resource are pre-populated.

- Note:** Resource names can contain the "*" and "?" pattern-matching characters. For example, *.html.

- iii) Select an issue that you want to modify and then click **Edit**.

- iv) On the Edit Custom Resource Issue page, select the action to take against this issue in the **Response** field.

- v) *Optional:* If **Quarantine** is selected as the event response in the previous step, specify the quarantine time in the **Quarantine Period** field.

- vi) Click **Save** on the Edit Custom Resource Issue page.

- vii) Click **Save** on the Add Custom Resource page.

- To edit a resource:

- i) Select the resource name to edit.

- ii) Click **Edit**.

- iii) On the Edit Custom Resource page, select the issue that you want to modify and then click **Edit**.

- iv) On the Edit Custom Resource Issue page, modify the event response and quarantine time as needed.

- v) Click **Save** on the Edit Custom Resource Issue page.

- vi) Click **Save** on the Edit Custom Resource page.

- To delete a resource:

- i) Select the resource name to delete.

- ii) Click **Delete**.

- Note:** There is no confirmation window for this delete operation. Make sure that the selected resource is the one you want to delete before you click **Delete**.

f) Under **Registered Resources**:

Note: The registered resources are used to designate the requests that are passed to the inspection engine. When a request is received by the Web reverse proxy, the entries in the list is sequentially searched until a match is found. The action that is assigned to the matching resource controls whether the inspection is enabled or disabled. The resources can contain wildcard characters for pattern matching.

- To add a registered resource:

- i) Click **New**.
- ii) On the Add Protected Resources page that pops up, provide the **Resource Name**. For example, index.html, *.html or *.gif.
- iii) Select **Enabled** or **Disabled** as needed.
- iv) Click **Save**.

- To edit a registered resource:

- i) Select the resource to edit from the list.
- ii) Click **Edit**.
- iii) On the Edit Protected Resources page that pops up, modify the resource name and whether it is enabled as needed.
- iv) Click **Save**.

- To delete a registered resource

- i) Select the resource to delete from the list.
- ii) Click **Delete**.

Note: There is no confirmation window for this delete operation. Make sure that the selected resource is the one you want to delete before you click **Delete**.

- g) Under **Injection Tuning Parameters**, modify the listed parameters by double-clicking a value in the **Units** column and editing inline as needed. To see a description of each parameter, hover your mouse cursor on that parameter and a pop-up message that contains the description is displayed.

5. On the Issues tab, you can enable or disable certain issues.

Note: The list of issues control the events that are monitored by the inspection engine. If an issue is disabled, the inspection engine no longer checks for this issue.

- Approach 1:
 - a. Select the event to edit.
 - b. Click **Edit**.
 - c. On the Edit Issue page, select **Enabled** or **Disabled** as needed.
 - d. Click **Save**.
- Approach 2:
 - Select or clear the **Enabled** check box to enable or disable a particular issue.
- Approach 3:
 - Click **Trust X-Force** to automatically disable all issues for which there is not a default response.

6. On the Audit tab, you can configure logging and auditing settings.

- a) Under **Log detailed audit events**, select the check box if you want to enable logging for detailed audit events.
- b) Under **Log Audit Events**, select one of the options to indicate where the audit events are sent.
- c) Under **Log Audit Config**, define the following parameters based on the selections made in the previous step.
 - If **Log to File** is selected:

Parameter	Description
File Name	The entry specifies the name of the log file.
Rollover Size	The maximum size to which a log file can grow before it is rolled over. The default value is 2000000 bytes.
Buffer Size	The maximum size of the message that is used when smaller events are combined.
Queue Size	There is a delay between events being placed on the queue and the file log agent removing them. This parameter specifies the maximum size to which the queue is allowed to grow.
High Water Mark	Processing of the event queue is scheduled regularly at the configured flush interval. It also is triggered asynchronously by the queue size reaching a high water mark on the event queue. The default value is two-thirds of the maximum configured queue size. If the maximum queue size is zero, the high water mark is set to a default of 100. If the event queue high water mark is set to 1, every event queued is relayed to the log agent as soon as possible.
Flush Interval	This entry controls the frequency with which the server asynchronously forces a flush of the file stream to disk. The value defined for this parameter is 0, < 0, or the flush interval in seconds.

- If **Log to Remote Authorization Server** is selected:

Parameter	Description
Compress	To reduce network traffic, use this parameter to compress buffers before transmission and expand on reception. The default value is no.
Buffer Size	To reduce network traffic, events are buffered into blocks of the nominated size before they are relayed to the remote server. This parameter specifies the maximum message size that the local program attempts to construct by combining smaller events into a large buffer. The default value is 1024 bytes.
Flush Interval	This parameter limits the time that a process waits to fill a consolidation buffer. The default value is 20 seconds. A flush interval of 0 is not allowed. Specifying a value of 0 results in the buffer being flushed every 600 seconds.
Queue Size	There is a delay between events being placed on the queue and the file log agent removing them. This parameter specifies the maximum size to which the queue is allowed to grow.
High Water Mark	Processing of the event queue is scheduled regularly at the configured flush interval. It also is triggered asynchronously by the queue size reaching a high water mark on the event queue. The default value is two-thirds of the maximum configured queue size. If the maximum queue size is zero, the high water mark is set to a default of 100. If the event queue high water mark is set to 1, every event queued is relayed to the log agent as soon as possible.
Error Retry Timeout	If a send operation to a remote service fails, the system tries again. Before the system tries again, it waits for the error retry timeout in seconds. The default value is 2 seconds.
Logging Port	Configure the port parameter to specify the port that the remote authorization server listens on for remote logging requests. The default value is port 7136.
Rebind Retry	If the remote authorization server is unavailable, the log agent attempts to rebind to this server at this frequency in number of seconds. The default rebind retry timeout value is 300 seconds.

Parameter	Description
Hostname	The remote logging services are offered by the authorization service. The server parameter nominates the hosts to which the authorization server process is bound for event recording.
DN	To establish mutual authentication of the remote server, a distinguished name (DN) must be configured. A distinguished name must be specified as a string that is enclosed by double quotation marks.

- If **Log to Remote Syslog Server** is selected:

Parameter	Description
Remote Syslog Server	The host to which the syslog server process is bound for event recording.
Port	The port on which the remote syslog server listens for remote logging requests.
Application ID	The name of the application, as it appears in the messages that are sent to the remote syslog server.
Error Retry Timeout	If a send operation to a remote service fails, the system tries again. Before the system tries again, it waits for the error retry timeout in seconds. The default value is 2 seconds.
Flush Interval	This parameter limits the time that a process waits to fill a consolidation buffer. The default value is 20 seconds. A flush interval of 0 is not allowed. Specifying a value of 0 results in the buffer being flushed every 600 seconds.
High Water Mark	Processing of the event queue is scheduled regularly at the configured flush interval. It also is triggered asynchronously by the queue size reaching a high water mark on the event queue. The default value is two-thirds of the maximum configured queue size. If the maximum queue size is zero, the high water mark is set to a default of 100. If the event queue high water mark is set to 1, every event queued is relayed to the log agent as soon as possible.
Queue Size	There is a delay between events being placed on the queue and the file log agent removing them. This parameter specifies the maximum size to which the queue is allowed to grow.
Rebind Retry	If the remote system log server is unavailable, the log agent attempts to rebind to this server at this frequency in number of seconds. The default rebind retry timeout value is 300 seconds.
Maximum Event Length	The maximum length of an event to be transmitted to the remote syslog server. If the event text is longer than the configured length, it is truncated to the maximum event length. If the maximum event length is zero, the event text is never truncated. If transmitting the event to the remote syslog server in clear text, set the maximum event length to less than the maximum transmission unit (MTU) for the network path to the server. This avoids fragmentation of the event.
Enable SSL Communication	Whether SSL is be used for communication.
SSL Keyfile	The name of the GSKit key database file that contains the CA certificate. It is used when the system establishes a secure connection with the remote syslog server over TLS. If the Enable SSL Communication check box is selected, this field is required.

Parameter	Description
SSL Certificate Label	The name of the certificate to be presented to the remote syslog server, upon request, when the system establishes a secure connection. If no value is set for this field, the default certificate from the key database is used.

7. On the Advanced Configuration tab, you can configure coalescer, inspection engine, issues, and custom actions.

a) Under **Coalescer Configuration**:

Note: The coalescer is used to correlate audit events. The administrator can use these configuration settings to fine-tune the processing of the coalescer and thus reduce the number of messages that are sent to the audit log.

- To add a coalescer parameter:

- i) Click **New**.
- ii) On the Add Coalescer Parameter page that pops up, provide the parameter name and value.
- iii) Click **Save**.

- To edit a coalescer parameter:

- i) Select the parameter to edit from the list.
- ii) Click **Edit**.
- iii) On the Edit Coalescer Parameter page that pops up, modify the parameter name and value as needed.
- iv) Click **Save**.

- To delete a coalescer parameter:

- i) Select the parameter to delete from the list.
- ii) Click **Delete**.

Note: There is no confirmation window for this delete operation. Make sure that the selected parameter is the one you want to delete before you click **Delete**.

b) Under **Inspection Engine Configuration**:

- To add a inspection engine configuration parameter:

- i) Click **New**.
- ii) On the Add Inspection Parameter page that pops up, provide the parameter name and value.
- iii) Click **Save**.

- To edit a inspection engine configuration parameter:

- i) Select the parameter to edit from the list.
- ii) Click **Edit**.
- iii) On the Edit Inspection Parameter page that pops up, modify the parameter name and value as needed.
- iv) Click **Save**.

- To delete a inspection engine configuration parameter:

- i) Select the parameter to delete from the list.
- ii) Click **Delete**.

Note: There is no confirmation window for this delete operation. Make sure that the selected resource is the one you want to delete before you click **Delete**.

8. Click **Save**.

Managing administration pages

To manage files and directories in the administration pages root with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Manage > Management Root**.

All current management files and directories are displayed. The default directories include:

management

The Web Reverse proxy management pages. For example, `login.html`

errors

The error pages that can be returned by the Web Reverse proxy.

oauth

The HTML files that can be returned by the oauth module.

junction-root

The static HTML files that are served by the local junction of the Web Reverse proxy.

Note: A fixed location is used as the document root. A local junction cannot run any CGI scripts. It can serve only static page content.

4. Work with all the management files and directories.

- **Create a new file in the administration pages root**

- a. Select the directory in which you want to create the file.
- b. Select **File > New > File**.
- c. Enter the file name.
- d. Optionally, you can add file contents in the **New File Contents** field.
- e. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

- **Create a new directory in the administration pages root**

- a. Select the directory in which to create the directory.
- b. Select **File > New > Directory**.
- c. Enter the directory name.
- d. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

- **View or update the contents of a file in the administration pages root**

- a. Select the file of interest.
- b. Select **File > Open**. You can then view the contents of the file.
- c. Optionally, edit the contents of the file. Then, click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

- **Export a file from the administration pages root**

- a. Select the file of interest.
- b. Select **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

c. Confirm the save operation when your browser displays a confirmation window.

- **Rename a file or directory in the administration pages root**

a. Select the file or directory of interest.

b. Select **Manage > Rename**.

c. Enter the new name of the file or directory in the **New Resource Name** field.

d. Click **Save**.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

- **Delete a file or directory in the administration pages root**

a. Select the file or directory of interest.

b. Select **Manage > Delete**.

c. Click **Yes** to confirm the delete operation.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

- **Import a file to administration pages root**

a. Select the directory that you want to import the file into.

b. Select **Manage > Import**.

c. Click **Browse**.

d. Browse to the file you want to import and then click **Open**.

e. Click **Import**.

- **Import the contents of a .zip file into the administration pages root**

a. Select **Manage > Import Zip**.

b. Click **Browse**.

c. Browse to the .zip file you want to import and then click **Open**.

d. Click **Import**.

- **Export the contents of the administration pages root as a .zip file**

a. Select **Manage > Export Zip**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

b. Confirm the save operation when your browser displays a confirmation window.

Renewing web reverse proxy management certificates

Renew the management certificate of a web reverse proxy instance.

About this task

An SSL certificate is used to authenticate the web reverse proxy instance to the policy server. Use this option to automatically generate a new certificate that can be used in this communication.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance to update the management certificate for.
3. Select **Manage > Renew Management Certificate**.

4. Enter your administrator name and password.
5. Click **Renew**.

Configuring Mobile Multi-Factor Authentication

Configure Mobile Multi-Factor Authentication (MMFA) for a specific Web Reverse Proxy instance.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance to configure Mobile Multi-Factor Authentication for.
3. Select **Manage > MMFA Configuration**.
4. On the **Main** tab, select the type of traffic you want to apply MMFA to.
5. On the **AAC LMI** tab, provide the following details and then click **Next**.

Host name

The host name or IP address of the LMI server. This field is automatically populated with values from the current browser window.

Port

The port number of the LMI server. This field is automatically populated with values from the current browser window.

Username

The user name that is used to authenticate with the LMI server. The default value is `admin`.

Password

The password that is used to authenticate with the LMI server.

6. On the **AAC Runtime** tab, provide the following details and then click **Next**.

Host name

The host name or IP address of the runtime server. The default value is `localhost`.

Port

The port number of the runtime server. The default value is `443`.

Username

The user name that is used to authenticate with the runtime server. The default value is `easuser`.

Password

The password that is used to authenticate with the runtime server.

7. On the **Reuse Options** tab, provide the following details and then click **Next**.

Reuse certificates

Select to reuse the SSL certificate if it was already saved. If this check box is not selected, the certificate is overwritten.

Reuse ACLs

Select to reuse any existing ACLs with the same name. If this check box is not selected, the ACLs are replaced.

8. Click **Finish**.

Chapter 16. Reverse proxy status

You can use the local management interface (LMI) to manage status and view statistics.

Showing the current state of all instances

To show the current state of all instances with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. You can view the current state and version information of all instances.

Modifying the statistics settings for a component

To modify the statistics settings for a particular component with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Troubleshooting > Statistics**.
4. Select the statistics component that you want to modify.
5. Click **Edit**.
6. Select the check box beside **Enabled** if it is not already checked.
7. Modify the **Interval**, **Count**, **Flush Interval**, **Rollover Size**, **Maximum Rollover Files**, and **Compress** fields as needed.

By default, the **Compress** option is set to **No**. To save disk space, set the **Compress** option to **Yes** so that all rollover files are automatically compressed.

8. Click **Save** to save your changes.

Managing statistics log files

To manage statistics log files with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Troubleshooting > Statistics**.
4. Select the statistics component of interest.
5. Click **Files**.

The file name, file size, and last modified time information of all statistics log files is displayed.

- **View a statistics log file or a snippet of a statistics log file**
 - a. Select the statistics log file that you want to view and then click **View**. The contents of the statistics log file are displayed.

- b. You can enter a value into the **Number of lines to view** field and then click **Reload** to get a customized snippet view of the log file. Optionally, you can provide a value in the **Starting from line** field to define the start of the lines. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

- **Export a statistics log file**

- a. Select the statistics log file that you want to export.
- b. Click **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

- c. Confirm the save operation in the browser window displayed.

- **Delete a statistics log file**

- a. Select the statistics log file or files that you want to delete and then click **Delete**.

Note: Only log files that are not in use can be deleted. To disable a log file, you can select the log file, click **Edit**, clear the **Enabled** check box, and then click **Save**.

- b. Click **Yes** to confirm the operation.

- **Delete all unused statistics log files**

- a. Click **Manage > Delete All**.

- b. Click **Yes** to confirm the operation.

Archiving and deleting reverse proxy log files with the command-line interface

Use the `logs` option in the command-line interface to archive Web Reverse Proxy log files to a USB device and then delete old log files to free up disk space.

Procedure

1. In the command-line interface, go to **isam > logs**.
2. *Optional:* Enter `help` to display all available commands.

```
Current mode commands:
archive           Archive the log files to a USB device.
delete           Delete the log files which have been rolled over by the system.
Global commands:
back             Return to the previous command mode.
exit            Log off from the appliance.
help            Display information for using the specified command.
reboot          Reboot the appliance.
shutdown        End system operation and turn off the power.
top             Return to the top level.
```

3. Archive or delete the log files.

- **Archive the log files to a USB device**

- a. Enter `archive` to save the log files to a USB device.
- b. Insert a USB device into the USB port of the appliance.

- c. Enter YES to start the archive operation. A list of archived files are displayed, along with a message that indicates when the archive operation has completed. Example output is shown as follows:

```
updating: var/PolicyDirector/log/ (stored 0%)
updating: var/PolicyDirector/log/msg__pdmgrd_utf8.log (deflated 85%)
updating: var/PolicyDirector/log/PDMgr_config_start.log (deflated 37%)
updating: var/PolicyDirector/log/ivmgrd.pid (stored 0%)
updating: var/pdweb/default/log/ (stored 0%)
updating: var/pdweb/default/log/iss-pam1.so (deflated 59%)
updating: var/pdweb/default/log/webseald-default.pid (stored 0%)
updating: var/pdweb/default/log/config_data__default
-webseald-felbb.wga.gc.au.ibm.com.log (deflated 92%)
updating: var/pdweb/default/log/referer.log (stored 0%)
updating: var/pdweb/default/log/msg__webseald-default.log (deflated 89%)
updating: var/pdweb/default/log/pam.log (deflated 98%)
updating: var/pdweb/default/log/agent.log (stored 0%)
updating: var/pdweb/default/log/request.log (stored 0%)
The log files have been successfully archived to the USB drive:
iswga_logs.zip. It is now safe to remove the USB drive.
```

- d. Remove the USB device from the USB port.

- **Delete the log files**

- a. Enter delete to purge all log files that are rolled over.
- b. Enter YES to confirm the delete operation.

Viewing reverse proxy traffic

To view flow data at an instance-specific level with the local management interface, use the Reverse Proxy Traffic management page.

Procedure

1. From the top menu, select **Monitor > Reverse Proxy Graphs > Reverse Proxy Traffic**.
2. On the Reverse Proxy Traffic page, specify the settings for the chart displayed.

Instance

The instance which the data displayed are specific to.

Aspect Type

The type of chart to display the data with. Select one from **Column and Lines**, **Column**, and **Lines**.

Start Date

The starting date.

Start Time

The starting time of the day.

Date Range

The duration over which data is collected and displayed. Select from **1 Hour** to **30 Days**.

For example, if the date and time that is chosen is 04.12.2012 10.00 and the duration is 12 Hours, the data that are collected between 10:00 a.m. and 10:00 p.m. on 12th April 2012 are displayed.

By default, data of the first instance in the instance list for the last 24 hours are displayed, grouped by junction.

Viewing reverse proxy throughput

To view flow data at an appliance-wide level with the local management interface, use the Reverse Proxy Throughput management page or the Reverse Proxy Throughput widget on the dashboard.

Procedure

1. To view the Reverse Proxy Throughput:

- From the dashboard, locate the Reverse Proxy Throughput widget.
 - From the top menu, select **Monitor > Reverse Proxy Graphs > Reverse Proxy Throughput**.
2. Specify the settings for the chart displayed.

- On the dashboard, select the duration over which data is collected and displayed with the **Data Range** list.
- On the Reverse Proxy Throughput page, use the following settings:

Chart Type

The type of chart to display the data with. Select one from **Column and Lines**, **Column**, and **Lines**.

Date Range

The duration over which data is collected and displayed. Select from **1 Hour** to **30 Days**.

Start Date

The starting date.

Start Time

The starting time of the day.

For example, if the date and time that is chosen is 04.12.2012 10.00 and the duration is 12 Hours, the data that are collected between 10:00 a.m. and 10:00 p.m. on 12 April 2012 are displayed.

By default, data of all configured WebSEAL instances on this appliance from the last 24 hours are displayed.

Viewing reverse proxy health status

The health status of a reverse proxy is determined by the state of instances, junctions, and junctioned servers. You can view the health status information with the Reverse Proxy Health dashboard widget.




Procedure

1. From the dashboard, locate the Reverse Proxy Health widget.

The health status of each instance, its junctions, and the junctioned servers are displayed in a hierarchical structure. Health status is determined by the health of all elements lower than the current element in the hierarchy.

- An instance is unhealthy if it is stopped or pdadmin cannot contact it.
- A junction is unhealthy if it is disabled or pdadmin cannot return information for it.
- A junctioned server is unhealthy if it is disabled or offline.

Each element can be in one of the three health states:

Icon	State	Description
	Healthy	All child elements are healthy.
	Warning	The element contains at least one unhealthy child element and at least one healthy child element.
	Unhealthy	All child elements are unhealthy.

2. *Optional:* Click **Refresh** to refresh the health data.




Viewing front-end load balancer health status

The health status of a front-end load balancer is determined by the state of the load balanced servers. You can view the health status information with the Load Balancer Health dashboard widget.

