IBM Security Access Manager
Version 9.0.7
June 2019

Troubleshooting Topics

IBM Security Access Manager
Version 9.0.7
June 2019

Troubleshooting Topics

IBM

Contents

Chapter 10. Trace event logs 67

Chapter 11. Common Security Access Manager problems. 83

Chapter 12. Common user registry problems. 89

Chapter 13. Single sign-on Issues: Windows Desktop single sign-on, Kerberos, and SPNEGO 93

Chapter 14. Common problems with WebSEAL servers 105

Part 1. IBM Security Access Manager troubleshooting

Chapter 1. Getting started with troubleshooting

Problem determination, or troubleshooting, is a process of determining why a product is not functioning in the expected manner.

These topics provide information to help you identify and resolve problems.

The troubleshooting process, in general, requires that you isolate and identify a problem, then seek a resolution. You can use a troubleshooting checklist to help you. If the checklist does not lead you to a resolution, you can collect additional diagnostic data and analyze it yourself. You can also submit the data to IBM® Software Support for analysis.

Avoiding potential problems

If you plan the deployment of your software, you can often prevent problems before they happen.

Before you install or upgrade Security Access Manager, review the product requirements. You can access these from a link on the Welcome page in the Knowledge Center. The link connects you with the following information:

- Supported operating system levels
- Prerequisite software requirements
- Required software patches
- Minimum memory requirements
- Disk space requirements

After you install Security Access Manager, ensure that you have a comprehensive backup and system recovery strategy in place. When you create your backup and recovery strategy, include the following information to help avoid the possibility of running into problems:

- Periodically back up the user registry by following the instructions that are provided by the user registry vendor.
- Maintain information about your environment, including system topology, IP addresses, host names, and which components are installed on each system.
- Maintain updated information that describes the key system resources that are being managed by Security Access Manager and the security policies that are being applied to them by Security Access Manager.
- Periodically check that all systems that are running Security Access Manager have sufficient disk space for runtime and problem determination data. As your security policy grows, and the number of users, groups, and protected objects increase, the space requirements for the policy databases, message logs, trace logs, and any auditing information can increase as well.
- Regularly check for the availability of fix packs and install them as they become available. Information about fix packs and other useful information can be found on the IBM Software Support Site.

Chapter 2. Diagnosing problems with diagnostic tools

When problems do occur, use the information about diagnostic tools to learn how to identify and possibly resolve them.

If you are unable to correct the problem, gather the relevant diagnostic information and then contact IBM Support to get further assistance.

- "Messages"
- "Logs" on page 6
- "Using the error messages to resolve errors" on page 6

Messages

All messages issued by IBM Security Access Manager adhere to a Message Standard.

The Message Standard specifies a standard format for all messages issued by this product. The standard, based on the IBM Message Standard, is intended to provide a consistent and meaningful way for identifying messages across the entire IBM product set. Messages issued by IBM Security Access Manager, along with detailed explanations and suggested actions, can be found in the Error messages section of the IBM Knowledge Center.

Message types

Security Access Manager is written in both the C and Java™ programming languages, with different types of messages for each programming language.

Applications that use the Security Access Manager APIs are also written in these programming languages.

Security Access Manager produces the following types of messages:

Runtime messages
Messages that are generated by applications, commands, and utilities that use the Security Access Manager Runtime component, and messages that are generated from the C language-based Security Access Manager components, such as WebSEAL. These messages are written to the runtime message logs based on their severity levels.

IBM Security Access Manager Runtime for Java messages
Messages that are generated by applications, commands, and utilities that use the IBM Security Access Manager Runtime for Java component, and messages that are generated from the Java language-based Security Access Manager components. These messages are written to the IBM Security Access Manager Runtime for Java message logs. These messages tend to provide exception and stack trace information from the JRE.

Server messages
Messages that are generated by the Security Access Manager daemons and servers. Messages from the policy server, authorization server, and WebSEAL servers are written to the server message logs.

Installation and configuration messages
Messages that are generated during installation and by the configuration utilities. Some of these messages follow the message standard and have an associated ID.

WebSEAL HTTP messages
WebSEAL provides the capability of logging HTTP messages. This message log capability is described in the Auditing topics in the Knowledge Center.

Message format

A message consists of a message identifier (ID) and message text and an error code. The error code is a unique 32-bit value. The error code is either a decimal or hexadecimal number and indicates that an operation was not successful.

All messages that follow the message standard are listed in the Error messages section of the IBM Knowledge Center. Each of these messages has a detailed explanation and suggested actions.

Message identifiers

A message ID consists of 10 alphanumeric characters that uniquely identify the message.

The message ID consists of the following parts:

- A 3-character product identifier
- A 2-character component or subsystem identifier
- A 4-digit serial or message number
- A 1-character type code that indicated one of the following message severities:

W	Warning
E	Error
I	Information

Logs

Use log files to retrieve information about a problem in your environment.

Enable the collection of detailed log and trace information to troubleshoot problems. You can collect and review standard informational log messages or detailed trace messages to help determine the root cause of a problem.

For information about viewing event logs, see “Viewing the event log” on page 19.

For information about viewing the appliance log files, see “Viewing application log files” on page 22

Using the error messages to resolve errors

Use the unique message ID associated with a message to locate detailed explanations and suggested operator responses.

The Error message section of the IBM Knowledge Center contains a list of messages in the IBM Security Access Manager logs, graphical user interfaces, and the command line.

For example, if you see the following error message in the message log:

```
CTGSI0311E The distributed session cache server was
unable to generate a new key.
```

Search for CTGSI0311E in the Error messages section for information about why the error occurred and how to resolve it. For example, the previous error message has the following information in the IBM Knowledge Center:

```
Explanation: The distributed session cache server was
unable to generate a new key.
Administrator response: Examine the distributed
session cache server logs for further details. It may be
necessary to restart the distributed session cache server
completely to correct this condition.
```

Chapter 3. Troubleshooting installation and uninstallation

This chapter describes problems that you might encounter while you install or uninstall Security Access Manager and provides information about how to determine the origin of the problem. After you determine what caused the problem, you can use the information that is provided to resolve this problem.

Before you list some of the common Security Access Manager problems that you might encounter during installation or uninstallation, it is worthwhile to mention that the cause of most common installation and uninstallation problems is one of the following failures:

- Failure to install the following prerequisite and corequisite software:
  - Operating system software
  - Operating system patches
  - Prerequisite software products
  - Prerequisite software product level and patches
- Failure to install all of the required software components for any type of Security Access Manager system
- Failure to install or configure any of the prerequisite and corequisite items properly
- Failure to adhere to all hardware prerequisites such as disk space and memory requirements

Installation directories

The installation directory for the Security Access Manager base components is specified during installation.

Security Access Manager

When you install Security Access Manager, one or more of the following components can be installed:

- Security Access Manager Runtime
- Security Access Manager Runtime for Java
- Security Access Manager Application Development Kit (ADK)

The default installation location for Security Access Manager files is platform-dependent:

Windows operating systems

```
C:\Program Files\Tivoli\Policy Director
```

AIX, Linux, and Solaris operating systems

```
/opt/PolicyDirector
```

During the installation of Security Access Manager, the PD_HOME environment variable is set to the installation directory on Windows operating systems. No environment variable is set on AIX, Linux, and Solaris operating systems. After installation, ensure that only trusted users and groups have access to this directory and its subdirectories.

Common installation problems

This section describes problems that you might encounter while you install Security Access Manager and provides information about how to manage the problem.

Insufficient disk space

Installation and use of Security Access Manager requires adequate disk space.

If you do not have sufficient disk space during installation, an error message stops the installation and alerts you that there is not enough space.

If you encounter this error message, see review the product requirements for detailed information about required disk space. See the Welcome page in the Knowledge Center for a link to the product requirements. Clear adequate space on the disk, or select a root directory on a partition with more space, and run the installation process again.