Procedure

1. From the dashboard, locate the Load Balancer Health widget.
 - Under **High Availability** (if high availability is configured):
 - The first row displays the health status of the self front-end load balancer and whether it is active or passive.
 - The second row displays the health status of the peer front-end load balancer and whether it is active or passive.
 - Under **Services** (if at least one service is configured):
 - The health status of the configured services and the load balanced servers are displayed in a hierarchical structure. You can expand a service to view the health status of the servers that are attached to this service.

Each element can be in one of the following health states:

Icon	State	Description
	Healthy	All child elements are healthy.
	Warning	The element contains at least one unhealthy child element and at least one healthy child element.
	Unhealthy	All child elements are unhealthy.

2. *Optional:* Click **Refresh** to refresh the health data.

Viewing average response time statistics

The Web Reverse Proxy can be configured to record transaction logs. One of the attributes that is recorded is the average request response time. This information is recorded at a per-junction level. To view a summary of the average response time that has been recorded, use the Average Response Time widget.

Procedure

1. From the dashboard, locate the Average Response Time widget.

The average response time for requests is displayed on a graph.

Note: The widget is only displayed if one or more Reverse Proxy instances have the Flow Data function enabled.
2. Under **Reverse Proxy Instances**, select the instance to view the average response time statistics for.
3. Under **Junctions**, select the junctions to display on the graph. Each junction is represented by a separate line on the graph.
4. Under **Date Range**, select the duration over which the response times are recorded.

Viewing security action statistics

The Web Reverse Proxy can be configured to perform inspections on web content, searching for potential malicious requests (known as issues). It can then take certain defensive actions against any discovered

issues. A summary of the defensive actions that have been taken can be viewed by using the Security Actions widget.

Procedure

1. From the dashboard, locate the Security Actions widget. The number of times each defensive action has been taken is displayed in a graph.

Note: The widget is only displayed if one or more instances have the security statistics function enabled.

2. Under **Reverse Proxy Instances**, select the instances to view action statistics for.

Note: Only instances that have security statistics function enabled are listed for selection.

3. Under **Actions**, select the actions to be included in the statistics. The number of actions that are displayed is the total of all selected actions.
4. Under **Date Range**, select the duration over which the actions are taken.

Chapter 17. Junctions

Creating virtual junctions

Use the Junction Management page to create one or more virtual junctions in your environment.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the reverse proxy to manage junctions for.
3. Select **Manage > Junction Management**.
4. Click **New > Virtual Junction**.
5. On the Junction tab page:
 - a) Enter the junction label in the **Junction Label** field.
 - b) Select the **Stateful Junction** check box if you want the junction to be stateful.
 - c) Select the **HTTP/2 Junction** check box if you want to enable HTTP/2 protocol to the junction server.
 - d) Select the **HTTP/2 Proxy** check box if you want to enable HTTP/2 protocol to the proxy server.
 - e) Specify the **Server Name Indicator (SNI)**.
 - f) Select a junction type from the listed options on the right.

Notes for HTTP/2 junctions:

 - The protected Web Server must serve HTTP/2 over both TCP and SSL for WebSEAL mutual junction type with HTTP/2 to work. For example, Microsoft IIS only serves HTTP/2 over SSL. So an HTTP/2 mutual junction type cannot be created to an IIS Web Server.
 - TCP HTTP/2 junction connections do not use HTTP/2 upgrade. They require the "Prior Knowledge" method to connect to an HTTP/2 Web Server over TCP. In Apache configuration terms, this is the "Direct mode".
6. On the Servers tab page:
 - a) Click **New** to add a target back-end server. At least one target back-end server must be added to create a junction.
 - b) Complete the fields displayed.
 - c) Click **Save**.
7. On the Basic Authentication tab page:
 - a) Select the **Enable Basic Authentication** check box if BA header information is to be used for authentication with the back-end server.
 - b) Enter the WebSEAL user name in the **Username** field.
 - c) Enter the WebSEAL password in the **Password** field.
 - d) Select the **Enable mutual authentication to junctioned WebSEAL servers** check box if mutual authentication is to be used between a frontend WebSEAL server and a back-end WebSEAL server.
 - e) Select the key file from the list to use for mutual authentication.
 - f) Select the key label from the list to use for mutual authentication.
8. On the Identity tab page:
 - a) Define how WebSEAL server passes client identity information in BA headers to the back-end server by selecting appropriate actions from the list under **HTTP Basic Authentication Header**.

- b) If **GSO** is selected in the previous step, enter the GSO resource or resource group name in the **GSO Resource or Group** field. If a value other than **GSO** is selected in the previous step, skip this step.
 - c) Select what HTTP header identity information is passed to the back-end server in the **HTTP Header Identity Information** field.
 - d) Select encoding from the list under **HTTP Header Encoding**.
 - e) Select the check box on the right as necessary.
9. On the SSO and LTPA tab page:
- a) Select the **Enable LTPA cookie Support** check box if the junctions are to support LTPA cookies.
 - b) If LTPA version 2 cookies (LtpaToken2) are used, select the **Use Version 2 Cookies** check box.
 - c) Select the LTPA keyfile from the list under **LTPA Keyfile**.
 - d) Enter the keyfile password in the **LTPA Keyfile Password** field.
10. On the General tab page:
- a) Specify the name of the form based single sign-on configuration file in the **FSSO Configuration File** field.
 - b) Define the hard limit for consumption of worker threads in the **Percentage Value for Hard Limit of Worker Threads** field.
 - c) Define the soft limit for consumption of worker threads in the **Percentage Value for Soft Limit of Worker Threads** field.
 - d) If you want denied requests and failure reason information from authorization rules to be sent in the Boolean Rule header, select the **Include authorization rules decision information** check box.
 - e) Click **Save**.

Creating standard junctions

Use the Junction Management page to create one or more standard junctions in your environment.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the reverse proxy to manage junctions for.
3. Select **Manage > Junction Management**.
4. Click **New > Standard Junction**.
5. On the Junction tab page:
 - a) Enter the junction point name. Names for standard junctions must start with a forward slash (/) character.
 - b) Select the **Create Transparent Path Junction** check box if the junction name must match the name of a subdirectory under the root of the back-end server document space.
 - c) Select the **Stateful Junction** check box if you want the junction to be stateful.
 - d) Select the **HTTP/2 Junction** check box if you want to enable HTTP/2 protocol to the junction server.
 - e) Select the **HTTP/2 Proxy** check box if you want to enable HTTP/2 protocol to the proxy server.
 - f) Specify the **Server Name Indicator (SNI)**.
 - g) Select a junction type from the listed options on the right.

Notes for HTTP/2 junctions:

- The protected Web Server must serve HTTP/2 over both TCP and SSL for WebSEAL mutual junction type with HTTP/2 to work. For example, Microsoft IIS only serves HTTP/2 over SSL. So an HTTP/2 mutual junction type cannot be created to an IIS Web Server.

- TCP HTTP/2 junction connections do not use HTTP/2 upgrade. They require the "Prior Knowledge" method to connect to an HTTP/2 Web Server over TCP. In Apache configuration terms, this is the "Direct mode".
6. On the Servers tab page:
 - a) Click **New** to add a target back-end server. At least one target back-end server must be added to create a junction. The options available when you add a server vary depending on the junction type selected.
 - b) Complete the fields displayed.
 - c) Click **Save**.
 7. On the Basic Authentication tab page:

Note: The properties on this tab are specific to SSL junctions. They are available only if you create an SSL junction.

 - a) Select the **Enable Basic Authentication** check box if BA header information is to be used for authentication with the back-end server.
 - b) Enter the WebSEAL user name in the **Username** field.
 - c) Enter the WebSEAL password in the **Password** field.
 - d) Select the **Enable mutual authentication to junctioned WebSEAL servers** check box if mutual authentication is to be used between a frontend WebSEAL server and a back-end WebSEAL server.
 - e) Select the key file from the list to use for mutual authentication.

Note: The options in the list include certificates from both the local and network key files. The certificates from the network key file are prefixed with the token label for the network HSM device.
 8. On the Identity tab page:
 - a) Define how WebSEAL server passes client identity information in BA headers to the back-end server by selecting appropriate actions from the list under **HTTP Basic Authentication Header**.
 - b) If **GSO** is selected in the previous step, enter the GSO resource or resource group name in the **GSO Resource or Group** field. If a value other than **GSO** is selected in the previous step, skip this step.
 - c) Select what HTTP header identity information is passed to the back-end server in the **HTTP Header Identity Information** field.
 - d) Select encoding from the list under **HTTP Header Encoding**.
 - e) Select an option from the list under **Junction Cookie Javascript Block**.
 - f) Select the check box on the right as necessary.
 9. On the SSO and LTPA tab page:
 - a) Select the **Enable LTPA cookie Support** check box if the junctions are to support LTPA cookies.
 - b) If LTPA version 2 cookies (LtpaToken2) are used, select the **Use Version 2 Cookies** check box.
 - c) Select the LTPA keyfile from the list under **LTPA Keyfile**.
 - d) Enter the keyfile password in the **LTPA Keyfile Password** field.
 10. On the General tab page:
 - a) Specify the name of the form based single sign-on configuration file in the **FSSO Configuration File** field.
 - b) Define the hard limit for consumption of worker threads in the **Percentage Value for Hard Limit of Worker Threads** field.
 - c) Define the soft limit for consumption of worker threads in the **Percentage Value for Soft Limit of Worker Threads** field.
 - d) If you want denied requests and failure reason information from authorization rules to be sent in the Boolean Rule header, select the **Include authorization rules decision information** check box.
 11. Click **Save**.

Managing standard and virtual junctions

To manage standard and virtual junctions with the local management interface, use the **Junction Management** page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the reverse proxy to manage junctions for.
3. Select **Manage > Junction Management**.
4. Perform junction-related tasks as needed.

- **Create standard junctions**

See [“Creating standard junctions”](#) on page 232.

- **Create virtual junctions**

See [“Creating virtual junctions”](#) on page 231.

- **Edit a standard or virtual junction**

- a. Select the junction to edit from the list.
- b. Click **Edit**.
- c. Modify the settings as needed.
- d. Click **Save**.

- **Delete a standard or virtual junction**

- a. Select the junction to delete from the list.
- b. Click **Delete**.
- c. In the confirmation window that pops up, click **Yes**.

Note: Some junction management tasks can be performed only with the web service, but not the local management interface. For example, functions achieved by using the following web service commands cannot be achieved by using the local management interface:

- jmt load
- jmt clear
- offline
- online
- throttle

Chapter 18. Federation management

Use the local management interface to configure your federations with a reverse proxy server.

Adding a federation for a reverse proxy server

Configure a federation on a reverse proxy server to set up access between the federation and reverse proxy appliances.

Before you begin

The reverse proxy server that you want to use for your federations must already be configured. See [“Configuring an instance” on page 206](#).

Procedure

1. From the local management interface, select **Web > Manage > Reverse Proxy**.
A list of reverse proxy instances displays.
2. Select the reverse proxy instance name from the list.
3. Select **Manage > Federation Management**.
A list of federations configured for this reverse proxy instance displays.
4. Click **Add**.
A window opens where you can add the configuration information.
5. Enter the configuration details for the federation.

The **Runtime** tab provides authentication information for the federation runtime:

Host name

The host name or IP address of the runtime server. This field is required.

Port

The SSL port number of the runtime server. This field is required.

User name

The user name that is used to authenticate with the runtime server. This field is required.

Password

The password that is used to authenticate with the runtime server. This field is required.

The **Federation** tab specifies the federation name:

Federation Name

The name that identifies the federation that you are configuring on this reverse proxy instance. Select the correct name from the list. If the federation name is not in the list, ensure that you set up the runtime configuration properly for that federation.

The **ACLs and Certificates** tab indicates reuse of existing access control lists (ACLs) and certificates:

Reuse ACLs

Select to reuse any existing ACLs with the same name. If this check box is not selected, the ACLs are replaced.

Reuse Certificates

Select to reuse the SSL certificate if it was already saved. If this check box is not selected, the certificate is overwritten.

6. Click **Submit**.

Removing a federation from a reverse proxy server

You can remove a federation that was configured for a reverse proxy server.

Procedure

1. From the local management interface, select **Web > Manage > Reverse Proxy**.
A list of reverse proxy instances displays.
2. Select the reverse proxy instance name from the list.
3. Select **Manage > Federation Management**.
A list of federations configured for this reverse proxy instance displays.
4. Select the federation name from the list.
5. Click **Remove**.
A pop-up window is displayed for confirmation.
6. Click **Yes**.

Chapter 19. Authorization servers

To manage IBM Security Verify Access authorization server instances, go to **Web > Manage > Authorization Server**.

Cleaning up authorization servers

After you import a migration bundle, some authorization server instances might no longer be relevant to your current environment. In such situation, you can use the cleanup function on the Runtime Component management page to remove these instances.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Click **Manage > Cleanup Servers**.
3. In the pop-up window, enter you IBM Security Verify Access administrator user name and password. These are the same user name and password you would use with the pdadmin utility.
4. Click **Login**.
5. From the list of authorization servers, select the one to be removed.

Note: A red icon indicates that the server is uncontactable. Stopping a server also renders it uncontactable. Make sure that you select only the instance that is no longer relevant in your current environment and thus should be removed.

6. Click **Delete**.

Note: The **Delete** button is only clickable when an uncontactable server with a red icon is selected. After you delete an instance, all knowledge of this instance is removed from the policy server including LDAP.

7. In the confirmation window, click **Yes** to confirm the operation.

Creating an authorization server instance

To create an authorization server instance, use the **Authorization Server** management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
The status of all authorization server instances is displayed.
2. Click **New**.
3. In the **New Authorization Server Instance** window, provide values for the displayed fields.
 - On the **Instance** tab, define the following fields.

Field	Description
Instance Name	Name of the authorization server instance.
Host Name	Name of the local host. The name is used during the construction of the authorization server instance name. The default value is the host name of the local system.
Authorization Port	The port over which authorization requests are received. The default value is the next available port from 7136.

Field	Description
Administration Port	The port over which Security Verify Access administration requests are received. The default value is the next available port after the authorization port value.
IP Addresses	The IP addresses on which the authorization server listens for requests. To add an IP address to the selected box, select the address from the list immediately under IP Addresses and then click Add . To remove an IP address from the selected list, select the address from the box and then click Remove .

- On the **IBM Security Verify Access** tab, define the following fields.

Field	Description
Administrator Name	The administrator user name of IBM Security Verify Access.
Administrator Password	The administrator user password of IBM Security Verify Access.
Domain	The domain name of IBM Security Verify Access.

- If you use an LDAP server that is external to the appliance, a **User Registry** tab is also displayed. On the **User Registry** tab, define the following fields.

Field	Description
Enable SSL	Specifies whether to enable SSL communication between the instance and the LDAP server.
Key File Name	The file that contains the LDAP SSL certificate. This field is only valid if the Enable SSL check box is selected.
Certificate Label	The LDAP client certificate label. This field is only valid if the Enable SSL check box is selected.
Port	The port number through which to communicate with the LDAP server. This field is only valid if the Enable SSL check box is selected.

4. Click **Finish**.

Deleting an authorization server instance

To delete an authorization server instance, use the **Authorization Server** management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
The status of all authorization server instances is displayed.
2. Select the instance to delete.
3. Click **Delete**.
4. In the **Delete Authorization Server Instance** window, enter the administrator name and password.
5. Optional: If you want to unconfigure the instance even if the policy server is unreachable, select the **Force** check box.
6. Click **Delete** to confirm the operation.

Stopping, starting, or restarting an authorization server instance

To stop, start, or restart an authorization server instance, use the Authorization Server management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
2. Select the instance of interest.

Stop an instance

- a. Click **Stop**.
- b. A message is displayed indicating that the instance is stopped successfully.

Start an instance

- a. Click **Start**.
- b. A message is displayed indicating that the instance is started successfully.

Restart an instance

- a. Click **Restart**.
- b. A message is displayed indicating that the instance is restarted successfully.

Editing an authorization server instance advanced configuration file

To edit an authorization server instance advanced configuration file, use the Authorization Server management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
2. Select the instance of interest.
3. Select **Manage > Configuration > Edit Configuration File**.
The configuration file contents are displayed.
4. In the **Advanced Configuration File Editor** window, modify the configuration file.
5. Click **Save** to save the changes. If you want to revert to the last successfully saved version of this file, click **Revert**. Or click **Cancel** if you do not want to save the changes.

Note: For the changes to take effect, the changes must be deployed and the running instance must be restarted.

Editing an authorization server instance tracing configuration file

To edit an authorization server instance tracing configuration file, use the Authorization Server management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
2. Select the instance of interest.
3. Select **Manage > Configuration > Edit Tracing Configuration File**.
The tracing configuration file contents are displayed.
4. In the **Tracing Configuration File Editor** window, modify the file.
5. Click **Save** to save the changes. Or click **Cancel** if you do not want to save the changes.

Note: For the changes to take effect, the changes must be deployed and the running instance must be restarted.

Renewing authorization server management certificates

Renew the management certificate of an authorization server instance.

About this task

An SSL certificate is used to authenticate the authorization server instance to the policy server. Use this option to automatically generate a new certificate that can be used in this communication.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
2. Select the instance to update the management certificate for.
3. Select **Manage > Renew Management Certificate**.
4. Enter your administrator name and password.
5. Click **Renew**.

Chapter 20. Clusters

Replicating runtime settings across the cluster

In a cluster environment, enable this option on the primary master to replicate the IBM Security Verify Access runtime settings to the non-primary nodes.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Select the **Replicate with Cluster** check box.

Note: This option is selectable on the primary master of the cluster only.

3. In the confirmation window, click **Yes** to confirm the operation.

The current IBM Security Verify Access runtime settings of the primary master and any future updates are automatically replicated to the non-primary nodes.

Note: After you enable this replication option, you can no longer update the IBM Security Verify Access runtime settings on the non-primary nodes of the cluster.

Managing Distributed Session Cache

In a clustered appliance environment, session information is stored in the Distributed Session Cache. To work with these sessions, use the Distributed Session Cache management page.

About this task

The Distributed Session Cache feature replaces the Session Management Server. The Session Management Server (SMS) is not supported on IBM Security Verify Access for Web Version 8 and later.

Procedure

1. From the top menu, select the menu for your activation level.
 - **Web > Manage > Distributed Session Cache**
 - **AAC > Global Settings > Distributed Session Cache**
 - **Federation > Global Settings > Distributed Session Cache**

All replica set names and the number of sessions in each replica set are displayed.

2. You can then view the replica set server list and manage sessions in a particular replica set.
 - a) To view a list of the servers that are registered with a replica set, select the replica set and then click **Servers**.
 - b) To manage the sessions in a replica set, select the replica set and then click **Sessions**.

Tip: Typically, the list of sessions contains many entries. You can locate a session or a user faster by using the filter in the upper left corner.

Delete a specific session

- i) Select the session to delete.
- ii) Click **Delete**.
- iii) In the confirmation window, click **Delete Session**.

Delete all sessions for a user

- i) Select any session for that user.

- ii) Click **Delete**.
- iii) In the confirmation window, click **Delete User**.

Chapter 21. Policy management with Web Portal Manager

Web Portal Manager is a graphical management console for managing domains, users, groups, permissions, policies, and other resources in your enterprise. The appliance provides an embedded version of Web Portal Manager.

To access Web Portal Manager from the appliance, go to **Web > Manage > Policy Administration**.

Note: The Web Portal Manager panels might carry a different appearance than the other appliance panels. This behavior is expected. It does not affect the performance of the embedded Web Portal Manager.

For more information about how to use Web Portal Manager, see [Web Portal Manager](#).

Chapter 22. Global settings

Managing dynamic URL configuration files

In the local management interface, go to **Web > Global Settings > URL Mapping**. A list of all dynamic URL (DynURL) configuration files is displayed. You can view individual file details, and create, import, export, update, rename, and delete DynURL files.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **URL Mapping**.
4. Perform any of the following actions:

Viewing details of a DynURL configuration file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a DynURL configuration file:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.
- d. Click **Save**.

Importing a DynURL configuration file:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a DynURL configuration file:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Modifying a DynURL configuration file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Enter the name for the file.
- e. Click **Save**.

Renaming a DynURL configuration file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a DynURL configuration file:

- a. Select the file that you want to delete.
 - b. Click **Delete**.
 - c. Click **Yes** when you are prompted to confirm the deletion.
5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing junction mapping JMT configuration files

In the local management interface, go to **Web > Global Settings > Junction Mapping**. A list of all files is displayed. You can view individual file details, and create, import, export, update, rename, and delete files.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **Junction Mapping**.
4. Perform any of the following actions:

Viewing details of a JMT configuration file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a JMT configuration file:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.
- d. Click **Save**.