Without adequate disk space, Security Access Manager cannot install or function as expected.

Windows disk space

On Windows operating systems, concerns for adequate disk space include the following directories:

- The Security Access Manager installation directory
- The Security Access Manager WebSEAL installation directory
- The Security Access Manager Plug-in for Web Servers installation directory

AIX, Linux, and Solaris disk space

On AIX, Linux, and Solaris operating systems, concerns for adequate disk space include the `/opt` and the `/var` directories.

Use the **df** command with the **–k** option to display the free disk space for each file system. The **–k** option causes the disk space to be displayed in kilobytes.

Installation of the license fails on Linux x86_64 systems

If the installation of the license fails when you use the isamLicense script on a 64-bit Linux system, ensure that the following 32-bit libraries are installed from the *i686.rpm packages:

```
ld-linux.so.2
libstdc++.so.6
```

Multiple network interfaces

Some operating systems can be configured with multiple network interfaces. When there are multiple network interface aliases, there might be more than one route to the policy server. In these situations, the operating system might choose a different route for each communication.

When the operating system routes each communication differently, the policy server might not be able to definitively identify the client. When the policy server cannot identify the client, the communication between the client and the policy server might fail with a message similar to the following error:

```
The server lost the client authentication, because of session expiration.
```

This communication failure can happen between the following components:

- An authorization API server in local mode with the policy server
- An authorization API server in remote mode with the policy server
- An authorization API server in remote mode with the authorization server
- The **pdadmin** utility with the policy server
- An administration API with the policy server
- The policy server with any authorization API server, such as the authorization server or WebSEAL
- The **svrsslcfg** utility with the policy server

To prevent this problem, use one of the following mechanisms:

- Change the operating system routing table so that the same route is always selected. For example, if there are three routes, two of these routes must be downgraded so that one route is always selected. For more information about route commands and metrics that are used in routing tables, see your operating system documentation.
- Set the PD_FIXED_CLIENT_IP environment variable to the IP address of a valid network interface on the operating system. This value must be in the IP version 4 (IPv4) or IP version 6 (IPv6) format. The PD_FIXED_CLIENT_IP environment variable can be set on all the supported operating systems. See RFC 2460 at the following website to determine what constitutes a valid representation of an IPv6 address:

  http://www.faqs.org/rfcs/rfc2460.html

Java error on Windows during Launchpad or script installation

On some Windows 2008 systems, the IBM Java Runtime fails to install during a silent installation. Security Access Manager component installations on Windows require IBM Java Runtime to complete.

If IBM Java Runtime failed to install during a Launchpad or script installation, Security Access Manager component installations fail with the following error message:

```
Unable to find Java executable file (javaw.exe or java.exe).
Install the IBM SDK for Java or make sure a Java version 1.4 or
higher is accessible from the current PATH.
```

To continue with the installation, ensure that IBM Java is available in the environment. Install IBM Java manually by completing the procedure for installing IBM Java Runtime on Windows in the ADK installation topics in the Knowledge Center.

GSKit installation failure on Windows

During installation on Windows, GSKit can fail to install if the user ID is not correct and UAC settings are set to notify you when a program tries to make changes to the computer.

When this issue occurs from Launchpad, the `ISAMGskitInstall.log` shows an error such as the following error:

```
C:\build\bin\../windows/GSKit/gsk8ssl64.exe /s /v/quiet
1625
```

This error can also be seen when you run the GSKit installation from the install_isam.bat script.

This issue can be caused by using the user ID "administrator" instead of "Administrator."

To resolve this issue, run the Launchpad installation with the user ID of "Administrator" with a capital "A."

Alternatively, you can do one of the following options:
- Run the gsk8ssl64.exe installer manually.
- Disable UAC before installing GSKit. This method can require a system restart.

Installation stalls on Windows

When you install components on Windows, the installation program might appear stalled on the Welcome page.

After the Welcome page is displayed, the license is displayed. Check all open windows. The license page is in a separate window that might be hidden behind the Welcome page window. Minimize or move the Welcome page window, read and accept the license, and then return to the installation window to continue the installation.