Importing a JMT configuration file:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a JMT configuration file:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Modifying a JMT configuration file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming a JMT configuration file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a JMT configuration file:

- a. Select the file that you want to delete.
 - b. Click **Delete**.
 - c. Click **Yes** when you are prompted to confirm the deletion.
5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing client certificate CDAS files

In the local management interface, go to **Web > Global Settings > Client Certificate Mapping**. A list of all client certificate CDAS files is displayed. You can view individual file details, and create, import, export, update, rename, and delete CDAS files.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **Client Certificate Mapping**.
4. Perform any of the following actions:

Viewing details of a client certificate CDAS file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a client certificate CDAS file:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.
- d. Click **Save**.

Importing a client certificate CDAS file:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a client certificate CDAS file:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Modifying a client certificate CDAS file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming a client certificate CDAS file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a client certificate CDAS file:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing user mapping CDAS files

You can use a user mapping CDAS file to map an authenticated user name to a different Security Verify Access user identity.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **User Name Mapping**.
4. Perform any of the following actions:

Viewing details of a user mapping CDAS file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a user mapping CDAS file:

- a. Click **New**.
- b. Enter the name for the file.
- c. Click **Save**.

Importing a user mapping CDAS file:

- a. Click **Manage > Import**.
- b. Click **Browse**.

- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a user mapping CDAS file:

- a. Select the file that you want to export.
- b. Click **Manage > Export**.
- c. Confirm that you want to save the file to your local workstation.

Modifying a user mapping CDAS file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming a user mapping CDAS file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a user mapping CDAS file:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing password strength rule files

You can use a password strength rule file to define the criteria for new passwords to be validated against.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **Password Strength**.
4. Perform any of the following actions:

Viewing details of a password strength rule file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a password strength rule file:

- a. Click **New**.
- b. Enter the name for the file.
- c. Click **Save**.

Importing a password strength rule file:

- a. Click **Manage > Import**.

- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a password strength rule file:

- a. Select the file that you want to export.
- b. Click **Manage > Export**.
- c. Confirm that you want to save the file to your local workstation.

Modifying a password strength rule file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming a password strength rule file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a password strength rule file:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing forms based single sign-on files

In the local management interface, go to **Web > Global Settings > Forms Based Single Sign-On**. A list of all files is displayed. You can view individual file details, and create, import, export, update, rename, and delete files.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **Forms Based Single Sign-On**.
4. Perform any of the following actions:

Viewing details of a forms based single sign-on file:

- a. Select the file to view.
- b. Click **Edit**. The file content is displayed.

Creating a forms based single sign-on file:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.

- d. Click **Save**.

Importing a forms based single sign-on file:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting a forms based single sign-on file:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Modifying a forms based single sign-on file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming a forms based single sign-on file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting a forms based single sign-on file:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing HTTP transformation files

In the local management interface, go to **Web > Global Settings > HTTP Transformation**. A list of all files is displayed. You can create, import, export, update, rename, and delete files.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **HTTP Transformation**.
4. Perform any of the following actions:

Creating an HTTP transformation rule file:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.

- d. Click **Save**.

Importing an HTTP transformation rule file:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting an HTTP transformation rule file:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Modifying an HTTP transformation rule file:

- a. Select the file that you want to modify.
- b. Click **Edit**.
- c. Modify the content of the file.
- d. Click **Save**.

Renaming an HTTP transformation rule file:

- a. Select the file that you want to rename.
- b. Click **Manage > Rename**.
- c. In the **New Resource Name** field, enter the new name for the file.
- d. Click **Save**.

Deleting an HTTP transformation rule file:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing RSA SecurID configuration

In the local management interface, go to **Web > Global Settings > RSA SecurID Configuration**. The status of the RSA server and node is displayed, as well as the option to upload or clear a RSA configuration, clear a node secret, and test a configuration.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **RSA SecurID Configuration**.
4. Perform any of the following actions:

Uploading a new RSA server configuration file

- a. Click **Upload** in the **Server Configuration File** section.
- b. Select the Server Configuration File to be uploaded.

Note: The RSA configuration file to be uploaded to the appliance must be generated by the RSA server.

- c. **Optional:** Select the Server Configuration Options File to be uploaded.

Note: This is a text file named `sdopts.rec`. It contains the **CLIENT_IP** parameter that specifies the IP address that the SecurID authentication method should use. For example, **CLIENT_IP=1.2.3.4**.

- d. Click **Submit**.

Uploading a new RSA server configuration options file

- a. Click **Upload** in the **Server Configuration Options File** section.
- b. Select the Server Configuration Options File to be uploaded.

Note: This is a text file named `sdopts.rec`. It contains the **CLIENT_IP** parameter that specifies the IP address that the SecurID authentication method should use. For example, **CLIENT_IP=1.2.3.4**.

- c. Click **Submit**.

Downloading the RSA server configuration options file

- a. Click **Download** in the **Server Configuration Options File** section.

Removing an RSA server configuration file:

- a. Click **Clear** under the **Server Configuration File** section.
- b. Confirm that you want to clear the configuration.

Removing the RSA server configuration options file

- a. Click **Clear** under the **Server Configuration Options File** section.
- b. Confirm that you want to clear the configuration and click **Clear**.

Testing a configuration

- a. After uploading a server configuration file, click **Test**.
- b. Enter a valid user.
- c. Enter a valid passcode.

Note: You might need to disable two-step authentication on the RSA server to successfully test the configuration, as the test function does not support two-step authentication.

Clearing a node secret

- a. Click **Clear** under the **Node Secret File** section.
- b. Confirm that you want to clear the secret.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing the Redis configuration

In the local management interface, go to **Web > Global Settings > Redis Configuration**. A list of all existing Redis collections is displayed. You can create, edit and delete collections.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web > Global Settings > Redis Configuration**
3. Perform one of the following actions:

- **Creating a new Redis Collection**

- a. Click **New**.
- b. Specify the configuration entries on the **General** tab.
- c. Select the **Server** tab and add the definition for any Redis servers to be included in this collection.
- d. Click **OK** to save the collection information.

Editing an existing Redis Collection

- a. Select the collection which you wish to edit.
- b. Click **Edit**.
- c. Update the configuration entries on the **General** tab.
- d. Select the **Server** tab and update the definition for the Redis servers.
- e. Click **OK** to save the collection information.

Deleting an existing Redis Collection

- a. Select the collection which you wish to delete.
- b. Click **Delete**.
- c. Click **OK** to confirm the deletion of the selected collection.

Fields**On the General tab**

Section	Field	Description
General Properties	Name	The name which is used to describe this collection of Redis servers.
	Matching Host	Any specific hosts (obtained from the Host header of the request to the Web Reverse proxy) for which this collection should be used. Shell-style pattern matching characters, including “*”, “?”, “\” and “[]”, can be used when matching the Host header.

Section	Field	Description
Cross Domain Support	Enable Cross Domain Support	Enable cross-domain support so that a single Redis session can be shared across multiple DNS domains.
	Master Authentication Server URL	The base URL of the master authentication server for this collection of Redis servers. The master authentication server, if specified, will be responsible for the generation of all new sessions for this collection. This field should be of the format: http{s}://<server>{:<port>}
	Session Code Lifetime	The maximum number of seconds that a session code, used when communicating the session information from the master authentication server, will remain valid.
Connection Properties	Maximum Pooled Connections	The maximum number of pooled connections to a Redis server within this collection.
	Idle Timeout	The maximum number of seconds a pooled connection can remain idle before the connection is closed.
	Connection Timeout	The maximum number of seconds to wait for a connection to be established with a server.
	IO Timeout	The maximum number of seconds to wait for a valid response from a Redis server.
	Health Check Interval	The interval (in seconds) between health check requests sent to the Redis server.

On the Server tab, when adding a new server or editing an existing server:

Section	Field	Description
Server	Name	The name which is used to describe this Redis server.
	Host	The server name or IP address of the Redis server.
	Port	The port on which the Redis server is listening for requests.
	Username	The name of the user which is used when authenticating to the Redis server.
	Password	The password which is used to access the Redis server.
SSL	Enable SSL	Enable SSL communication with the Redis server.
	Key File	The name of the key database which is to be used when accessing this server. The key database should contain the CA certificate for the Redis server certificate, and if mutual authentication is in use, any intermediate certificates used to sign the client certificate, and the client key itself.
	Client Certificate Label	The label associated with the client key which is used to perform mutual authentication with the Redis server.
	SNI	The Server Name Indication (SNI) value which is provided when establishing the SSL connection with the Redis server.

4. Deploy the changes as described in [Configuration changes commit process](#).

Chapter 23. Global keys

Managing SSO keys

In the local management interface, go to **Web > Global Settings > SSO Keys**. A list of all keys is displayed. You can create, import, export, and delete keys.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **SSO Keys**.
4. Perform any of the following actions:

Creating an SSO key:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.
- d. Click **Save**.

Importing an SSO key:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting an SSO key:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Deleting an SSO key:

- a. Select the file that you want to delete.
 - b. Click **Delete**.
 - c. Click **Yes** when you are prompted to confirm the deletion.
5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Managing LTPA keys

In the local management interface, go to **Web > Global Settings > LTPA Keys**. A list of all keys is displayed. You can create, import, export, and delete keys.

Before you begin

Ensure that your browser allows pop-up windows to be displayed.

Procedure

1. Log in to the local management interface.
2. Click **Web**.
3. Under **Global Settings**, click **LTPA Keys**.
4. Perform any of the following actions:

Creating an LTPA key:

- a. Click **New**.
- b. Modify the content of the file.
- c. Enter the name for the file.
- d. Click **Save**.

Importing an LTPA key:

- a. Click **Manage > Import**.
- b. Click **Browse**.
- c. Select the file that you want to import.
- d. Click **Import**.

Exporting an LTPA key:

- a. Click **Browse**.
- b. Select the file that you want to export.
- c. Click **Manage > Export**.
- d. Confirm that you want to save the file to your local workstation.

Deleting an LTPA key:

- a. Select the file that you want to delete.
- b. Click **Delete**.
- c. Click **Yes** when you are prompted to confirm the deletion.

5. Deploy the changes as described in [“Configuration changes commit process”](#) on page 36.

Kerberos configuration

You can create, edit, delete, and test the following Kerberos settings from the local management interface.

Setting	Description
libdefault	Contains default values that are used by the Kerberos library.
realms	Contains subsections that are keyed by Kerberos realm names. Each subsection describes realm-specific information, which includes where to find the Kerberos servers for that realm.
domain realms	Contains relations that map domain names and subdomains to Kerberos realm names. These relations are used by programs to determine what realm a host is in, given its fully qualified domain name.

Table 27. Manage Kerberos configuration settings (continued)

Setting	Description
CA paths	Contains the authentication paths that are used with direct (non-hierarchical) cross-realm authentication. Entries in this section are used by the client to determine the intermediate realms that can be used in cross-realm authentication. It is also used by the end-service when it checks the transited field for trusted intermediate realms.
keytab files	Contains the keytab files that are used for Kerberos authentication. The files contain pairs of Kerberos principals and encrypted keys.

Managing the default values used by Kerberos

Use the Defaults tab on the Kerberos Configuration management page in the LMI to manage these settings. These settings are used as default values by the Kerberos library.

About this task

The Defaults tab contains settings for the **libdefault** section of the corresponding Kerberos configuration file. You can create, edit, and delete properties in this section. You can also test authentication with your web server principal name and password.

Procedure

1. From the top menu, select **Web > Global Settings > Kerberos Configuration**.

The current Kerberos configuration is displayed.

2. On the Defaults tab, take actions as needed.

- Create a property
 - a. Click **New**.
 - b. In the **Create New Property** window, select a name from the **Pre-Defined Names** list or enter a name in the **Name** field as the name of the new property.
 - c. Provide the value of the new property in the **Value** field.
 - d. Click **Save**.
- Edit a property
 - a. Select the property to edit from the table.
 - b. Click **Edit**.
 - c. In the **Edit Property** window, modify the value of the property as needed.
 - d. Click **Save**.
- Delete a property
 - a. Select the property to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
- Test authentication with principal and password
 - a. Click **Test**.
 - b. In the **Test Kerberos Authentication** window, enter the name of the user that is created as the web server principal in the **Username** field.
 - c. Enter the password in the **Password** field.

- d. Click **Test**.

Managing realms

Use the Realms tab on the Kerberos Configuration management page in the LMI to manage these settings. These settings describe realm-specific information.

About this task

The Realms tab contains settings for the **realms** section of the corresponding Kerberos configuration file. You can create, edit, and delete realms, configuration subsections, and properties in this section. You can also test authentication with your web server principal name and password.

Procedure

1. From the top menu, select **Web > Global Settings > Kerberos Configuration**.

The current Kerberos configuration is displayed.

2. On the Realms tab, take actions as needed.

- Create a realm
 - a. Click **New > Realm**.
 - b. In the **Create New Realm** window, enter the name of the new realm in the **Realm** field.
 - c. Click **Save**.
- Create a configuration subsection
 - a. Select the realm in which to create the subsection.
 - b. Click **New > Subsection**.
 - c. In the **Create New Subsection** window, select a name from the **Pre-Defined Names** list or enter a name in the **Subsection** field.
 - d. Click **Save**.
- Create a property
 - a. Select the realm or subsection in which to create the property.
 - b. Click **New > Property**.
 - c. In the **Create New Property** window, select a name from the **Pre-Defined Names** list or enter a name in the **Name** field.
 - d. Enter the value of the property in the **Value** field.
 - e. Click **Save**.
- Edit a property
 - a. Select the property to edit.
 - b. Click **Edit**.
 - c. In the **Edit Property** window, modify the value as needed.
 - d. Click **Save**.
- Delete a realm
 - a. Select the realm to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
- Test authentication with principal and password
 - a. Click **Test**.
 - b. In the **Test Kerberos Authentication** window, enter the name of the user that is created as the web server principal in the **Username** field.

- c. Enter the password in the **Password** field.
- d. Click **Test**.

Managing domain realm properties

Use the Domains tab on the Kerberos Configuration management page in the LMI to manage these settings. These settings describe relations that map domain names and subdomains to Kerberos realm names.

About this task

The Domains tab contains settings for the **domain_realm** section of the corresponding Kerberos configuration file. You can create, edit, and delete properties in this section. You can also test authentication with your web server principal name and password.

Procedure

1. From the top menu, select **Web > Global Settings > Kerberos Configuration**.
The current Kerberos configuration is displayed.
2. On the Domains tab, take actions as needed.
 - Create a domain realm property
 - a. Click **New**.
 - b. In the **Create New Translation** window, enter the local DNS address in the **Local DNS Value** field.
 - c. Select a realm from the **Realm** list.
 - d. Click **Save**.
 - Edit a domain realm property
 - a. Select the domain realm property to edit from the table.
 - b. Click **Edit**.
 - c. In the **Edit Property** window, modify the realm as needed.
 - d. Click **Save**.
 - Delete a domain realm property
 - a. Select the domain realm property to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
 - Test authentication with principal and password
 - a. Click **Test**.
 - b. In the **Test Kerberos Authentication** window, enter the name of the user that is created as the web server principal in the **Username** field.
 - c. Enter the password in the **Password** field.
 - d. Click **Test**.

Managing CA paths

Use the CA Paths tab on the Kerberos Configuration management page in the LMI to manage these settings. These settings contain the authentication paths that are used with direct (non-hierarchical) cross-realm authentication.

About this task

The CA Paths tab contains settings for the **capaths** section of the corresponding Kerberos configuration file. You can create, edit, and delete properties and CA paths in this section. You can also test authentication with your web server principal name and password.

Procedure

1. From the top menu, select **Web > Global Settings > Kerberos Configuration**.

The current Kerberos configuration is displayed.

2. On the CA Paths tab, take actions as needed.

- Create a CA path
 - a. Click **New > Client Realm**.
 - b. In the **Create Client Realm** window, enter the realm name in the **Client Realm** field.
 - c. Click **Save**.
- Create a property
 - a. Select the client realm in which to create the property.
 - b. Click **New > Property**.
 - c. In the **Create New Property** window, provide a value for the **Server Realm** and **Intermediate Realm**.
 - d. Click **Save**.
- Edit a property
 - a. Select the property to edit from the table.
 - b. Click **Edit**.
 - c. In the **Edit Property** window, modify the value as needed.
 - d. Click **Save**.
- Delete a CA path
 - a. Select the CA path to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
- Delete a property
 - a. Select the property to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
- Test authentication with principal and password
 - a. Click **Test**.
 - b. In the **Test Kerberos Authentication** window, enter the name of the user that is created as the web server principal in the **Username** field.
 - c. Enter the password in the **Password** field.
 - d. Click **Test**.

Managing keytab files

Use the Keyfiles tab on the Kerberos Configuration management page in the LMI to manage these settings.

About this task

The Keyfiles tab contains settings for the keytab files that are used for Kerberos authentication. You can import, combine, and delete keytab files. You can also test authentication with a Kerberos principal name and keytab file.

Procedure

1. From the top menu, select **Web > Global Settings > Kerberos Configuration**.

The current Kerberos configuration is displayed.

2. On the Keyfiles tab, take actions as needed.

- Import a keytab file
 - a. Click **Import**.
 - b. In the **Import Keytab File** window, click **Browse**.
 - c. Select the keytab file to be imported and then click **Open**.
 - d. Click **Import**.
- Delete a keytab file
 - a. Select the file to delete from the table.
 - b. Click **Delete**.
 - c. In the **Confirm Action** window, click **Yes**.
- Combine keytab files
 - a. Select the keytab files to be combined from the table.
 - b. Click **Combine**.
 - c. In the **Combine Keytab Files** window, enter the name for the combined file in the **New Resource Name** field.
 - d. Click **Save**.
- Verify authentication with a keytab file
 - a. Select the keytab file to test from the table.
 - b. Click **Test**.
 - c. In the **Test Keytab Authentication** window, provide the value of the Kerberos principal in the **Username** field.
 - d. Click **Test**.

Chapter 24. Trace data

You can use the local management interface (LMI) to control tracing.

Trace data is intended primarily for use by IBM Software Support. Trace data might be requested as part of diagnosing a reported problem. However, experienced product administrators can use trace data to diagnose and correct problems in an IBM Security Verify Access environment. For more information about trace event logging, see [Troubleshooting](#).

Note: Use trace with caution. It is intended as a tool to use under the direction of IBM Software Support. Messages from tracing are sometimes cryptic, are not translated, and can severely degrade system performance.

Modifying the tracing settings for a component

To modify the trace level, flush interval, rollover size, maximum rollover files, and whether rollover files are automatically compressed for a component, use the Reverse Proxy management page or the Authorization Server management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy** if you want to manage tracing for a reverse proxy instance. Or select **Web > Manage > Authorization Server** if you want to manage tracing for an authorization server instance.
2. Select the instance of interest.
3. For reverse proxy, select **Troubleshooting > Tracing**. For authorization server, select **Manage > Tracing**.
4. Select the component to be modified and then click **Edit**.
5. Modify the trace level, flush interval, rollover size, maximum rollover files, and whether rollover files are automatically compressed.
By default, the **Compress** option is set to **No**. To save disk space, set the **Compress** option to **Yes** so that all rollover files are automatically compressed.
6. Click **Save**.

Managing the trace files for a component

To manage the trace files and rollover files for a component, use the Reverse Proxy management page or the Authorization Server management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy** if you want to manage tracing for a reverse proxy instance. Or select **Web > Manage > Authorization Server** if you want to manage tracing for an authorization server instance.
2. Select the instance of interest.
3. For reverse proxy, select **Troubleshooting > Tracing**. For authorization, select **Manage > Tracing**.
4. Select a component and then click **Files** to view a list of all its trace and rollover files.
The file name, file size, and last modified time of each file is displayed.

View or export a trace file or rollover file

- a. Select the file of interest.
- b. Click **View**. The content of the trace files is then displayed. To view a particular number of lines of trace, provide a value in the **Number of lines to view** field and then click **Reload**. Optionally,

you can provide a value in the **Starting from line** field to define the start of the lines. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

- c. Click **Export** if you want to export the file.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

- d. Confirm the save operation when the browser prompts you to save the file.

Delete a trace file or rollover file

- a. Select the file or files of interest.

Note: Only a file that is not in use can be deleted.

- b. Click **Delete**.
- c. Click **Yes** to confirm the operation.

Export a trace file or rollover file

- a. Select the file of interest.
- b. Click **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

- c. Confirm the save operation when the browser prompts you to save the file.

Delete all trace files and rollover files that are not in use

- a. Click **Manage > Delete All**.
- b. Click **Yes** to confirm the operation.

Editing the tracing configuration file for the runtime environment

To edit the tracing configuration file for the runtime environment, use the Runtime Component management page.

Procedure

1. From the top menu, select **Web > Manage > Runtime Component**.
2. Select **Manage > Configuration Files > Tracing Configuration Files**.

The tracing configuration file contents are displayed.

Note: The **Tracing Configuration File** menu item is available only when a local policy server is configured. When a remote policy server is configured, this menu item is disabled. In that case, you must directly edit the file on the machine where the policy server is installed.