"Stop the script" message on Windows

When you install components on Windows using Launchpad, a message might be displayed that asks if you want to stop the running script.

The message indicates that the running script is causing the web browser to run slowly. Answer No to this message to continue installing IBM Security Access Manager.

Installation logs

When you install and configure Security Access Manager components, log files are created. If you natively install and configure, there are separate log files for installation and configuration.

Command line installation log files

Security Access Manager native installation log files contain the completion status for the installation tasks performed. If the components are installed with the installation programs that are provided with the operating system, these are the only installation log files that are created. These native installation log files contain messages that are generated during the installation of the product.

The names and locations of the native installation log files are shown in Table 1.

Table 1. Native installation log files

Operating system	Command line installation log file
AIX®	/smit.log
Linux	/tmp/install.log or /var/tmp/install.log
Solaris	See the contents of the pkginfo files that are stored in the subdirectories of the /var/sadm/pkg directory.

Table 1. Native installation log files (continued)

Operating system	Command line installation log file
Windows	`%PD_HOME%\log\msg__PDInstall.log` **Note:** There are two underscore characters (_) in the file name.

Command-line configuration log files

The Security Access Manager **pdconfig** command is used to configure Security Access Manager. Similar configuration commands are used for the Web Security components. Messages that are generated during the configuration process are stored within Security Access Manager configuration log files.

The locations of these configuration log files are listed in Table 2 for Windows operating systems and Table 3 for AIX, Linux, and Solaris operating systems.

Table 2. Default locations for command-line configuration log files on Windows

Component	Command-line configuration log file on Windows
C runtime	`%PD_HOME%\log\msg__config.log`
Runtime for Java	`%PD_HOME%\log\msg__PDJrteCfg1.log`

Table 3. Default locations for command-line configuration log files on AIX, Linux, or Solaris

Component	Command-line configuration log file on AIX, Linux, or Solaris
C runtime	none
Runtime for Java	`/var/PolicyDirector/log/msg__PDJrteCfg1.log`

Registry entries

This section lists the registry entries that are created when you install Security Access Manager on Windows.

Attention: Do not modify any registry entry unless directed by IBM Support.

Registry entries for Security Access Manager are stored under the following directory:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\
```

If part of a registry key is missing or a key is not removed during uninstallation, there can be problems with installing or upgrading the product.

Registry entries

Security Access Manager creates the following registry entries:

Access Manager Web Security Runtime:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Access Manager Web
Security Runtime
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Access Manager Web
Security Runtime\8.0.0
      Path  install_location
      Version 8.0.0.2
```

License

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director License
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director License\7.0.0
      Path  install_location
      Version 8.0.0.2
```

Runtime:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director Runtime
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director Runtime\7.0.0
      Configured   Yes/No
      Path  install_location
      Version 8.0.0.2
```

Application Development Kit:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director
Authorization Toolkit
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director
Authorization Toolkit\8.0.0
      Path  install_location
      Version 8.0.0.2
```

Java Runtime:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director
Java Runtime
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Policy Director
Java Runtime\8.0.0
      Path  install_location
      Version 8.0.0.2
```

Security Utilities:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Tivoli\Tivoli
Security Utilities
     MajorVersion 8.0
     Path  install_location
     Version 8.0.0.2
```

Common configuration problems

This section details common problems that you might encounter during configuration of Security Access Manager.

Invalid LDAP management domain location DN causes error

Security Access Manager uses the LDAP client `ldapsearch` command to verify the validity of the DN location.

Security Access Manager uses the LDAP client `ldapsearch` command to verify the validity of the DN location. Under the following condition, `ldapsearch` generates a success message instead of an error:

- The error that is returned from the LDAP server is related to a referral chasing error.

As a result, Security Access Manager generates the following misleading error message during configuration:

```
A policy server is already configured to this LDAP server.
A second would be used as standby server only.
```

To resolve this issue, enter a valid value for the LDAP management domain location DN during the Security Access Manager policy server configuration.

Tivoli Directory Server configuration fails contacting LDAP server

If automated configuration scripts for IBM Tivoli Directory Server fail or if configuration fails from the Launchpad, review the logs to examine the cause of failure.

For example, the log can show the following error:

```
/opt/IBM/ldap/V6.3/bin/idsldapadd -p 389 -D cn=root -w \? -f /tmp/org_ldif
Enter password ++>
ldap_simple_bind: Can't contact LDAP server
The return code is : 81
ERROR: Problem adding the sample ldif file :/tmp/org_ldif
```

If the log shows a return code of 81 from an automated configuration scripts command such as **`IDSConfigServerSSL.sh`**, this code indicates that the LDAP server cannot be contacted. From the Launchpad, you might see the message `Error: The configuration failed with return code: 5` displayed. This error is often a timing issue.

For more information about IBM Tivoli Directory return codes, see the IBM Tivoli Directory Server documentation.

To resolve the issue, run the command that is listed in the log as failing from command line. If the error happens from Launchpad, retry the command by clicking the **Configure IBM Tivoli Directory Server** button.

Recovering from failed Security Directory Server automated script configuration on Solaris

The **`idsdefinst`** script might fail to create the Security Directory Server default instance and suffix if you did not properly set the kernel parameters. DB2® requires sufficient memory to complete the request.

About this task

If DB2 has insufficient memory, then an error, similar to the following error, is in the `db2cli.log` file:

```
2012-08-29-09:08:53.native retcode = -1084; state = "57019"; message = "SQL1084C
Shared memory segments cannot be allocated.  SQLSTATE=57019
```

You must clean your system before you run the **`idsdefinst`** script again.

Procedure

1. Verify the installation path of DB2. For example:

```
# /usr/local/bin/db2ls

Install Path       Level     Fix Pack    Install Date          Installer UID
-----------------------------------------------------------------------------
/opt/ibm/db2/V9.7   9.7.0.5      5    on Dec 12 16:25:10 2011 CST       0cd
```

2. Run the **db2ilist**, **db2idrop**, and **db2iset** commands from the DB2 instance directory to remove the instance. The following example uses:
   - The installation path from step 1.
   - dsrdbm01 as the instance.

   ```
   # cd /opt/ibm/db2/V9.7/instance
   # ./db2ilist
   dsrdbm01