3. In the **Tracing Configuration File Editor** window, modify the file.
4. Click **Save** to save the changes. Or click **Cancel** if you do not want to save the changes.

Note: For the changes to take effect, the changes must be deployed and the runtime environment must be restarted.

Updating a tracing configuration file

To update a tracing configuration file with the local management interface, use the Reverse Proxy Instances management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Manage > Configuration > Edit Tracing Configuration File**.

The tracing configuration file contents are displayed.

4. Modify the file.
5. Click **Save** to save the changes. Or click **Close** if you do not want to save the changes.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Chapter 25. Logging

You can use the local management interface (LMI) to manage the reverse proxy log files.

Note: The web reverse proxy log files record the events and activities of the web reverse proxies during the daily operation of the appliance. There are two ways to reduce the disk space that is used by these files.

1. Configure the web reverse proxy to send the log information to a remote server. For more information about the remote logging options, see [“Configuring Web Application Firewall ”](#) on page 216.
2. Clear the unused log files regularly. For details, see [“Managing reverse proxy log files”](#) on page 271. Alternatively, use the command-line interface to back up the log files to a USB device, and to purge all log files that were rolled over. For details, see [“Archiving and deleting reverse proxy log files with the command-line interface”](#) on page 226.

Listing the names of all log files and file sizes

To list the names of all log files and file size with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. *Optional:* If instance-specific log files are of interest, select the instance.
3. Select **Troubleshooting > Logging**.

If an instance is selected, details of all common log files and instance-specific log files are displayed. If no instance is selected, only details of the common log files are displayed.

You can use the filter bar under **Name** to filter entries that meet specific conditions. Click **Clear filter** to return to the full list.

Viewing a snippet of or export a log file

To view a snippet of a log file or export a log file with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. *Optional:* If instance-specific log files are of interest, select the instance.
3. Select **Troubleshooting > Logging**.
4. Select the log file that you want to view.
5. Click **View**.

The content of the log file is displayed. By default, the last 100 lines of a log file is displayed if the file is longer than 100 lines. You can define the number of lines to display by entering the number in the **Number of lines to view** field and then click **Reload**. Optionally, you can provide a value in the **Starting from line** field to define the start of the lines. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

6. Click **Export** to download the log file.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

You can also export a file by selecting it and then clicking **Manage > Export**.

Clearing a log file

To clear a log file and turn its size to 0 with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. *Optional:* If instance-specific log files are of interest, select the instance.
3. Select **Troubleshooting > Logging**.
4. Select the log file or files that you want to clear.
5. Click **Clear**.
6. On the Confirm Action confirmation page, click **Yes**.

Managing transaction logging components and data files

To manage transaction logging components and data files with the local management interface, use the Reverse Proxy management page.

Procedure

1. From the top menu, select **Web > Manage > Reverse Proxy**.
2. Select the instance of interest.
3. Select **Troubleshooting > Transaction Logging**.

All transaction logging components and their status, total file size, and rollover size are displayed.

- **Enable or disable a transaction logging component**
 - a. Select the transaction logging component of interest.
 - b. Click **Edit**.
 - c. Select or clear the **Enabled** check box to enable or disable the transaction logging component.
 - d. Optionally, define the rollover size by providing a value in the **Rollover Size** field. If no value is provided, the default rollover size is used.
 - e. Optionally, define the maximum number of rollover files by providing a value in the **Maximum Rollover Files** field. If no value is provided, no rollover files will be deleted.
 - f. Optionally, set the **Compress** option to **Yes** so that all rollover files are automatically compressed to save disk space. By default, the **Compress** option is set to **No**.
 - g. Click **Save** to save your changes.
- **Rollover the data file of a transaction logging component**
 - a. Select the transaction logging component of interest.
 - b. Click **Manage > Rollover**.
 - c. Click **Yes** to confirm the operation.
- **Manage transaction logging data files**
 - a. Select the transaction logging component of interest.
 - b. Click **Files**.
 - **Export a transaction logging data file**

i) Select the transaction logging data file of interest.

ii) Click **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

iii) Confirm whether to open or save the exported file in the browser window.

– **Delete a transaction logging data file**

Note: Only transaction logging data files that are not in use can be deleted.

i) Select the transaction logging data file or files of interest.

ii) Click **Delete**.

iii) Click **Yes** to confirm the operation.

– **Delete all unused transaction logging data files**

i) Click **Manage > Delete All**.

ii) Click **Yes** to confirm the operation.

Managing reverse proxy log files

Use the Manage Reverse Proxy Log Files management page to work with reverse proxy log files.

Procedure

1. From the top menu, select **Monitor > Logs > Manage Reverse Proxy Log Files**.

Details of all common log files are displayed under **Log Files for Selected Instance**.

You can use the filter bar under **Name** to filter entries that meet specific conditions. Click **Clear filter** to return to the full list.

2. *Optional:* If instance-specific log files are of interest, select the instance from the list under **Reverse Proxy Instances**.

Details of all common log files and instance-specific log files are displayed under **Log Files for Selected Instance**.

3. Work with the reverse proxy log files.

• **View the content of a reverse proxy log file**

a. Select the log file that you want to view.

b. Click **View**. The content of the log file is displayed. By default, the last 100 lines of a log file are displayed if the file is longer than 100 lines. You can define the number of lines to display by entering the number in the **Number of lines to view** field and then click **Reload**. Optionally, you can provide a value in the **Starting from line** field to define the start of the lines. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

c. *Optional:* Click **Export** to download the log file.

Note: You must configure the software to block pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

• **Export a reverse proxy log file**

a. Select the log file that you want to export.

b. Click **Manage > Export**.

Note: You must configure the software to block pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

c. Confirm the save operation in the browser window to export the file to a local location.

- **Clear a reverse proxy log file**

a. Select the log file or files that you want to clear.

b. Click **Clear**.

c. On the Confirm Action confirmation page, click **Yes**.

Managing authorization server log files

To work with authorization server log files, use the **Manage Authorization Server Log Files** management page.

Procedure

1. From the top menu, select **Web > Manage > Authorization Server**.
2. Select the instance of interest.
3. Select **Manage > Logging**.
4. Work with the authorization server log files as needed.

View the content of an authorization server log file

a. Select the log file that you want to view.

b. Click **View**. The content of the log file is displayed. By default, the last 100 lines of a log file are displayed if the file is longer than 100 lines. You can define the number of lines to display by entering the number in the **Number of lines to view** field and then click **Reload**. Optionally, you can provide a value in the **Starting from line** field to define which line in the log file to start viewing from. If the **Starting from line** field is set, then the **Number of lines to view** field determines how many lines to view forward from the starting line. If the **Starting from line** field is not set, then the **Number of lines to view** field determines how many lines to view from the end of the log file.

Note: The maximum size that can be returned is 214800000 lines. If a size greater than that is specified, then the maximum (214800000 lines) is returned.

c. *Optional:* Click **Export** to download the log file.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

Clear an authorization server log file

a. Select the log file that you want to clear.

b. Click **Clear**.

c. On the Confirm Action confirmation page, click **Yes**. A system notification is displayed to indicate that the log file is successfully cleared. The original log file with empty content remains in the log list. Any rollover log files (for example, *xxx.log.1* and *xxx.log.2*) are deleted.

Export an authorization server log file

a. Select the log file that you want to export.

b. Click **Manage > Export**.

Note: You must configure the software that blocks pop-up windows in your browser to allow pop-up windows for the appliance before files can be exported.

c. Confirm the save operation in the browser window to export the file to a local location.

Chapter 26. Front-end load balancer

The appliance provides front-end load balancing function to automatically assign client requests to the appropriate reverse proxy server based on the scheduling specified algorithm.

In an IBM Security Verify Access environment, you can have many services. Each *service* has a virtual IP address and a port. Every service is available on one or more real servers. Each *server* is defined by IP address and a port. The front-end load balancer maps incoming service requests to real servers.

A front-end load balancer is a server that uses a virtual IP address to accept requests from a client. It determines which reverse proxy server is most suitable to handle the request and forwards it to the appropriate reverse proxy server.

Incoming requests from the same client are forwarded to the same server. That is, the front-end load balancer provides *stickiness* or *persistence* for existing sessions. The load balancer uses a scheduling algorithm to forward requests from clients that are not already assigned to a back-end server.

In a typical setup, there are two front-end load balancer servers and multiple reverse proxy servers. Configuring two front end load balancers in the environment provides high availability for the front-end load balancing service.

A heartbeat is transmitted between the two front-end load balancers so that the state of each front-end load balancer is known. The load balancer that is actively receiving and processing requests is known as the *active* load balancer. The other load balancer is known as the *passive* load balancer.

When available, the primary front-end load balancer acts as the active load balancer. It is assigned the virtual IP address for the load balancing service and awaits incoming client requests.

If the primary front-end load balancer becomes unavailable, the backup load balancer can no longer detect heartbeats. In this situation, the backup load balancer assumes the virtual IP address and starts accepting requests from clients. That is, the backup load balancer becomes the active load balancer until the primary load balancer is restored.

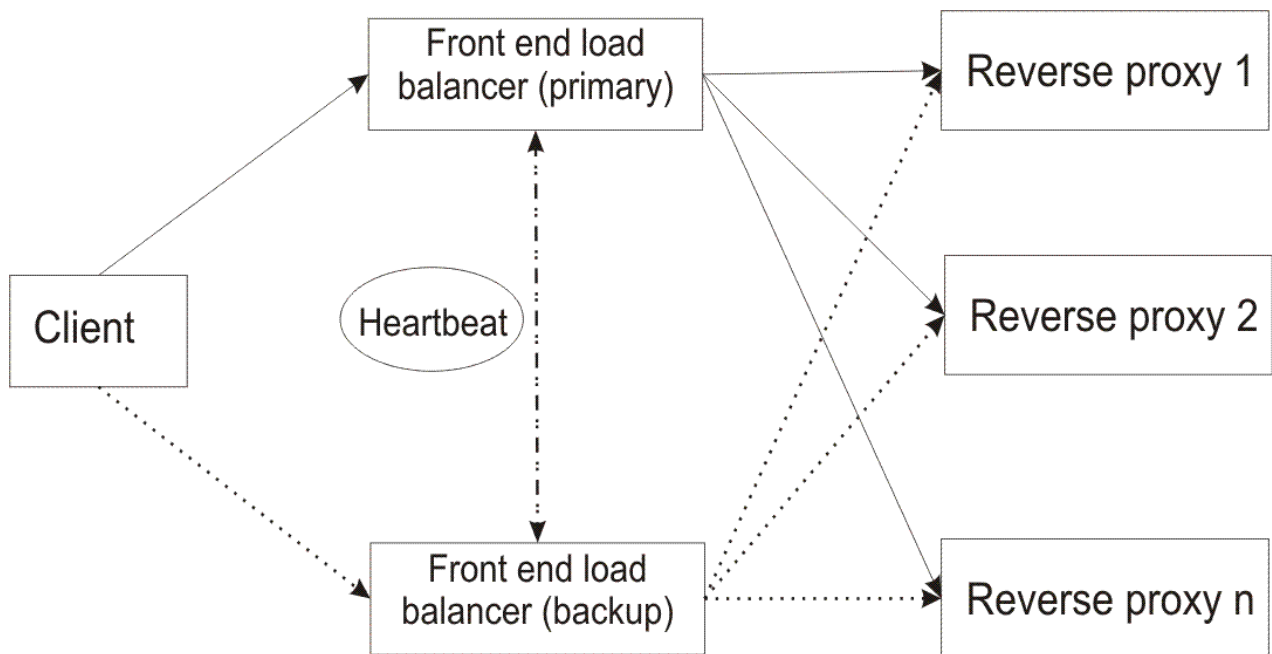


Figure 6. Front-end load balancer

Note: You can have only two front-end load balancers in your environment.

It is possible to configure the reverse proxy functionality on an appliance that is also acting as a front-end load balancer. However, this configuration might have a negative impact on the performance of the front-

end load balancer. If you decide to use such setting, you must take the resources that are used by the reverse proxy into consideration.

You must make sure that the front-end load balancer still has enough resources to perform routing effectively.

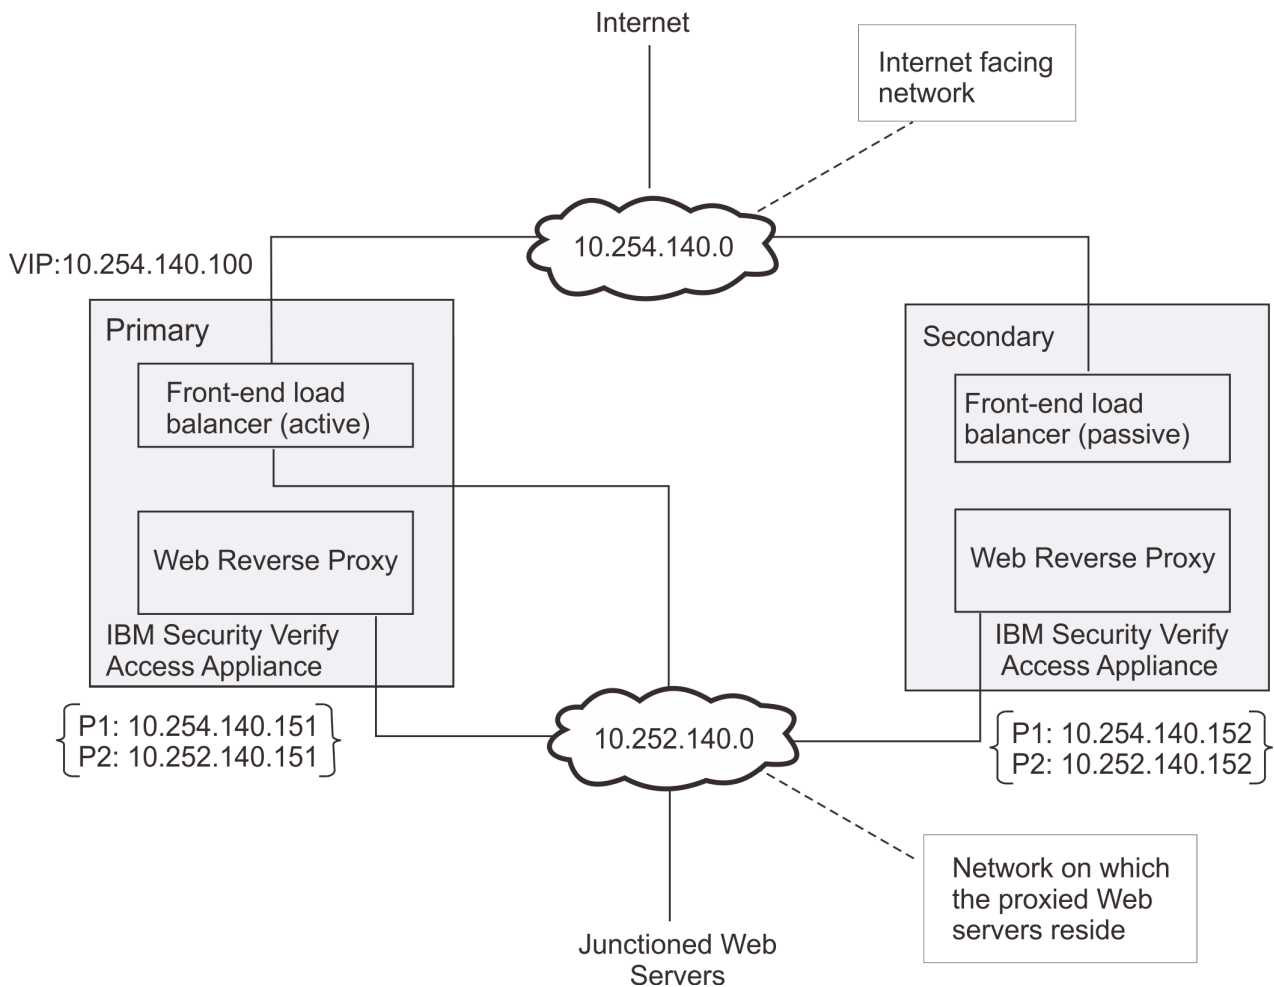


Figure 7. Example high availability environment

You can configure a highly available web reverse proxy environment with as few as two appliances, as shown in Figure 7 on page 274. The active load balancer is on the primary appliance. This load balancer assumes the virtual IP address for the load balancing service. Client requests are received from the Internet-facing network, 10.254.140.0. The load balancer distributes these requests between the web reverse proxy servers, which are on the 10.254.140.0 network.

Scheduling

The front-end load balancing function of the appliance supports several types of scheduling.

In your environment, you might have some servers that are more powerful than others. You can configure the front-end load balancer to respect the relative performance of each server by setting a **weight** value for each server. You can assign weights between 1 and 256, with 256 indicating the most powerful server.

For more information about how to configure the **weight** of each server and select the scheduling algorithm, see “Configuring front-end load balancer” on page 277.

The following scheduling types are supported:

lc

Least connection. The server with the lowest number of connections receives the request. This algorithm is dynamic so you can update the weight ratios in real time.

rr

Round robin. Requests are rotated between the servers. This algorithm is dynamic and uses the weight parameter that is assigned to each server.

srr

Static round robin. Each server is used in turn according to the defined weight for the server. This algorithm is static so you cannot dynamically change the weight ratio for a server.

sh

Source hashing. A hash of the source IP is divided by the total weight of the running servers to determine which server receives the request. This algorithm inherently sends requests from the same IP address to the same server provided that the available servers remains unchanged.

Load balancing layer

Security Verify Access supports load balancing at layer 4 or layer 7 of the Open Systems Interconnection (OSI) network model.

For each service, you can configure either of the following load balancing layers:

TCP Layer (Layer 4)

At this layer, the load balancer can use the TCP header information to determine how to process the request.

Application Layer (Layer 7)

At this layer, the load balancer can recognize application requests (for example, HTTP requests) and process these requests accordingly.

Note: The appliance load balancer does not support HTTP/2 at Layer 7.

Layer 7 offers the following extra features when compared to layer 4 load balancing:

- Ability to use an HTTP cookie to provide *stickiness*. For more information, see [“Persistence” on page 275](#).
- Ability to use and manipulate the headers in HTTP requests and responses. For more information, see [“Benefits of layer 7 load balancing” on page 276](#).

If you do not require these features, use layer 4 load balancing. Layer 4 load balancing is the most efficient type of load balancing. Layer 7 load balancers incur extra processing costs as they need to complete the following extra tasks:

- SSL termination.
- HTTP packet inspection.
- HTTP header manipulation (as required).

For more information about configuring the load balancing layer, see [“Configuring front-end load balancer” on page 277](#).

Persistence

Session persistence, also known as *stickiness*, is a mechanism that ensures a client is connected to the same reverse proxy server during a session.

Layer 4 load balancers can extract the client IP address from the TCP header to maintain persistence. Layer 7 load balancers can use an HTTP cookie to provide stickiness. Subsequent requests from a particular client are routed through the same processing path and use the same WebSEAL session.

Network termination

The front-end load balancer that is provided in Security Verify Access is a network terminating load balancer.

Clients send requests directly to the virtual IP address of the front-end load balancer. The front-end load balancer processes each request.

The load balancer terminates the network connection of the request from the client. It then creates a new network connection to forward the load-balanced request to the appropriate backend server.

The Web Reverse Proxy server receives the request from the front-end load balancer and processes it. The Web Reverse Proxy server sends its response back to the front-end load balancer. The load balancer acts as a proxy and sends the information back to the original client.

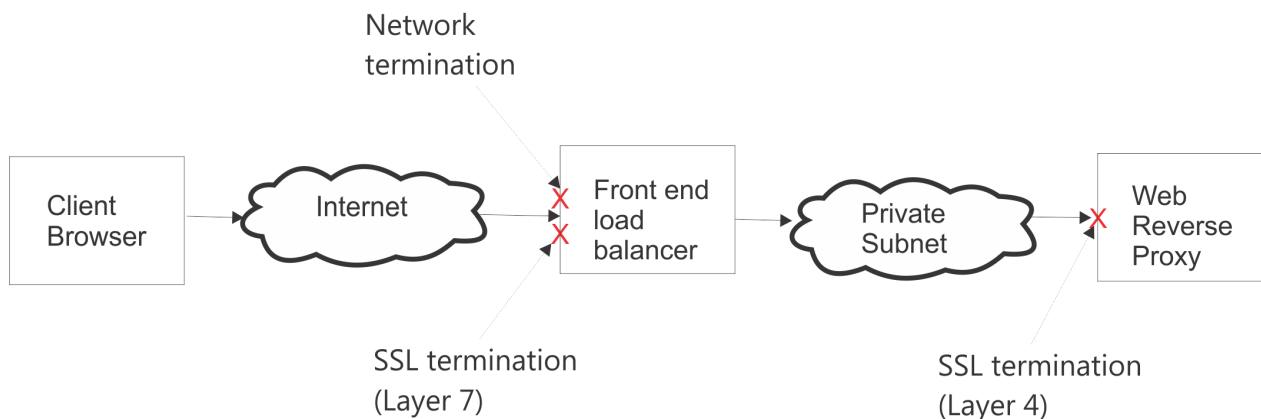


Figure 8. Network termination

The point of SSL termination depends on the load balancing layer. In a layer 4 configuration, WebSEAL is responsible for the SSL termination. In a layer 7 configuration, SSL is terminated by the load balancer.

Benefits of layer 7 load balancing

The main benefit of layer 7 load balancing is the ability to use and manipulate the HTTP headers in requests and responses.

When a layer 7 load balancer processes a request from a particular client for the first time, it adds a load balancer cookie to the HTTP request. The front-end load balancer checks for this load balancer cookie in each subsequent request to provide persistence, or *stickiness*. When you configure Security Verify Access version 7.0 to use layer 7 load balancing, you must specify the name of this cookie for each service.

If you use a layer 7 load balancer, you have access to extra attributes that you can use to manipulate the HTTP requests and responses. For example, you can use the replace attributes, such as **reqrep**, to rewrite URLs or domain names in "Host" headers.

The available attributes for header manipulation are as follows:

reqadd

Adds a header to the end of the HTTP request.

reqdel

Headers that match a specified regular expression are deleted from the request.

reqrep

(Case-sensitive) Search the HTTP request line for a specified regular expression and replace any instances with a specified string.

requirep

(Case-insensitive) Search the HTTP request line for a specified regular expression and replace any instances with a specified string.

rspadd

Adds a header to the end of the HTTP response.

rspdel

Headers that match a specified regular expression are deleted from the response.

rsprep

(Case-sensitive) Search the HTTP response line for a specified regular expression and replace any instances with a specified string.

rspirep

(Case-insensitive) Search the HTTP response line for a specified regular expression and replace any instances with a specified string.

The available attributes to assist with HTTP header based balancing are as follows:

balance

`hdr(<name>)` Overrides the standard scheduler to enable balancing on the specified HTTP request header.

There are also generic attributes to configure connection properties for the front-end load balancer. For example, you can set values for the connection timeout, number of retries, and number of concurrent connections. For a complete list of the available attributes, see [“Configuring front-end load balancer” on page 277](#).

Configuring front-end load balancer

To configure the front end load balancer with the local management interface, use the Front End Load Balancer management page.

Procedure

1. From the top menu, select **System > Network Settings > Front End Load Balancer**.
2. On the **General** tab page:
 - a) Select **Enabled** if you want to enable this front-end load balancer.
 - b) Select **Debug** if you want more debug messages to be sent to the security log.
 - c) Select **Enable SSL** if you plan to enable SSL communication for any Layer-7 services.
 - d) In the **SSL Key File** list, select the key file that contains the certificates to be used in the Layer-7 SSL communication.
3. Optional: On the **Advanced Tuning** tab page, modify global level parameters to fine tune the configuration.
 - a) Click **Add**.
 - b) In the Add New Parameter window, select the desired parameter from the **Name** list.
 - c) Enter a value for the selected parameter in the **Value** field.
 - d) Click **Save**.
4. On the **Servers** tab page, you can work with virtual servers and real servers. Each virtual server corresponds to an interface (virtual IP address and port) that is load balanced. Each real server corresponds to a load balanced server.

- **Add a virtual server**

- a. Click **New**.
- b. On the **Add Virtual Server** page, define settings of the virtual server to be added.

On the **General** tab page:

Field	Description
Enabled	Specifies whether the new virtual server is active.
Name	Name of the virtual server, which is used to uniquely identify this server. Note: The syntax for the virtual server name must be treated as if it were a server host name. It must not contain any space characters.
Virtual Address	Specifies the IP address that connects this virtual server to the public network.
Port	Specifies the port on which this virtual server listens.
Mask	Specifies the network mask to be applied to the IP address for the virtual server.
Interface	Specifies the appliance interface on which the new virtual server connects to the public network.
Layer 4 or Layer 7	The load balancing layer for the server. Layer 4 indicates TCP level load balancing. Layer 7 indicates application level load balancing.
Cookie used in Layer 7	The name of the cookie to be used in Layer 7 load balancing. Note: This field is available only when Layer 7 load balancing has been selected.
Layer 7 SSL Enabled	Whether SSL is used to terminate the connection. Note: This field is available only when Layer 7 load balancing has been selected.
Layer 7 SSL Certificate Label	The label of the certificate to be used when terminating the connection. Note: This field is available only when Layer 7 load balancing has been selected.

On the **Scheduler** tab page:

Field	Description
Scheduler	<p>Specifies the scheduling algorithm for distributing jobs to the real servers. Available choices are:</p> <p>lc Least connection. The server with the lowest number of connections receives the request. This algorithm is dynamic so you can update the weight ratios in real time.</p> <p>rr Round robin. Requests are rotated between the servers. This algorithm is dynamic and uses the weight parameter that is assigned to each server.</p> <p>srr Static round robin. Each server is used in turn according to the defined weight for the server. This algorithm is static so you cannot dynamically change the weight ratio for a server.</p> <p>sh Source hashing. A hash of the source IP is divided by the total weight of the running servers to determine which server receives the request. This algorithm inherently sends requests from the same IP address to the same server provided that the available servers remains unchanged.</p> <p>For Layer 4 operations, only a scheduler setting of sh (source hash) specifies to use all CPUs available on the appliance. If other scheduler settings are used for Layer 4 operation, then the load balancer process operates that particular virtual server by using one CPU. This behavior might impact performance of the front end load balancer for the virtual server, particularly if the back-end servers are using SSL.</p> <p>For Layer 7 operations, all CPUs available are always used regardless of the scheduler setting.</p>
Health Check Interval	Number of seconds between health check messages that are sent to the real servers.
Rise	The number of successful health checks before a server is considered active.
Fall	The number of unsuccessful health checks before a server is considered inactive.

Optional: On the **Advanced Tuning** tab page, add, edit, or delete any service level advanced configuration parameters as needed. See [“Front-end load balancer advanced tuning parameters” on page 281](#) for the available parameters. See [“Benefits of layer 7 load balancing” on page 276](#) for descriptions of the advanced tuning attributes available.

- c. Click **Save**.
- **Delete a virtual server**
 - a. Select the virtual server to delete from the list.
 - b. Click **Delete**.
 - c. On the confirmation page, click **Yes**.
- **Edit a virtual server**
 - a. Select the virtual server to edit from the list.
 - b. Click **Edit**.
 - c. On the **Edit Virtual Server** page, modify the settings as needed.

d. Click **Save**.

- **Manage real servers**

a. From the list of virtual servers, select the virtual server to associate the real servers with.

b. Click **Real Servers**. The **Real Servers** page is displayed.

– To add a real server:

i) Click **New**.

ii) On the **Add Real Server** page that pops up, define settings for the server to be added.

Field	Description
Enabled	Specifies whether the new real server is active.
Address	Specifies the IP address for the real server.
Weight	Specifies an integer that represents this processing capacity of the server relative to that of other real servers. For example, a server assigned 2000 has twice the capacity of a server assigned 1000. The weighted scheduling algorithms adjust this number dynamically based on workload.
SSL Enabled	Specifies whether to use an SSL connection between the load balancer and the back-end server.
SSL Certificate Label	Specifies the SSL certificate label.

iii) Click **Save**.

– To delete a real server:

i) Select the real server to delete from the list.

ii) Click **Delete**.

iii) On the confirmation page, click **Yes**.

– To edit a real server:

i) Select the real server to edit from the list.

ii) Click **Edit**.

iii) On the **Edit Real Server** page, modify the settings as needed.

iv) Click **Save**.

c. Click **Close** to return to the Front End Load Balancer main page.

5. On the **High Availability** tab page, you can define the settings that enable high availability of the front-end load balancer function. For example, configure a second front-end load balancer as either a primary or a back-up load balancer for the environment.

a) Select the **Enable High Availability** check box to enable this feature.

b) Select **Primary** or **Backup** to designate this system as the primary or backup front-end load balancer.

c) For the **Local Interface - Primary** field, select the local IP address of the front-end load balancer.

d) For the **Remote Address - Backup** field, specify the IP address that is used by this system to communicate with the other front-end load balancer. This field is required if a backup load balancer is in use.

e) For the **Remote Port** field, specify the port to be used for high availability communication.

f) In the **Health Check Interval** field, specify in seconds the interval of the heartbeat messages that are sent between the primary and backup front-end load balancers.

g) In the **Health Check Timeout** field, specify in seconds the time to wait before the system declares a non-responsive router unavailable and initiating failover.

6. On the **Logging** tab page, configure the local or remote logging options.
 - If you select **Log to local**, no additional configuration is required on this page.
 - If you select **Log to remote**, you must provide values for **Syslog facility**, **Remote syslog server address**, and **Remote syslog server port**.
7. On the **Error Pages** tab page, customize the error pages (200, 400, 403, 408, 500, 502, 503, and 504) that are returned by the software. These error pages are returned when the layer-7 load balancing function encounters a problem.
 - To edit an existing error page:
 - a. Select the error page to customize.
 - b. Click **Edit**.
 - c. In the Edit File window, modify the error page as needed.
 - d. Click **Save**.
 - To import a new page to replace an existing error page:
 - a. Select the error page to be replaced.
 - b. Click **Import**.
 - c. In the Import Error Page window, click **Browse**.
 - d. Select the new page.
 - e. Click **Save**.
 - To export an error page:
 - a. Select the error page to export.
 - b. Click **Export**.
 - c. Specify the destination location to export the file to.
 - d. Click **Export** to confirm the operation.
8. Click **Save** to save all changes that are made on the **Front End Load Balancer** management page.

Note: For the changes to take effect, they must be deployed as described in [“Configuration changes commit process”](#) on page 36.

Front-end load balancer advanced tuning parameters

Use these parameters to tune the front-end load balancer configuration.

For detailed descriptions of these parameters, see the HAProxy documentation at <http://www.haproxy.org/download/1.8/doc/configuration.txt>.

Note:

- When you configure an option that does not contain any parameters (for example, **disable-on-404**), the contents of the **Value** field in the UI will be ignored.
- If you experience difficulty when configuring the front-end load balancer, examine the front-end load balancer log file to help with troubleshooting.

Chapter 27. dscadmin command

Use the **dscadmin** command option from the command-line interface (CLI) to administer the distributed session cache.

To access this command, log onto the command-line interface (either by logging onto the appliance console, or performing an ssh into the machine), and then enter the **isam** menu, followed by the **dscadmin** sub-menu.

The **dscadmin** command supports the following operations:

- replica set show *replica_set_name*
- replica set list
- session terminate all_sessions *user_id replica_set_name*
- session terminate session *session-id replica-set-name*
- session list *pattern maximum_return replica_set_name*
- exit
- quit

replica set show

Lists all session management replicas in the specified replica set. A *replica* is a client that has registered with the distributed session cache.

Syntax

```
replica set show replica_set_name
```

Options

replica_set_name

Specifies the name of the replica set.

Examples

The following example returns details about the `ibm.com` replica set:

```
dscadmin> replica set show ibm.com
```

replica set list

Lists all session management replica sets in the domain.

Syntax

```
replica set list
```

Options

N/A

Examples

The following example lists all the replica sets:

```
dscadmin> replica set list
```

session terminate all_sessions

Terminates all user sessions for a specific user within the specified replica set.

Syntax

```
session terminate all_sessions user_id replica-set-name
```

Options

user_id

Specifies the name of the user. An example of user name is `sec_master`. Pattern matching can be used when specifying the user name.

replica_set_name

Specifies the name of the replica set.

Examples

The following example terminates all sessions for the `sec_master` user in the `ibm.com` replica set:

```
dscadmin> session terminate all_sessions sec_master ibm.com
```

The following example terminates all sessions whose user names start with `sec_m` in the `ibm.com` replica set:

```
dscadmin> session terminate all_sessions sec_m* ibm.com
```

session terminate session

Terminates a user session using a session ID within the specified replica set.

Syntax

```
session terminate session session-id replica-set-name
```

Options

session-id

Specifies the ID of a user session.

replica_set_name

Specifies the name of the replica set.

Examples

The following example terminates session 678 in the `ibm.com` replica set:

```
dscadmin> session terminate session 678 ibm.com
```


session list

Lists all session management sessions within the specified replica set.

Syntax

```
session list pattern maximum_return replica_set_name
```

Options

pattern

Specifies the pattern for returning user names. The pattern can include a combination of wild card and string constant characters. The pattern is not case-sensitive. For example, you can specify **luca** or **LUCA** as the pattern to find all users that contain the substring *luca* in the user name.

Note: Only the asterisk (*) character can be used as wild card.

maximum_return

Specifies the maximum number of sessions to return. When there are more matches than designated by this option, the output contains the number of matches.

replica_set_name

Specifies the name of the replica set.

Examples

The following example (entered as one line) lists the user sessions in the *ibm.com* replica set for users that contains the string *ons* and limits the number of matches to 100:

```
dscadmin> session list *ons* 100 ibm.com
```

exit or quit

Use either the **exit** command or the **quit** command to exit from the **dscadmin** utility interactive command-line mode.

Syntax

```
exit  
quit
```

Options

N/A

Examples

The following example displays how to exit the **dscadmin** utility:

```
dscadmin> exit
```

The following example displays how to quit the **dscadmin** utility:

```
dscadmin> quit
```


Chapter 28. API Access Control

Use the local management interface to manage the API Access Control configuration

Overview of the API Access Control

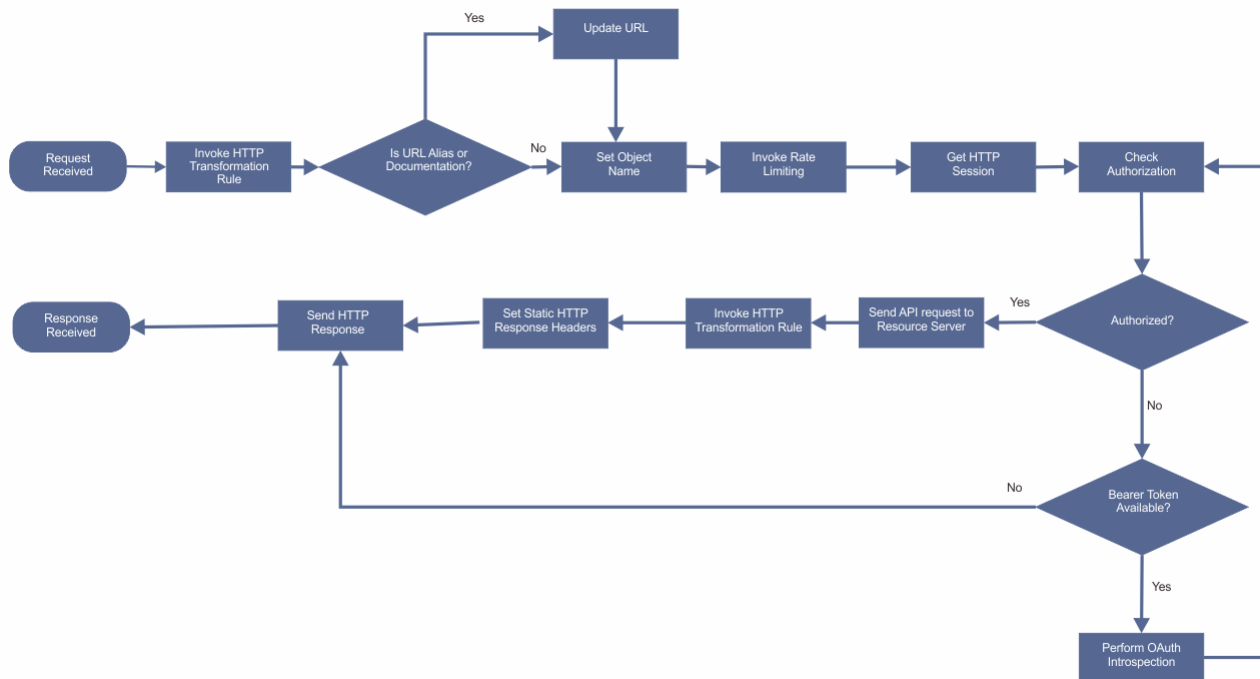
IBM Security Verify Access provides many capabilities which can be configured independently to protect a RESTful API.

The following are capabilities that are included:

- Junctions
- Access Control Lists (ACLs)
- Protected Object Policy (POP)
- HTTP Transformation Rules
- Rate Limiting Policy
- Static Response Headers
- OAuth Validation

The API Access Control component provides a simple way for these various capabilities to be configured in order to protect a RESTful API.

The following diagram shows a high level overview of the flow of a request when the API Access Control has been configured.



Components

The API Access Control component is broken into three separate sub components.

Resource Servers

Each of these defines a single server that hosts the API that is being protected.


Resources

Each of these defines a single API.

Resource Access Policies

These are the authorisation policies for accessing the resource.

The following diagram shows the hierarchy of the sub components.



Authorization

The API Access Control component introduces a new objectspace named `/WebSEAL_API`.

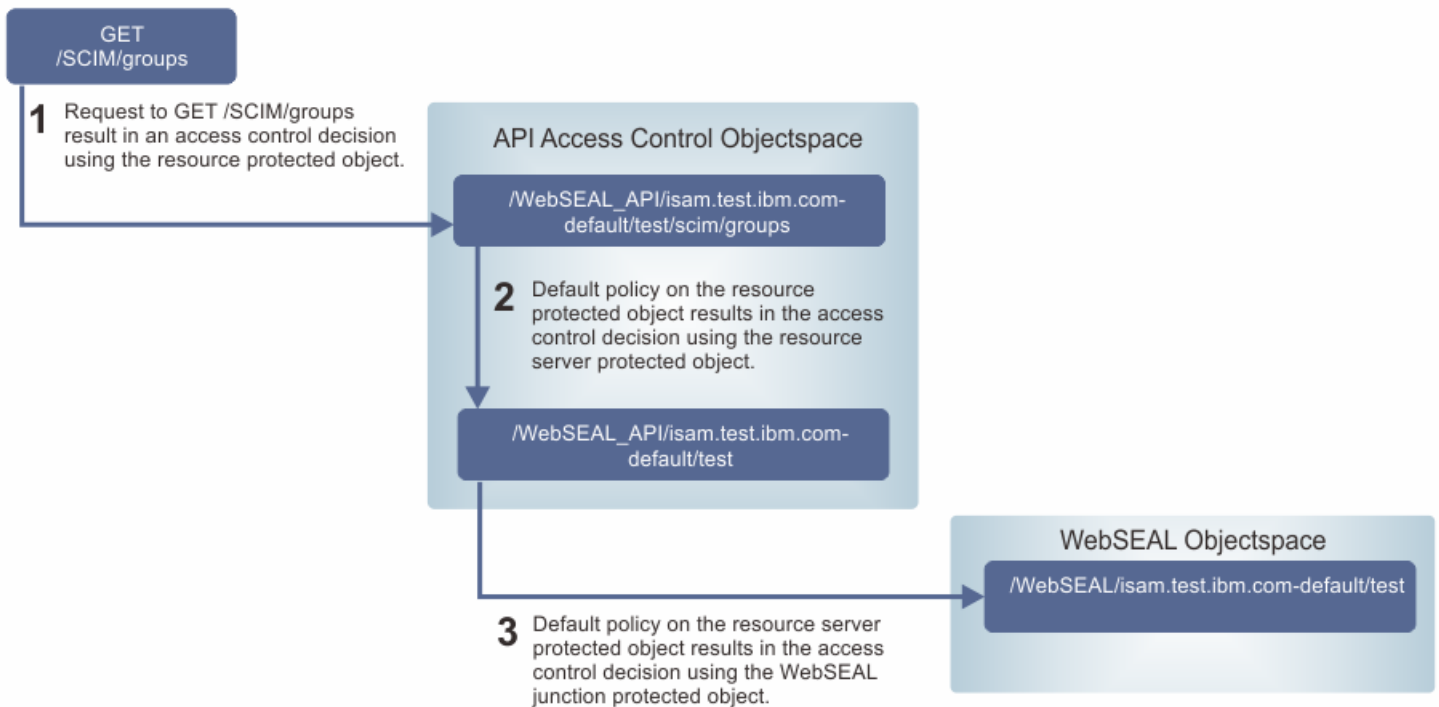
The new objectspace is used to contain the protected objects that represent both resource servers and resources (described below). The hierarchy of the protected objects in this objectspace resembles the WebSEAL objectspace:

```
/WebSEAL_API/<hostname>-<instance_name>/<resource_server>/<resource>
```

The objectspace is managed by the API Access Control component and any requests made to a protected API will use it in the authorization decision. The protected object used in the access control decision depends upon the configuration of the API Access Control policy.