   #./db2idrop dsrdbm01

   # ./db2iset -d dsrdbm01
   ```
3. Confirm that the instance is removed by issuing the following command:

   ```
   # ./db2ilist
   ```
4. Open the system variables file in the /etc/system directory.
5. Add the following lines to the end of the file to set the kernel parameters appropriately. The following values are suggested as starting values:

   ```
   set msgsys:msginfo_msgmax = 65535
   set msgsys:msginfo_msgmnb = 65535
   set shmsys:shminfo_shmmax = 2134020096
   ```

   For more information, see the Solaris tuning documentation.
6. Remove any Security Directory Server instance by using the following command:

   ```
   idsidrop -I instance_name -r
   ```

   where *instance_name* is the name of the instance. The idsidrop file is in the /opt/IBM/ldap/V6.3/sbin directory.
7. Run the **idsdefinst** script again to define the default database instance.

Upgrade common problems

This section details common problems that you might encounter when you upgrade Security Access Manager.

Cannot create users or groups after upgrade

Security Access Manager does not have authority to create users and groups after Security Directory Server is upgraded. If Security Directory Server is your user registry, and you are upgrading Security Access Manager, then the Security Directory Server component must be migrated first, if all components are on the same machine.

Complete the migration of Security Directory Server by following the instructions in the *IBM Security Access Manager for Web Upgrade Guide*.

These instructions guide you through the process of backing up the current data with the **db2ldif** utility, upgrading the Security Directory Server, and restoring the data with the **bulkload** utility.

When you use the **bulkload** utility, specify the **–A yes** option to have it properly process Access Control List (ACL) updates. If the ACLs are not loaded properly, Security Access Manager does not have the authority to complete the needed tasks to create and maintain user and group information.

If bulkload fails to update the ACLs properly and these symptoms occur, you can create the ACLs manually by following the "Applying Access Manager ACLs to new LDAP suffixes" procedure in the Administering topics in the Knowledge Center. Using the Web Administration Tool, apply the ACLs to all existing LDAP suffixes and `secAuthority=Default` entries below all defined users in the LDAP server. Applying these ACLs restores the correct authority to allow Security Access Manager to continue.

Chapter 4. Troubleshooting on the appliance

Use appliance options to view system resources and logs.

You can use the **Monitoring: Analysis and Diagnostics** appliance menus to view the use of resources such as memory, CPU, and storage. Also, review the known problems and limitations with deployment on the appliance.

Running self-diagnostic tests (hardware appliance only)

The hardware appliance provides a self-diagnostic program to assist with troubleshooting. This feature is not available with the virtual appliance.

Procedure

1. Reboot the appliance.
2. From the console, select the **Hardware Diagnostics** option from the GNU GRUB menu. The diagnostics program starts running after this option is selected. After the appliance finishes booting, the diagnostics program is accessible from the console.
3. To run a diagnostic test, enter `phdiag <test_name>`, where `<test_name>` is one of the following values in bold:

   all Run all standard tests (No storage bad blocks test, default)

   lcd Run LCD tests

   system
   Run standard system tests (MTM-Serial, Inventory, PSU, FAN, SEL)

   network
   Run network port tests (Selftest, Traffic)

   storage
   Run storage tests (SMART, FSCK)

   badblocks
   Run storage bad blocks test

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

1. Click **Monitor Analysis and Diagnostics** > **Logs** > **Event Log**. The system events displayed.
2. Click **Pause Live Streaming** to stop the live updating of the event log.
3. Click **Start Live Streaming** to resume live updating of the event log.

Viewing memory statistics

View the memory graph to see the memory utilization of the IBM Security Access Manager appliance.

Procedure

1. Click **Monitor Analysis and Diagnostics** > **System Graphs** > **Memory**.
2. Select a **Date Range**:

Option	Description
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 IBM Security Access Manager.

Procedure

1. Click **Monitor Analysis and Diagnostics** > **System Graphs** > **CPU**.
2. Select a **Date Range**:

Option	Description
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 IBM Security Access Manager appliance.

Procedure

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

Option	**Description**
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 interface statistics

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

Procedure

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

Option	**Description**
1 Day	Displays data for every 20-minute interval in one day.
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 Access Manager.

Procedure

1. From the top menu, select **Monitor Analysis and Diagnostics** > **Application Log Files**. The displayed directories contain the application log files that can be viewed and downloaded:
   - **access_control**: Log files for the Advanced Access Control Module.
   - **cluster**: Logs files for the cluster manager.
   - **dsc**: Log files for distributed session cache.
   - **federation**: Log files for the Federation Module.
   - **felb**: Log files for the front-end load balancer.
   - **isam_runtime**: Log files for the runtime.
   - **management_ui**: Log files for the management interface.
   - **snmp**: Log files for the SNMP manager.