1. If a non default policy is applied to the API Access Control resource, the ACL and/or POP that is attached to the resource protected object is used for the access control decision.
2. If a non default policy is applied to the API Access Control resource server, the ACL and/or POP that is attached to the resource server protected object is used for the access control decision.
3. If the default policy is specified for the resource and resource server the ACL and/or POP that is attached to the WebSEAL junction protected object is used for the access control decision.

The following diagram shows the flow of the protected objects that are used in an access control decision when default policy is applied to the resource and resource server.



Resource Servers

A resource server is the definition of the server that provides access to the RESTful API that is being protected.

Each resource server corresponds to a Reverse Proxy junction. The API Access Control component provides an extended configuration mechanism that allows more than just the standard junction management. It provides a way to:

1. Create a basic standard junction with only the minimal required configuration or an advanced junction creation that provides all of the standard junction create options.
2. Set the authentication options for incoming requests to the resource server.
3. Set the default authentication policy for all requests to the resource server.
4. Set static response headers that are set on every response to requests to the resource server.

The base level operations that occur internally when you creating a new resource server include the following:

- The junction is created.
- The Reverse Proxy administrative pages root directories are updated to include new directories that are specific to the new resource server. Pages that are specific to the resource server can be placed in these locations.
- The Reverse Proxy configuration file is updated to include any OAuth introspection configuration as well as static response header definitions.
- A new IBM Security Verify Access protected object is created to represent the new resource server.
- A new Access Control List (ACL) might be created to represent the authentication policy. This ACL is then attached to the new protected object.

Note: Due to the different mix of IBM Security Verify Access operations that are performed, there are some tasks which are completed immediately (for example, pdadmin tasks) and some tasks that are not completed until the next commit operation is executed. This means that there might be orphaned protected objects, ACLs and POPs if the creation is followed by a rollback of the pending changes. Therefore, take precaution if the administrator chooses to roll back the pending changes instead of deploying them.

Take the following actions, if a cleanup or audit of orphaned API Access Control artifacts is required:

- Objects under the /WebSEAL_API object space must be reviewed and unnecessary objects removed;
- ACLs that have a suffix of "_resource_access_control_policy" must be reviewed and unnecessary ACLs must be removed.

Resources

A resource is the definition of the RESTful API that is being protected.

The simplistic view of an API is the combination of a HTTP method and path. For example, GET /scim/groups.

The API Access Control component provides an extended configuration mechanism which allows more than just the method and path to be specified. It provides a way to:

1. Create the protected object that represents the method and path.
2. Specify any URL aliases that should map to the specified path. This includes the ability to use a wildcard character (*). These are applied through the use of a HTTP transformation rule.
3. Specify a rate limiting policy that is applied to any requests to the resource.
4. Set the authentication policy for all requests to the resource.
5. Set static response headers that are set on every response to requests to the resource. These are applied through the use of a HTTP transformation rule.
6. Set a documentation file that can be returned when requests to the resource specify a certain content type in the accept header. This is applied through the use of a HTTP transformation rule. See ["Resource Documentation" on page 292.](#)

The base level operations that occur internally when creating a new resource include:

- The protected object is created.
- A new ACL may be created representing the authentication policy. This ACL is then attached to the new protected object.
- The Reverse Proxy configuration file is updated with:
 - A mapping of the resource description to the protected object name.
 - Any new HTTP transformation rule specifications.
 - Activation of the specified rate limiting policy.
- The rate limiting policy is updated to include the new method and path as a resource.
- A new request HTTP Transformation rule is created. This rule:
 - Has a name formatted as : "Request_" + instance name + resource server path + method + path. For example, Request_default_testGET_test.
 - Contains a rule to set the protected object used for the authorisation check to the required API Access Control protected object name.
 - Map any URL aliases back to the main resource path.
 - Forward the request onto the documentation file if the accept header matches the specified documentation content type.
- A new response HTTP transformation rule is created. This rule:
 - Has a name formatted as : "Response_" + instance name + resource server path + method + path. For example, Response_default_testGET_test.
 - Sets any static response headers.

The following is an example API Access Control Request Transformation Rule (Request_default_scimGET_scim_groups):

```
<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0"
```

```

xmlns:external="http://xsltfunctions.isam.ibm.com"
xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

<!--This stylesheet is used to set static response headers in the HTTP Response for the GET /test resource.-->

<!--Firstly, strip any space element-->
<xsl:strip-space elements="*" />

<xsl:template match="/">
  <HTTPRequestChange>
    <xsl:apply-templates
      select="//HTTPRequest/Headers/Header[@name = 'accept']"/>
    <xsl:apply-templates select="//HTTPRequest/RequestLine"/>

    <!--Set the ACL bits that will be used for authorisation for this resource.-->
    <AclBits>T</AclBits>

    <!--Set the object name which will be used in the authorization decision.-->
    <ObjectName>/WebSEAL_API/isam.test.ibm.com-default/scim/GET/scim/groups</ObjectName>
  </HTTPRequestChange>
</xsl:template>
<!--Handle any URL aliasing-->
<xsl:template match="//HTTPRequest/RequestLine">
  <xsl:choose>
    <xsl:when test="external:matches(URI, '^/scim/groups/.*)">
      <URI>
        <xsl:value-of
          select="external:replace(URI, '/scim/groups/(.*)', '/scim/groups$1)"/>
      </URI>
    </xsl:when>
  </xsl:choose>
</xsl:template>
<!--Manage any requests for documentation using the accept header.-->
<xsl:template match="//HTTPRequest/Headers/Header[@name = 'accept']">
  <xsl:if test="node()='application/swagger.json'">
    <URI>/apiac/scim/groups.json</URI>
  </xsl:if>
</xsl:template>
</xsl:stylesheet>

```

The following is an example API Access Control Response Transformation Rule (Response_default_scimGET_scim_groups):

```

<?xml version="1.0" encoding="UTF-8"?>
<xsl:stylesheet version="1.0"
  xmlns:external="http://xsltfunctions.isam.ibm.com"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform">

<!--This stylesheet is used to set static response headers in the HTTP Response for the GET /test resource.-->
<!--Firstly, strip any space element-->

<xsl:strip-space elements="*" />
<xsl:template match="/">
  <HTTPResponseChange>
    <Header action="add" name="Strict-Transport-Security">true</Header>
  </HTTPResponseChange>
</xsl:template>
</xsl:stylesheet>

```

Note: Due to the different mix of IBM Security Verify Access operations that are performed, there are some tasks which are completed immediately (for example, pdadmin tasks) and some tasks that are not completed until the next commit operation is executed. This means that there might be orphaned protected objects, ACLs and POPs if the creation is followed by a rollback of the pending changes. Therefore, take precaution if the administrator chooses to roll back the pending changes instead of deploying them.

Take the following actions, if a cleanup or audit of orphaned API Access Control artifacts is required:

- Objects under the /WebSEAL_API object space must be reviewed and unnecessary objects removed;
- ACLs that have a suffix of "_resource_access_control_policy" must be reviewed and unnecessary ACLs must be removed.

Deletion of a resource that has a rate limiting policy applied results in the method and path being removed from the rate limiting policy file but the policy is not removed from the Reverse Proxy configuration file. If the rate limiting policy is no longer required, it can be manually removed from the '[rate-limiting]' stanza of the Reverse Proxy configuration file.

Resource Documentation

IBM Security Verify Access provides the capability for a RESTful API to be extended such that a request can be made to retrieve documentation for the API.

This is achieved through the use of a HTTP Transformation rule. The rule will look for a certain accept header value and if it exists the documentation file which was provided during the configuration of the API will be returned.

For example an API Access Control resource GET /scim/groups is configured with the documentation content type set to application/swagger.json and the documentation file set as scimgroups.json. The following CURL request returns the contents of the documentation file scimgroups.json:

```
curl -H "Accept: application/swagger.json" -k --user
easuser:password https://isam.test.com/scim/groups
```

API Access Control allows the API documentation HTTP Transformation rule to be created and enabled when you are creating or modifying a resource. Set the documentation content type and file as part of the resource data.

The documentation files are served by the Web Reverse proxy local junction and are stored within the /apiac directory.

The new location can be managed by using the new API Access Control management capabilities or as part of the existing Reverse Proxy management root. The directory is created when an API resource is first created and cannot be deleted. Files can be added directly to the /apiac directory or a complex sub directory structure can be created.

Resource Access Policies

A resource access policy is the authorisation policy that can be applied to resource servers and/or resources.

The underlying resource access policy is represented by an ACL and/or a POP.

There are 5 different types of resource access policy that can be applied to a resource server or resource:

Default Verify Access Policy

- No ACL or POP is attached to the protected object. Instead the parent object is used to check access.
- As stated in “Authorization” on page 288, protected objects exist in a new objectspace / WebSEAL_API. If both the resource and resource server have **Default Verify Access Policy** enabled, the authorization check reverts to use the standard /WebSEAL objectspace.

No Access Permitted

An ACL is created to not allow access to anyone.

Unauthenticated Access Allowed

An ACL is created to allow access to unauthenticated users.

Any Authenticated

An ACL is created to allow access to any authenticated users.

Custom

A custom ACL and/or POP is attached to the protected object.

API Access Control has the capability of defining a custom policy that can be used for resource server or resource access control. The custom policy can define:

- Groups that are allowed access.
- Credential attributes that are checked for access control.

Any groups that are specified are added to an ACL where:

- The format of the ACL name is “resource_access_control_policy_” + custom policy name.
- If the user is a member of any of the groups access is allowed.

The credential attributes are specified as part of an attribute check like:

```
<attribute_name> = <attribute_value> { OR <attribute_name> = <attribute_value>}
```

The attribute check is added to a POP where:

- The format of the POP name is “resource_access_control_policy_” + custom policy name.
- The POP has the attribute “eas-trigger” set to “trigger_attr_eas”. This enables the attribute check.
- The POP has one or more requires attributes set. These are used for credential attribute checks.
 - Multiple credential attributes can be specified in a single POP attribute. In this case only one of the specified attributes need to match. (OR condition).
 - Multiple POP attributes can be specified and in this case every POP attribute must evaluate to true (AND condition).
- If all of the attribute checks are true access is allowed

For example:

A policy named "test" is created with the following groups:

- admin
- privileged

and the following attribute checks:

- scope=usr:write OR scope=usr:admin
- AuthenticationLevel=2

This results in the creation of the ACL and the POP:

```
resource_access_policy_test

ACL Name: resource_access_control_policy_test
Description: API Access Control Policy ACL- Do not modify
Entries:
  Group admin T
  Group privileged T
  User sec_master TcmdbsvaBR1
```

```
resource_access_policy_test

Protected object policy:
resource_access_control_policy_test
eas-trigger
  trigger_attr_eas
requires
  test='scope=usr:write OR scope=usr:admin'
requires
  test='AuthenticationLevel=2'
```

When the policy is used in an authorisation check for access to be granted:

1. The user must be in at least one of the admin or privileged groups.
2. The credential attributes must include:
 - AuthenticationLevel=2; AND
 - scope set as either usr:write OR usr:admin

For a description of the POP requires attribute, see [Using credential attributes in authorization decisions](#).

Cross-Origin Resource Sharing (CORS) Policies

A CORS policy specifies the settings that can be applied to resources to allow Cross-Origin Resource Sharing.

CORS is a mechanism that uses additional HTTP header to inform a browser to allow a web application running at one origin (domain) have permission to access selected resources from a server at a different origin.

Each IBM Security Verify Access API Access Control resource can be configured with a CORS policy where each policy defines:

- Whether or not the reverse proxy should perform the pre-flight check. For example, OPTIONS check.
- The origins that are permitted to make requests to this resource.
- Whether or not to set the Access-Control-Allow-Credentials header.
- The headers that are added to a pre-flight check response.
- The methods that are allowed in requests to this resource.
- The maximum time a client should cache the pre-flight check response.
- The headers that a client should expose.

Note: Attaching a CORS policy to an API Access Control resource results in a new stanza being added to the reverse proxy configuration file. This new stanza is marked with a comment stating that the contents are machine generated and should not be modified manually. This is to ensure that the API Access Control management component is not effected by any manual changes. Any changes that are made by an administrator are overwritten by CORS policy updates. An example of the new stanza:

```
[cors-policy:apiac_policyA]
# *****
*****#
THIS STANZA IS AUTO GENERATED. PLEASE DO NOT UPDATE AS IT MAY CAUSE PROBLEMS WITH THE API
ACCESS CONTROL COMPONENT
# *****
*****
handle-pre-flight = false
max-age = 0
allow-credentials = false
allow-origin = http://test.com
request-match = GET /application/endpointA HTTP/*
```

To create a new CORS policy, see [“Creating a CORS policy”](#) on page 307.

To add a CORS policy to an API Access Control resource, see [“Create a new Resource”](#) on page 303 and [“Modify an Existing Resource”](#) on page 304

For more information on how the Reverse Proxy handles CORS processing, see [\[cors-policu:<policy-name>\]](#) stanza and [Cross-Origin Resource Sharing \(CORS\) support](#).

Configuration Auditing

When you are managing API Access Control resource servers, resources, or resource access policies there are a number of IBM Security Verify Access operations that are performed internally.

A log file is used to audit the type of operation and who it was performed by. This file is located in the application log files directory `/isam_runtime/policy_server/api_access_control.log`.

The following is an example of a log entry. It shows:

- The date and time the operation was performed.
- The name of the LMI user that performed the operation (admin).
- As this is a pdadmin operation the name of the IBM Security Verify Access user that ran the command (sec_master).

- The actual command that was run.

```
Apr 01, 2019 9:40:48 PM ApiAccessControl User:admin
pdadmin> : sec_master -> server task default create -t ssl -h 192.168.42.131 -p 443 -x /test
```

Storing the IBM Security Verify Access operations for managing Access Control Policies

The IBM Security Verify Access administrator credentials are required for all of the API Access Control configuration tasks. Rather than providing them for each task individually they need to instead be stored once.

Procedure

1. In the appliance top menu, **Web > API Access Control > Policies or Resources**

Note:

- If the credentials are not already set, a dialog box is shown prompting for the username and password. Proceed to [Step 3](#).
 - If the credentials are already set and need to be updated, proceed to [Step 2](#).
2. Click the **Set Credentials** button on the toolbar to launch the dialog box.
 3. Enter the username and password.
 4. Optional: Enter the domain.

Note: If not specified, the domain is set to **Default**.

5. Click **Save**.

Note: The credentials are stored in memory for the life of the API Access Control servlet. The password is obfuscated.

Auditing the Verify Access operations that are performed when managing API Access Control components

To view the Verify Access operations that are performed when you are managing API Access Control components, use one of the following methods:

Procedure

1. Use the API Access Control page:
 - a) From the appliance dashboard, select **Web > API Access Control > Resources, Policies or CORS Policies**.
 - b) Click the **Logging** button in the toolbar to view the list of internal Verify Access operations that have been run.
2. Use the **Application Log Files** page:
 - a) From the appliance dashboard, select **Monitor > Application Log Files**
 - b) In the tree open **isam_runtime > policy_server**.
 - c) Select the `api_access_control.log` file.
 - d) Click **View** to see the list of internal Verify Access operations that have been run.

Manage Access Control Policies

In the local management interface, go to **Web > API Access Control > Policies**. A list of all current policies is displayed.

Create a new Access Control Policy

To create a new API Access Control policy with the local management interface, use the API Access Control policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Policies**
2. Click **Add**.
A dialog box is displayed prompting for policy details.
3. Enter the name for the new policy in the **Name** field.
4. Add any new groups to the policy criteria by clicking the **Add** button in the groups toolbar.
A new dialog box is shown.
 - a. Select the new group from the drop-down list of all available groups.
 - b. Click **Save**.
5. Click the **Remove** button to remove any groups by selecting the group in the groups list.
6. Add any new attributes to the policy criteria by clicking the **Add** button in the attributes toolbar.
A new dialog is shown.
 - a. Enter the attribute name in the **Name** field.
 - b. Enter the attribute value in the **Value** field
 - c. Multiple attributes can be combined together as a single attribute definition. They are combined using an OR condition.
 - i) To add more attributes click the **Add** button.
 - ii) To delete an attribute click the **Delete** button.
 - d. Once the attribute definition is correct, click the **Save** button.
7. Click the **Remove** button to remove any attributes by selecting the attribute in the attribute grids.
8. Click **Save**.

Note:

- For the policy to allow access the user must be in at least one of the groups.
- For the policy to allow access all of the separate attribute definitions must be met.
- To view a list of all of the internal Verify Access operations that are run to create a new policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Modify an existing Access Control Policy

To modify an existing API Access Control policy with the local management interface, use the API Access Control policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Policies**
2. Select the policy to edit from the list of displayed policies and click **Edit**.
3. Enter the name for the new policy in the **Name** field.
4. Add any new groups to the policy criteria by clicking the **Add** button in the groups toolbar.
A new dialog box is shown.
 - a. Select the new group from the drop-down list of all available groups.
 - b. Click **Save**.
5. Click the **Remove** button to remove any groups by selecting the group in the groups list.

6. Add any new attributes to the policy criteria by clicking the **Add** button in the attributes toolbar. A new dialog is shown.
 - a. Enter the attribute name in the **Name** field.
 - b. Enter the attribute value in the **Value** field
 - c. Multiple attributes can be combined together as a single attribute definition. They are combined using an OR condition.
 - i) To add more attributes click the **Add** button.
 - ii) To delete an attribute click the **Delete** button.
 - d. Once the attribute definition is correct, click the **Save** button.
7. Click the **Remove** button to remove any attributes by selecting the attribute in the attribute grids.
8. Click **Save**.

Note:

- For the policy to allow access the user must be in at least one of the groups.
- For the policy to allow access all of the separate attribute definitions must be met.
- To view a list of all of the internal Verify Access operations that are run to create a new policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Delete one or more existing Access Control Policies

To delete one or more existing API Access Control policies with the local management interface, use the API Access Control policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Policies**
2. To delete all existing policies, follow the steps:
 - a) Select all policies or select no policies.
 - b) Click the **Delete** button.
3. To delete a selection of policies, follow the steps:
 - a) Select the policies to delete.
 - b) Click the **Delete** button.

Note:

- Deletion of a policy deletes the ACL and POP that is associated with the policy and removes all resource attachments.
- To view a list of all of the internal Verify Access operations that are run to create a new policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Manage Resource Servers and Resources

In the local management interface, go to **Web > API Access Control > Resources**. A list of all current reverse proxy instances is displayed in a tree view. To view the existing resource servers in each instance click the icon next to the instance. To view the existing resources for a resource server click the icon next to the resource server.

Resource Servers

This section describes the procedures to manage resource servers.

Create a new resource server

To create a new Resource Server with the local management interface, use the API Access Control resources page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).

3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.

4. Click **Add**.

A dialog box is displayed prompting for the resource server details.

5. In the API Host tab enter the details for the API host server.

- a) Enter the server details by using the basic or advanced data.

For Basic Data, click the Basic Radio button

- i) Enter the path prefix in the Path Prefix field. This becomes the standard junction point to the server.
- ii) Enter a user friendly description for this resource server in the **Description** field.
- iii) Enter the hostname or IP address for this resource server in the **Server** field.
- iv) Enter the port for this resource server in the **Port** field.
- v) If SSL is required check the SSL checkbox, otherwise leave it unchecked.
- vi) After the server and port are entered optionally, click the **Load Key** button to load the CA certificate from the server into the reverse proxy keyfile.
- vii) For server authentication data click **None** if not required.
- viii) For server authentication data, click **Client Certificate** for certificate authentication and select the certificate from the drop-down list.
- ix) For server authentication data, click **Basic Authentication** for basic authentication and enter the username and password

For Advanced Data, click the Advanced Radio button

- i) Select the **Standard junction** radio button to create a new standard junction.
- ii) Select the **Virtual junction** radio button to create a new virtual junction.
- iii) Click the **Create** button. This changes the dialog box to allow the advanced junction data to be entered. See [“Creating virtual junctions” on page 231](#) or [“Creating standard junctions” on page 232](#).
- iv) Once all the values are entered, click **OK** to return to the previous dialog box.

Note: At this stage the advanced junction is not yet created. Values can be changed by clicking **Create** on the API Host tab again.

6. In the **Authentication** tab enter the details for how the OAuth token is validated.

- a) To use the existing reverse proxy configuration select the **Current Reverse Proxy Authentication** radio button.
- b) To use an external OAuth introspection endpoint select the **OAuth Introspection** radio button and enter the details.
 - i) Enter the URL for the introspection endpoint in the **Introspection URL** field.

- ii) After the URL is entered optionally, click the **Load Key** button to load the CA certificate from the endpoint into the reverse proxy keyfile.
 - iii) Choose the method by which the authentication data is presented to the introspection endpoint by selecting either Basic Authentication or POST parameter from the drop-down list.
 - iv) If the authentication data is client ID and/or client secret, click the **Client Credentials** radio button and enter the Client Id and/or Client secret.
 - v) If the authentication data is a client ID header name, select the HTTP Header and enter the Header Name.
 - vi) If the mapped identity must correspond to an existing Verify Access identity, select the **OAuth Identity must correspond to a known Verify Access identity** radio button. If the mapped identity is not required to correspond to an existing Verify Access identity, select **OAuth Identity does not need to correspond to a known Verify Access identity** radio button.
 - vii) To add a new Introspection attribute definition, click the **Add** button in the Introspection Response Attributes toolbar.
 - a) Choose whether this definition is to include or not include this attribute in the response.
 - b) Enter the Attribute name.
 - c) Click **OK**.
 - viii) Click **Delete** to remove an Introspection attribute definition.
 - ix) Click **Move Up** to move an attribute definition up in the ordered list.
7. In the policy tab select the policy that is to be attached to this resource server.
- a) Use the parent policy and select the default **Verify Access Policy** radio button. Do not attach any policy directly.
 - b) Click the **No Access Permitted (disabled)** radio button to not allow access.
 - c) Click the **Unauthenticated Access Allowed** radio button to allow unauthenticated access.
 - d) Click the **Any Authenticated** radio button to allow any authenticated access.
 - e) Click the **Custom** radio button to use a custom Access Control Policy. Select the custom policy name from the drop-down list.
8. In the **Response** tab, set any static response headers to **Create**.
- a) Click **Add** to add a new response header.
 - i) In the dialog box, select the Header Name or enter a new value in the Header Name field.
 - ii) Enter the header value.
 - iii) Click **Save**.
 - b) Click **Delete** to delete a response header from the header list.
9. In the new **Identity** tab, set the JWT configuration:
- a) Check the **Enable JWT** check-box to enable JWT generation.
 - b) Specify the HTTP header name for the generated JWT in the Header Name field.
 - c) Select the certificate that is used to sign the generated JWT from the **Certificate** dropdown. This dropdown is populated with the available personal certificates from either:
 - i) The keystore configured in the `jct-cert-keyfile` entry of the junction stanza in the reverse proxy configuration file.
 - ii) If the `jct-cert-keyfile` entry is not configured, the keystore configured in the `webseal-cert-keyfile` entry of the `ssl` stanza in the reverse proxy configuration file.
 - d) Set the list of claims to add to the generated JWT by using the toolbar for the claims table.
 - i) Click the **Add** button to create a new claim.
 - a) Click the **Literal claim** radio button if the claim is a literal text value.

- b) Click the **Credential attribute claim** radio button if the claim value is retrieved from a credential attribute.
- c) Enter the value for a literal claim or the attribute name for a credential attribute claim. The attribute name can include wildcard characters “*” or “?” if a pattern of attributes is to be included in the generated JWT.
- d) Enter the name for the claim in Claim Name field. This field is optional when the claim is a credential attribute claim.

Note: If the claim is a credential attribute claim and the attribute name includes a wildcard this field is not valid. Instead the claim name for each matched attribute is set as the name of the matched attribute.

If the claim is a credential attribute claim and this field is not set, the claim name is set to the attribute name.

- e) Click the **Save** button to add the new claim to the list of claims.
 - ii) Select the claim to edit and click the Edit button to update an existing claim.
 - iii) Select the claim to remove and click the Delete button to remove an existing claim.
10. Once all of the data is set, click **Save**, to create the new resource server.

Note:

- When a new resource server is created the junction specific management and error pages directories are created.
- To view a list of all of the internal Verify Access operations that are run to create a new policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Modify an existing resource server

To modify an existing Resource Server with the local management interface, use the API Access Control resources page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Resources**.
2. Click the icon next to the Reverse Proxy instance that holds the Resource Server.
These are the root level objects in the tree.
3. Select the resource server to modify.
4. Click **Edit**.
5. The API Host tab only allows the host server to be edited in advanced mode. If any of the data needs to be edited, click the **Edit** button.
This changes the dialog box to allow the advanced junction data to be entered. See [“Creating virtual junctions”](#) on page 231 or [“Creating standard junctions”](#) on page 232.
6. In the **Authentication** tab enter the details for how the OAuth token is validated.
 - a) To use the existing reverse proxy configuration select the **Current Reverse Proxy Authentication** radio button.
 - b) To use an external OAuth introspection endpoint select the **OAuth Introspection** radio button and enter the details.
 - i) Enter the URL for the introspection endpoint in the **Introspection URL** field.
 - ii) After the URL is entered optionally, click the **Load Key** button to load the CA certificate from the endpoint into the reverse proxy keyfile.
 - iii) Choose the method by which the authentication data is presented to the introspection endpoint by selecting either Basic Authentication or POST parameter from the drop-down list.

- iv) If the authentication data is client ID and/or client secret, click the **Client Credentials** radio button and enter the Client Id and/or Client secret.
 - v) If the authentication data is a client ID header name, select the HTTP Header and enter the Header Name.
 - vi) If the mapped identity must correspond to an existing Verify Access identity, select the **OAuth Identity must correspond to a known Verify Access identity** radio button. If the mapped identity is not required to correspond to an existing Verify Access identity, select **OAuth Identity does not need to correspond to a known Verify Access identity** radio button.
 - vii) To add a new Introspection attribute definition, click the **Add** button in the Introspection Response Attributes toolbar.
 - a) Choose whether this definition is to include or not include this attribute in the response.
 - b) Enter the Attribute name.
 - c) Click **OK**.
 - viii) Click **Delete** to remove an Introspection attribute definition.
 - ix) Click **Move Up** to move an attribute definition up in the ordered list.
7. In the policy tab select the policy that is to be attached to this resource server.
- a) Use the parent policy and select the default **Verify Access Policy** radio button. Do not attach any policy directly.
 - b) Click the **No Access Permitted (disabled)** radio button to not allow access.
 - c) Click the **Unauthenticated Access Allowed** radio button to allow unauthenticated access.
 - d) Click the **Any Authenticated** radio button to allow any authenticated access.
 - e) Click the **Custom** radio button to use a custom Access Control Policy. Select the custom policy name from the drop-down list.
8. In the **Response** tab, set any static response headers to **Create**.
- a) Click **Add** to add a new response header.
 - i) In the dialog box, select the Header Name or enter a new value in the Header Name field.
 - ii) Enter the header value.
 - iii) Click **Save**.
 - b) Click **Delete** to delete a response header from the header list.
9. In the new **Identity** tab, set the JWT configuration:
- a) Check the **Enable JWT** check-box to enable JWT generation.
 - b) Specify the HTTP header name for the generated JWT in the Header Name field.
 - c) Select the certificate that is used to sign the generated JWT from the **Certificate** dropdown. This dropdown is populated with the available personal certificates from either:
 - i) The keystore configured in the `jct-cert-keyfile` entry of the junction stanza in the reverse proxy configuration file.
 - ii) If the `jct-cert-keyfile` entry is not configured, the keystore configured in the `webseal-cert-keyfile` entry of the `ssl` stanza in the reverse proxy configuration file.
 - d) Set the list of claims to add to the generated JWT by using the toolbar for the claims table.
 - i) Click the **Add** button to create a new claim.
 - a) Click the **Literal claim** radio button if the claim is a literal text value.
 - b) Click the **Credential attribute claim** radio button if the claim value is retrieved from a credential attribute
 - c) Enter the value for a literal claim or the attribute name for a credential attribute claim. The attribute name can include wildcard characters “*” or “?” if a pattern of attributes is to be included in the generated JWT.

- d) Enter the name for the claim in Claim Name field. This field is optional when the claim is a credential attribute claim.

Note: If the claim is a credential attribute claim and the attribute name includes a wildcard this field is not valid. Instead the claim name for each matched attribute is set as the name of the matched attribute.

If the claim is a credential attribute claim and this field is not set, the claim name is set to the attribute name.

- e) Click the **Save** button to add the new claim to the list of claims.
 - ii) Select the claim to edit and click the Edit button to update an existing claim.
 - iii) Select the claim to remove and click the Delete button to remove an existing claim.
10. Once all of the data is set, click **Save**, to update the resource server.

Delete an existing resource server

To delete an existing Resource Server with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, select **Web > API Access Control > Resources**.
2. Click the icon next to the Reverse Proxy instance that holds the Resource Server.
3. Select the resource server to delete.
4. Click the **Delete** button.

Export the configuration for an existing resource server

To export the configuration of an existing Resource Server or servers with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, **Web > API Access Control > Resources**.
2. If all of the resource servers for a reverse proxy are to be exported, perform the following steps:
 - a) Select the reverse proxy instance. These are the root level objects in the tree.
 - b) Click **Manage > Export**.

This exports a zip file that contains the API Access Control configuration for the selected resource.
3. If a single resource server is to be exported, perform the following steps:
 - a) Click the icon next to the Reverse Proxy instance that holds the Resource Server. These are the root level objects in the tree.
 - b) Select the resource server to export.
 - c) Click **Manage > Export**.

This exports a zip file that contains the API Access Control configuration for the selected resource.

Note: The exported configuration does not contain the junction data. It only contains the API Access Control data that is associated with a junction. For example, OAuth authentication, policy, static response headers, and junction specific management and error pages.

Import the configuration of the resource server

To import the configuration of an existing Resource Server or servers with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, select **Web > API Access Control > Resources**.
2. Select the reverse proxy instance. These are the root level objects in the tree.
3. Click **Manage > Import**.
4. In the dialog box, select the ZIP file containing the resource server data to import.
5. Click **Import**.

Note: The import does not create the resource server junction(s). For each resource server in the ZIP file configuration an existing junction of the same name must exist in the target reverse proxy instance.

View the management and error pages for a resource server

To view the junction specific management and error pages with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, **Web > API Access Control > Resources**.
2. Select the reverse proxy instance. These are the root level objects in the tree.
3. Click **Manage Management Root**.
4. In the dialog box, browse to the type of page (error or management) and the locale.
In the local, there is a directory with the same name as the resource server.
5. Manage the files in the directory. See [“Managing administration pages” on page 222](#).

Resources

This section describes the procedures to manage resources.

Create a new Resource

To create a new Resource with the local management interface, use the API Access Control resources page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).
3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.
4. Select the required resource server to create the Resource.
5. Click **Add**.
A dialog box is displayed prompting for the resource server details.
6. In the **Identification** tab, enter the details for the new API resource.
 - a) Enter the user friendly description for this resource in the **Name** field.
 - b) Enter the path for this resource in the **Path** field.

- c) Select the HTTP method for this resource from the **Method** drop-down list.
 - d) Add any path aliases that can be used for this resource to the list of URL aliases.
 - i) Click **Add**. A dialog is displayed prompting for the new alias.
 - ii) Enter the new alias.

Note: A wildcard character (*) might be used if required.
 - iii) Click **Save**.
 - e) Remove any path aliases by selecting the alias or aliases in the list and click the **Delete** button.
7. In the **Policy** tab, select the policy to be attached to this resource
- a) Use the parent policy and select the default Verify Access Policy radio button. Do not attach any policy directly. This will use the policy set in the resource server.
 - b) Click the **No Access Permitted (disabled)** radio button to not allow access.
 - c) Click the **Unauthenticated Access Allowed** radio button to allow unauthenticated access.
 - d) Click the **Any Authenticated** radio button to allow any authenticated access.
 - e) Click the **Custom** radio button to use a custom Access Control Policy.
 - f) Select the custom policy name from the drop-down list.
 - g) Select the rate limiting policy to be attached to this resource by using the **Rate Limiting Policy** dropdown.
 - h) Select the CORS policy to be attached to this resource by using the **CORS Policy** dropdown.
8. In the **Responses** tab, set any static response headers to create.
- a) Click **Add** to add a new response header.
 - i) In the dialog box, select the Header Name or enter a new value in the **Header Name** field.
 - ii) Enter the Header Value.
 - iii) Click **Save**.
 - b) To delete a response header select the header from the list and click **Delete**.
9. In the **Documentation** tab, enter the settings to allow the resource documentation to be returned.
- a) Enter the content type that can be used to retrieve the documentation. For example, `application/swagger.json`. If a request to the resource contains this value in the accept header, the documentation file is returned.
 - b) Select the documentation file from the drop-down list. If it does not exist click the upload button to select and upload the file. These files can be managed by using [“Manage the API documentation root”](#) on page 307.
10. Once all of the data are set, click **Save** to create the new resource.

What to do next

To view a list of the operations that are run to create a new resource, see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Modify an Existing Resource

To modify an existing Resource with the local management interface, use the API Access Control resources page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.

See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295.](#)

3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.
4. Expand the resource server to show the list of existing resources.
5. Select the resource to modify.
6. Click **Edit**.
A dialog box is displayed showing the existing resource details.
7. In the **Identification** tab, enter the details for the new resource.
 - a) Enter the user friendly description for this resource in the **Name** field.
 - b) Add any path aliases that can be used for this resource to the list of URL aliases.
 - i) Click **Add**. A dialog is displayed prompting for the new alias.
 - ii) Enter the new alias.
Note: A wildcard character (*) might be used if required.
 - iii) Click **Save**.
 - c) Remove any path aliases by selecting the alias or aliases in the list and click the **Delete** button.
8. In the **Policy** tab, select the policy to be attached to this resource
 - a) Use the parent policy and select the default Verify Access Policy radio button. Do not attach any policy directly. This will use the policy set in the resource server.
 - b) Click the **No Access Permitted (disabled)** radio button to not allow access.
 - c) Click the **Unauthenticated Access Allowed** radio button to allow unauthenticated access.
 - d) Click the **Any Authenticated** radio button to allow any authenticated access.
 - e) Click the **Custom** radio button to use a custom Access Control Policy.
 - f) Select the custom policy name from the drop-down list.
 - g) Select the rate limiting policy to be attached to this resource by using the **Rate Limiting Policy** dropdown.
 - h) Select the CORS policy to be attached to this resource by using the **CORS Policy** dropdown.
9. In the **Responses** tab, set any static response headers to create.
 - a) Click **Add** to add a new response header.
 - i) In the dialog box, select the Header Name or enter a new value in the **Header Name** field.
 - ii) Enter the Header Value.
 - iii) Click **Save**.
 - b) To delete a response header select the header from the list and click **Delete**.
10. In the **Documentation** tab, enter the settings to allow the resource documentation to be returned.
 - a) Enter the content type header value. If a request to the resource contains this value in the content-type header, the documentation file will be returned.
 - b) Select the documentation file from the drop-down list. If it does not exist click the upload button to select and upload the file. These files can be managed by using [“Manage the API documentation root” on page 307.](#)
11. Once all of the data are set, click **Save** to update the selected resource.

Note:

- The resource path and method can not be updated.
- To view a list of all of the internal IBM Security Verify Access operations that are run to create a new resource see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components” on page 295.](#)

Deleting an existing Resource

To delete an existing Resource Server with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, select **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).
3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.
4. Expand the resource server to show the list of existing resources.
5. Select the resource to modify.
6. Click **Delete**.

What to do next

To view a list of the IBM Security Verify Access operations that are run to create a new resource, see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components” on page 295](#).

Export the configuration for an existing resource

To export the configuration of an existing resource or resources with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).
3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.
4. Perform one of the following options:

If all of the resources for a reverse proxy are to be exported

- a. Select the resource server.
- b. Click **Manage > Export**. This exports a zip file containing the API Access Control configuration for the resource server and all of its resources.

If a single resource is to be exported

- a. Expand the Resource Server to show the existing resources.
- b. Select the resource to export
- c. Click **Manage > Export**. This exports a zip file containing the API Access Control configuration for the selected resource.

Import the configuration for a resource

To import the configuration of an existing resource or resources with the local management interface, use the API Access Control resources page.

Procedure

1. In the appliance top menu, **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).
3. Expand the Reverse Proxy instance to show the list of available resource servers.
These are the 2nd level objects in the tree.
4. Select the target resource.
5. Click **Manage > Import**.

What to do next

To view a list of the IBM Security Verify Access operations that are run to create a new resource, see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components” on page 295](#).

Manage the API documentation root

To manage the API documentation root with the local management interface, use the API Access Control resources page.

Procedure

1. From the appliance top menu, select **Web > API Access Control > Resources**.
2. The user is prompted to set the user name, password, and domain for the Policy Server if these credentials are not already set.
See [“Storing the IBM Security Verify Access operations for managing Access Control Policies” on page 295](#).
3. Select the reverse proxy instance.
These are the 2nd level objects in the tree.
4. Click **Manage > API Documentation Root**.
5. In the dialog, create, delete, rename, import files, and directories.

Manage Cross-Origin Resource Sharing (CORS) Policies

In the local management interface, go to **Web > API Access Control > CORS Policies**. A list of all current CORS policies is displayed.

Note: To view the details of any of the CORS policies expand the tree to see the current configuration for that policy.

Creating a CORS policy

To create a new API Access Control CORS policy with the local management interface, use the API Access Control CORS policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > CORS Policies**

2. Click **Add**.

A dialog box is displayed prompting for policy details.

3. Enter the name for the new CORS policy in the **Name** field.

4. Enter the Access Control settings in the **Access Control** tab.

a) Select whether or not to set the Access-Control-Allow-Credentials header by using the **Allow Credentials** checkbox.

b) Add any allowed origins to the policy by clicking the **Add** button in the Allowed Origins toolbar. A new dialog box is shown.

i) Enter the new origin in the **Name** field. This value can be either "*" to allow all origins or an individual origin of the form

```
<protocol>://<hostname>:<port>
```

, where the port is optional. Do **not** enter a path in this field.

ii) Click **Save**.

c) Add any exposed headers to the policy by clicking the **Add** button in the Exposed Headers toolbar. A new dialog box is shown.

i) Enter the new header in the **Name** field.

ii) Click **Save**.

d) Remove any allowed origins or exposed headers by selecting the item in their respective lists and click the **Remove** button.

5. Enter the pre-flight check settings in the **Pre-flight Check** tab.

a) Select whether or not to enable the pre-flight check by using the **Handle pre-flight check** checkbox.

Note: If this is not checked the remainder of the fields in this tab are not shown.

b) Enter the maximum age of the pre-flight check response in the **Maximum age** field.

c) Add any allowed methods to the policy by clicking the Add button in the Allowed Methods toolbar. A new dialog box is shown.

i) Enter the new method in the **Name** field.

ii) Click **Save**.

d) Add any allowed headers to the policy by clicking the **Add** button in the Allowed Headers toolbar. A new dialog box is shown.

i) Enter the new header in the **Name** field.

ii) Click **Save**.

e) Remove any allowed methods or allowed headers by selecting the item in their respective lists and click the **Remove** button.

6. Click **Save**.

Note:

- For a policy to be created there must be a unique name and at least one allowed origin specified.
- To view a list of all of the internal Verify Access operations that are run to create a new CORS policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Modifying an existing CORS policy

To modify an existing API Access Control CORS policy with the local management interface, use the API Access Control CORS policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > CORS Policies**
2. Select the policy to modify from the list.
3. Click **Edit**.

A dialog box is displayed showing the current settings.

4. Enter the Access Control settings in the **Access Control** tab.
 - a) Select whether or not to set the Access-Control-Allow-Credentials header by using the **Allow Credentials** checkbox.
 - b) Add any allowed origins to the policy by clicking the **Add** button in the Allowed Origins toolbar. A new dialog box is shown.
 - i) Enter the new origin in the **Name** field. This value can be either "*" to allow all origins or an individual origin of the form.

```
<protocol>://<hostname>:<port>
```

, where the port is optional. Do **not** enter a path in this field.

- ii) Click **Save**.
- c) Add any exposed headers to the policy by clicking the **Add** button in the Exposed Headers toolbar. A new dialog box is shown.
 - i) Enter the new header in the **Name** field.
 - ii) Click **Save**.
 - d) Remove any allowed origins or exposed headers by selecting the item in their respective lists and click the **Remove** button.
5. Enter the pre-flight check settings in the **Pre-flight Check** tab.
 - a) Select whether or not to enable the pre-flight check by using the **Handle pre-flight check** checkbox.

Note: If this is not checked the remainder of the fields in this tab are not shown.
 - b) Enter the maximum age of the pre-flight check response in the **Maximum age** field.
 - c) Add any allowed methods to the policy by clicking the Add button in the Allowed Methods toolbar. A new dialog box is shown.
 - i) Enter the new method in the **Name** field.
 - ii) Click **Save**.
 - d) Add any allowed headers to the policy by clicking the **Add** button in the Allowed Headers toolbar. A new dialog box is shown.
 - i) Enter the new header in the **Name** field.
 - ii) Click **Save**.
 - e) Remove any allowed methods or allowed headers by selecting the item in their respective lists and click the **Remove** button.
 6. Click **Save**.

Note:

- The policy name cannot be modified.
- For a policy to be created there must be a unique name and at least one allowed origin specified.

- When an existing policy is updated all of the API Access Control resources that are using the policy are also updated. This means the reverse proxy configuration files will be updated with the new settings. This overwrites all of the existing CORS policy settings.
- To view a list of all of the internal Verify Access operations that are run to create a new CORS policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Deleting one or all existing CORS policies

To delete one or all existing API Access Control CORS policies with the local management interface, use the API Access Control CORS policies page.

Procedure

1. In the appliance top menu, **Web > API Access Control > CORS Policies**.
2. To delete all existing policies, click **Delete** without selecting a policy from the list. This opens a confirmation dialog.
 - a) Confirm the deletion by clicking the **Delete** button.
 - b) Cancel the deletion by clicking the **Cancel** button.
3. To delete one existing policy, select the policy to from the list and click the **Delete** button. This opens a confirmation dialog.
 - a) Confirm the deletion by clicking the **Delete** button.
 - b) Cancel the deletion by clicking the **Cancel** button.

Note:

- Deleting a CORS policy results in all the API Access Control resource CORS policy references being deleted also. This means the reverse proxy configuration files are updated to remove the CORS policy stanzas for the deleted policies.
- To view a list of all of the internal Verify Access operations that are run to create a new CORS policy see the `api_access_control.log` as described in [“Auditing the Verify Access operations that are performed when managing API Access Control components”](#) on page 295.

Accessibility features for Security Verify Access

Accessibility features assist users who have a disability, such as restricted mobility or limited vision, to use information technology content successfully.

Accessibility features

Security Verify Access includes the following major accessibility features:

Accessibility features
Supports interfaces commonly used by screen readers. This feature applies to applications on Windows operating systems only.
Can be operated by using only the keyboard.
Allows the user to request more time to complete timed responses.
Supports customization of display attributes such as color, contrast, and font size.
Communicates all information independently of color.
Supports interfaces commonly used by screen magnifiers. This feature applies to applications on Windows operating systems only.
Allows the user to access the interfaces without inducing seizures due to photosensitivity.

Security Verify Access uses the latest W3C Standard, WAI-ARIA 1.0 (<http://www.w3.org/TR/wai-aria/>), to ensure compliance to US Section 508 (<http://www.access-board.gov/guidelines-and-standards/communications-and-it/about-the-section-508-standards/section-508-standards>), and Web Content Accessibility Guidelines (WCAG) 2.0 (<http://www.w3.org/TR/WCAG20/>). To take advantage of accessibility features, use the latest release of your screen reader in combination with the latest web browser that is supported by this product.

The Security Verify Access online product documentation in IBM Knowledge Center is enabled for accessibility. The accessibility features of IBM Knowledge Center are described at <https://www.ibm.com/support/knowledgecenter/help?view=kc#accessibility>.

Keyboard navigation

This product uses standard navigation keys.

Interface information

The Security Verify Access user interfaces do not have content that flashes 2 - 55 times per second.

The Security Verify Access web user interfaces and the IBM Knowledge Center rely on cascading style sheets to render content properly and to provide a usable experience. The application provides an equivalent way for low-vision users to use a user's system display settings, including high-contrast mode. You can control font size by using the device or web browser settings.

The Security Verify Access web user interface includes WAI-ARIA navigational landmarks that you can use to quickly navigate to functional areas in the application.

Related accessibility information

In addition to standard IBM help desk and support websites, IBM has established a TTY telephone service for use by deaf or hard of hearing customers to access sales and support services:

TTY service 800-IBM-3383 (800-426-3383) (within North America)

IBM and accessibility

For more information about the commitment that IBM has to accessibility, see [IBM Accessibility \(www.ibm.com/able\)](http://www.ibm.com/able).

Index

A

- accessibility features for this product [311](#)
- account management [95](#)
- activation [25](#)
- administration [3](#)
- administration pages root
 - directories [222](#)
 - files [222](#)
- advanced configuration files [239](#)
- advanced tuning [96](#)
- advanced tuning parameters [198](#), [201](#)
- alerts
 - emails [102](#)
 - remote syslog [103](#)
 - SNMP [101](#)
- analysis [43](#)
- appliances
 - change commit process [36](#), [191](#)
 - clusters [241](#)
 - disk space usage [3](#)
 - front-end load balancers [273](#), [274](#)
 - hardware [5](#)
 - installation [6](#)
 - management [31](#), [35](#)
 - management interface [201](#)
 - RESTful web services [35](#)
 - runtime environment component [198](#)
 - setup wizard [25](#)
 - snapshots [172](#)
 - tasks [5](#)
 - updating [36](#), [191](#)
- application layer [275](#)
- applications
 - databases [105](#)
 - interface [46](#), [57](#)
 - locale [106](#)
 - logs [47](#), [106](#)
 - statistics [46](#)
- applying changes [195](#)
- architecture rules [172](#)
- authentication
 - enabling [234](#)
 - mutual [201](#)
 - paths [258](#), [262](#)
 - settings [263](#)
 - testing [258](#)
- authorization servers
 - configuration files [239](#)
 - deleting [238](#)
 - editing [239](#)
 - log files [272](#)
 - management [265](#)
 - restarting [239](#)
 - starting [239](#)
 - stopping [239](#)
- auto updating [105](#)

- availability [164](#), [166](#), [173](#)
- average response time statistics [229](#)

B

- backup [3](#), [172](#)

C

- CA paths [258](#), [262](#)
- cable connections [5](#)
- certificates
 - databases
 - adding description [109](#)
 - creating [110](#)
 - deleting [112](#)
 - exporting [111](#)
 - importing [111](#)
 - listing [109](#)
 - naming [111](#)
 - replicating [112](#)
 - expiry [40](#)
 - personal [113](#)
 - requests [115](#)
 - signer [112](#)
- changes
 - configuration [36](#), [191](#), [215](#)
 - debug levels [202](#)
 - federated directory configuration [203](#)
 - front-end load balancers [273](#)
 - passwords [201](#)
 - runtime configuration [195](#)
- CLI
 - logging [269](#)
- clusters
 - data
 - loss [174](#)
 - replication [163](#), [166](#)
 - Distributed Session Cache [241](#)
 - external reference entity (ERE) [166](#)
 - failure [165](#), [168](#)
 - maintenance [171](#)
 - masters [161](#), [164](#), [173](#)
 - nodes
 - availability [173](#)
 - identifiers [173](#)
 - master nodes [168](#), [170](#), [172](#)
 - registration [69](#), [170](#), [173](#)
 - rules [172](#)
 - services [164](#)
 - support [161](#)
- command-line interface [24](#), [31](#)
- commit process [36](#), [191](#)
- common tasks [24](#)
- components
 - compress [265](#)
 - flush intervals [265](#)

- components (*continued*)
 - modifying statistics settings [225](#)
 - rollover size [265](#)
 - runtime [195](#), [201](#)
 - trace levels [265](#)
- configuration
 - authentication [90](#)
 - changes [36](#), [191](#)
 - databases [161](#), [164](#)
 - hardware appliances [5](#)
 - host names [24](#)
 - initial [29](#)
 - local management interface [170](#)
 - passwords [24](#)
 - restricted nodes [171](#)
 - system alerts [101](#)
- configuration files
 - updating [267](#)
- control
 - access lists [202](#)
 - FIPS setting [201](#)
 - statistics [225](#)
 - tracing [265](#)
- CPU graph [45](#)
- creating clusters [68](#)

D

- data
 - administration [201](#)
 - files [270](#)
 - flow
 - appliance-wide [227](#)
 - instance-specific [227](#)
 - instances [227](#)
 - management [201](#)
 - range [227](#)
 - tracing [265](#)
 - users [201](#)
- data loss [174](#)
- databases
 - LDAP keys [201](#)
- date and time [87](#)
- debugging
 - Java extension points [198](#)
 - JVM [198](#)
 - log levels [202](#)
 - messages [277](#)
 - security log [277](#)
- default
 - certificates [203](#)
 - directories [222](#)
 - display [227](#)
 - passwords [201](#)
 - ports [201](#)
 - rollover size [270](#)
 - suffixes [202](#), [203](#)
 - values [258](#), [259](#)
- default gateway [63](#)
- default route [63](#)
- demote master [69](#)
- details
 - common log files [271](#)
 - logging [269](#)

- details (*continued*)
 - replica set [283](#)
 - statistics components [225](#)
- diagnostics
 - health [43](#)
 - statistics [43](#)
 - support files [100](#)
- disk space [39](#), [46](#), [269](#)
- distributed session cache
 - cluster service [161](#), [164](#)
 - failover [165](#)
 - policy data [174](#)
- Distributed Session Cache (DSC)
 - managing [241](#)
 - registered clients [283](#)
- DN
 - local hierarchies [202](#)
- domains
 - management [243](#)
 - names [276](#)
 - realms [258](#), [261](#)
 - replica sets [283](#)
 - TCP/IP [202](#)
- dscadmin command [285](#)

E

- email response objects [102](#)
- embedded LDAP servers
 - administration [201](#)
 - log levels [202](#)
 - management [201](#)
 - SSL interface [201](#)
 - suffixes [202](#)
- event logs [43](#)
- exiting [285](#)
- external reference entity (ERE) [164](#), [166](#)
- external user registries [201](#), [203](#)

F

- failure [164](#), [165](#), [168](#)
- federated directories [203](#)
- federations
 - management [235](#)
 - reverse proxy configuration [235](#), [236](#)
- file sizes [269](#)
- files [100](#), [116](#)
- firmware
 - settings [55](#)
 - updates [171](#)
- first management interface [173](#)
- fix packs [54](#)
- flow data
 - appliance-wide [227](#)
 - instance-specific [227](#)
- flush intervals [265](#)
- front-end load balancers
 - configuration [277](#)
 - functions [273](#)
 - health status [229](#)
 - layers [275](#), [276](#)
 - network termination [276](#)

front-end load balancers (*continued*)
 persistence [275](#)
 scheduling [274](#)
 servers [273](#)
FSSO
 configuration [234](#)
functionality [273](#)

G

geolocation data [161](#), [163](#), [165](#)
getting started [5](#)
groups [95](#)
gw_net.tuning.downdelay [96](#)
gw_net.tuning.miimon [96](#)
gw_net.tuning.updelay [96](#)
gw_net.tuning.use_carrier [96](#)
gw_net.tuning.xmit_hash_policy [96](#)

H

hardware appliances [5](#), [24](#)
headers [35](#)
health status [228](#), [229](#)
high availability [166](#), [229](#), [273](#)
history [54](#)
hosts
 files [65](#)
 names [24](#)
HTTP headers
 identities [234](#)
 requests [276](#)
 responses [276](#)
Hyper-V [10](#)

I

IBM Security Access Manager
 environments [273](#)
IBM Security Verify Access
 environment [265](#)
installation
 fix packs [54](#)
 license [54](#)
instances
 deleting [238](#)
 editing [239](#)
 restarting [205](#), [239](#)
 showing current state [225](#)
 starting [205](#), [239](#)
 stopping [205](#), [239](#)
 unconfiguring [207](#)
intermediate files [100](#)
IP addresses [40](#), [173](#)

J

Java
 extension points [198](#)
junctions
 graphs [229](#)
 health status [228](#)
 standard [234](#)

junctions (*continued*)
 virtual [234](#)
JVM debugging [198](#)

K

Kerberos
 CA paths [262](#)
 default values [259](#)
 domains [261](#)
 keytab files [258](#), [263](#)
 realms [260](#)
 settings [258](#)
keytab files [258](#), [263](#)
KVM [8](#)

L

layer 7
 benefits [276](#)
 HTTP cookies [275](#)
 load balancing [275](#)
 SSL proxy [277](#)
 SSL termination [276](#)
 stickiness [275](#)
LDAP
 administration [201](#)
 log levels [202](#)
 management [201](#)
 servers [201](#)
 SSL
 settings [203](#)
 support [201](#)
 suffixes [202](#)
LDAP PIP [109](#)
LDAP policy information point [109](#)
LDAP server [109](#)
license
 agreement [24](#)
 calculation [21](#)
 installation [54](#)
 metric tool [21](#)
 overview [51](#)
 support [25](#)
 updates [51](#)
listing
 file sizes [269](#)
 names of log files [269](#)
 session management sessions [283](#), [285](#)
LMI
 appliance setup wizard [25](#)
 cluster configuration [163](#)
 commit process [36](#), [191](#)
 configure hardware appliances [5](#)
 Kerberos configuration [258](#)
 license installation [54](#)
 logging on [31](#)
 restricted nodes [171](#)
 reverse proxy management [205](#)
 supported browsers [31](#)
load balancers
 configuration [274](#), [277](#)
 front-end [273](#), [277](#)

- load balancers (*continued*)
 - health status [229](#)
 - layer 7 [276](#)
 - layers [275](#)
 - network termination [276](#)
- local management interface
 - cluster configuration [168](#), [170](#)
 - See also* LMI
- log files
 - authorization [272](#)
 - clearing [270](#)
 - exporting [269](#)
 - management [269](#)
 - names [269](#)
 - reverse proxy [271](#)
 - statistics [225](#)
 - viewing snippets [269](#)
- log levels
 - customizing [202](#)
 - LDAP servers [202](#)
- logging [269](#), [270](#)
- logs
 - files [100](#), [103](#), [106](#)
 - response objects [103](#)
- LTPA keys
 - files [234](#)

M

- maintenance [171](#)
- management
 - applications
 - logs [47](#)
 - authentication [90](#)
 - certificates
 - personal [113](#)
 - requests [115](#)
 - signer [112](#)
 - SSL [95](#), [109](#)
 - file downloads [116](#)
 - groups [95](#)
 - hosts files [65](#)
 - interface [57](#)
 - packet tracing [67](#)
 - settings
 - secure [109](#)
 - systems [51](#)
 - users [95](#)
- master ERE [166](#)
- master nodes [69](#), [161](#)
- memory statistics [44](#)
- metadata
 - user registries [201](#)

N

- network
 - settings [57](#), [277](#)
 - termination [276](#)
 - traffic [41](#)
- nist.sp800-131a.strict [96](#)
- nodes
 - clusters [161](#)

- nodes (*continued*)
 - failure [168](#)
 - promotion [69](#), [168](#)
 - restricted [171](#)
- notifications [39](#), [101](#), [102](#)
- NTP servers [87](#)

O

- objects
 - email alerts [102](#)
 - log alerts [103](#)
- offline [104](#)
- OSI network model [275](#)
- overview [51](#)

P

- packet tracing [67](#)
- partitions [41](#), [46](#), [55](#)
- password management [201](#)
- passwords [24](#)
- patches [54](#)
- persistence [273](#), [275](#)
- policy servers
 - editing [266](#)
 - management [195](#)
- primary masters [112](#), [161](#)
- Process Value Unit (PVU) report [21](#)
- product activation [25](#)
- promotion [69](#), [168](#)

Q

- quaternary masters [161](#)
- quitting [285](#)

R

- realms
 - authentication [258](#), [262](#)
 - domains [258](#), [261](#)
 - management [260](#)
 - names [258](#)
 - properties [261](#)
 - servers [258](#)
- redirection [63](#)
- references
 - cluster configuration [72](#)
 - databases [79](#)
 - session cache [73](#)
- registration [69](#), [173](#)
- replica sets
 - listing [283](#), [285](#)
 - management [241](#)
 - session termination [284](#)
 - showing [283](#)
- replicating SSL certificates [112](#), [163](#)
- replication
 - syncing [202](#)
- response objects
 - emails [102](#)
 - logs [103](#)

response objects (*continued*)

SNMP [101](#)

restarting [104](#)

RESTful web services [35](#)

restricted nodes [171](#)

reverse proxy

configuration [215](#)

federation configuration [235](#), [236](#)

graphs [227](#)

health status [228](#)

instances

restarting [205](#)

showing current state [225](#)

starting [205](#)

stopping [205](#)

unconfiguring [207](#)

log files [271](#)

management [205](#)

throughput [227](#)

traffic [227](#)

RHEV [9](#)

roles [161](#)

rollover

compress [270](#)

files [265](#)

sizes [265](#), [270](#)

root [46](#)

rules [172](#)

runtime

cluster services [164](#)

components [195](#)

configuration [198](#), [266](#)

data [163](#), [174](#)

databases [161](#), [165](#)

failure [165](#)

JVM debugging [198](#)

profiles [198](#)

restarting [195](#)

settings [166](#)

starting [195](#)

stopping [195](#)

S

scheduling [274](#), [277](#)

scheduling updates [52](#)

secondary masters [161](#)

Security Access Manager

environments [273](#)

security actions [229](#)

Security Verify Access

authorization servers [239](#)

editing tracing configuration file [266](#)

environments [265](#)

external user registries [203](#)

serial consoles [5](#)

servers

authorization

deleting [238](#)

editing [239](#)

log files [272](#)

restarting [239](#)

starting [239](#)

stopping [239](#)

servers (*continued*)

authorization (*continued*)

trace files [265](#)

definition [273](#)

LDAP [201](#)

load-balanced [229](#)

real [277](#)

virtual [277](#)

WebSEAL [234](#)

services

health status [229](#)

IP addresses [273](#)

layer 7 [277](#)

load balancing [273](#)

ports [273](#)

requests [273](#)

web [36](#), [191](#)

sessions

cache reference [73](#)

information [241](#)

listing [285](#)

management [285](#)

management replicas [283](#)

persistence [275](#)

termination [284](#)

sets [283](#)

settings

appliances [24](#)

configuration [100](#)

debug levels [202](#)

firmware [55](#)

Kerberos [258](#)

management

port [24](#)

network [57](#)

policy [100](#)

secure [109](#)

snapshots [100](#)

statistics [225](#)

systems [87](#), [163](#)

updates [52](#)

setup [5](#), [25](#)

shutting down [104](#)

signature files [69](#), [170](#)

signer certificates [112](#)

simple network management protocol (SNMP) [101](#)

snapshots [100](#)

SNMP

configuring [106](#)

software

support [265](#)

SSL

support [201](#)

termination [276](#)

SSL certificates

adding description [109](#)

exporting [111](#)

importing [111](#)

management [95](#), [109](#)

naming [109](#), [111](#)

replicating [112](#), [163](#)

signer certificates [112](#)

updating [95](#)

viewing details [95](#)

- SSL connections
 - configuring [109](#)
- stand-alone clusters [161](#)
- standard junctions [234](#)
- static routes
 - configuration [63](#)
- statistics
 - average response times [229](#)
 - control [225](#)
 - log files [225](#)
 - security actions [229](#)
 - settings [225](#)
- status column [163](#)
- stickiness [273](#), [275](#)
- storage [46](#)
- suffix management [202](#)
- support
 - clusters [161](#)
 - files [100](#)
 - license [25](#)
- syslog [103](#)
- systems
 - alerts [101](#)
 - events [43](#)
 - notifications [39](#)
 - settings [51](#), [87](#), [163](#)

T

- TCP
 - header information [275](#)
 - layers [275](#)
 - levels [277](#)
- temporary files [100](#)
- terminal emulation [5](#)
- termination
 - networks [276](#)
 - sessions [284](#)
- tertiary masters [161](#)
- threat protection [52](#)
- time zones [87](#)
- trace
 - components [265](#)
 - configuration files
 - editing [239](#), [266](#)
 - updating [267](#)
 - control [265](#)
 - data [265](#)
 - files [265](#)
 - function calls [202](#)
 - levels [265](#)
- traffic
 - reverse proxy [227](#)
- transaction logging components [229](#), [270](#)
- troubleshooting support files [100](#)
- tuning parameters [163](#)

U

- updates
 - application databases [105](#)
 - auto updating [105](#)
 - changes [172](#)

- updates (*continued*)
 - firmware [51](#), [55](#), [171](#)
 - history [54](#)
 - intrusion prevention [51](#)
 - licensing [51](#)
 - overview [51](#)
 - schedules [52](#)
 - servers [52](#)
- URL categorization [105](#)
- USB devices [269](#)
- user registries
 - embedded [201](#)
 - external [203](#)
 - log files [195](#)
 - management [201](#)
- users
 - names [234](#)
 - sessions [284](#), [285](#)

V

- virtual
 - IP addresses [273](#)
 - junctions [234](#)
 - servers [277](#)
- virtual appliances
 - installation [6](#), [8–10](#)
 - tasks [6](#), [24](#)
- VMware [6](#)

W

- web
 - content [229](#)
 - servers
 - principals [258](#)
 - services [36](#), [191](#)
- Web Portal Manager [243](#)
- Web Reverse Proxy
 - configuration
 - environments [273](#)
 - files [215](#)
 - inspections [229](#)
 - log files
 - management [269](#)
 - malicious requests [229](#)
 - management [222](#)
 - servers [273](#), [276](#)
 - transaction logs [229](#)
- Web Security Gateway Appliance
 - reverse proxy log files [269](#)
 - statistics [225](#)
 - tracing [265](#)
- web services
 - error responses [35](#)
 - required headers [35](#)
- WebSEAL
 - instances
 - display [227](#)
 - servers [234](#)
 - sessions [275](#)

X

X-Force signatures [52](#)