2. Optional: Click **Refresh** to get the most up-to-date data.
3. 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 **Download** to save the file to your local drive.

   c. Confirm the save operation in the browser window that pops up.

Tuning runtime application parameters and tracing specifications

To manually tune selected runtime application parameters and tracing specifications, use the Runtime Parameters management page.

Procedure

1. From the top menu, select **Secure Access Control** > **Global Settings** > **Runtime Parameters** or **Secure Federation** > **Global Settings** > **Runtime Parameters**. This page contains three panels: **Runtime Status**, **Runtime Tuning Parameters**, and **Runtime Tracing**.
2. Perform one or more of the following actions to tune your runtime.

   Note: Certain changes might require a restart of the runtime before they can take effect.

   Disable automatic restart of the runtime

   By default, the runtime is automatically restarted after certain changes are made. You can disable this automatic restart function if you prefer manual restarts.

   a. On the **Runtime Tuning Parameters** panel, select **Auto Restart**.
   b. Click **Edit**.
   c. In the Auto Restart window, define the value as **False**.
   d. Click **OK**.

   View the status of the runtime and restart the runtime

   a. Select the **Runtime Status** panel. The status of local and clustered runtimes are displayed.
      - Under **Local Runtime Status**, you can view the runtime operational status, when it was last started, and whether a restart is outstanding. If the value of the **Restart Required** field is **True**, it means that the runtime must be restarted for some changes to take effect.
      - Under **Clustered Runtime Status**, all nodes in the cluster are listed.
        - The **Master** column indicates whether a node is the cluster master.
        - The **Runtime Status** column indicates whether a node is running or stopped.
        - The **Changes Active** column indicates whether changes made to the cluster configuration are active on this node. Having a green indicator in this column means that all changes made are already active. Having a yellow indicator in this column means that this node must be restarted before some changes can take effect.
   b. Depending on which runtime you want to restart, click **Restart Local Runtime** or **Restart All Clustered Runtimes**.

   Modify the maximum or initial heap size

   These parameters indicate the maximum and initial heap size in megabytes for the runtime Java virtual machine.

   a. On the **Runtime Tuning Parameters** panel, select **Max Heap Size** or **Initial Heap Size**.
   b. Click **Edit**.
   c. In the Max Heap Size or Initial Heap Size window, enter the heap size value as needed.
   d. Click **OK**.

Modify the minimum or maximum threads

These parameters indicate the minimum number of core threads that the runtime server starts with and the maximum number of threads that can be associated with the runtime server.

If the minimum value is not set or is set as -1, a default value is calculated based on the number of hardware threads on the system.

If the maximum value is not set or is set as 0 or less, a default value of unbounded is used.

The minimum **cannot** be set to a value larger than the maximum.

a. On the **Runtime Tuning Parameters** panel, select **Min Threads** or **Max Threads**.
b. Click **Edit**.
c. In the Min Threads or Max Threads window, enter the required value.
d. Click **OK**.

Modify whether to suppress sensitive trace

Enabling this parameter prevents sensitive information from being exposed in log and trace files. Examples of such sensitive information include bytes received over a network connection.

a. On the **Runtime Tuning Parameters** panel, select **Suppress Sensitive Trace**.
b. Click **Edit**.
c. In the Suppress Sensitive Trace window, select or clear the check box as needed.
d. Click **OK**.

Modify console log level

Console log level controls the granularity of messages that go to the `console.log` file.

a. On the **Runtime Tuning Parameters** panel, select **Console Log Level**.
b. Click **Edit**.
c. In the Console Log Level window, select the new value from the list.
d. Click **OK**.

Set whether to accept client certificates

This parameter controls whether the server accepts client certificates as a form of authentication.

a. On the **Runtime Tuning Parameters** panel, select **Accept Client Certificates**.
b. Click **Edit**.
c. In the Accept Client Certificates window, select or clear the check box as needed.
d. Click **OK**.

Set session invalidation timeout

This parameter defines the amount of time a session can remain unused before it is no longer valid.

Note: The default setting is **Unset**. When this setting is used, the session invalidation timeout is 1800 seconds.

a. On the **Runtime Tuning Parameters** panel, select **Session Invalidation Timeout**.
b. Click **Edit**.
c. In the Session Invalidation Timeout window, define the value in seconds.
d. Click **OK**.

Set session reaper poll interval

This parameter defines the wake-up interval in seconds for the process that removes invalid sessions. The minimum value is 30 seconds.

The default setting is **Unset**. When this setting is used, or if a value less than the minimum is entered, an appropriate value is automatically determined and used. This value overrides the default installation value, which is 30 - 360 seconds, based on the session invalidation timeout value. Because the default session invalidation timeout is 1800 seconds, the reaper interval is usually between 120 and 180 seconds.

a. On the **Runtime Tuning Parameters** panel, select **Session Reaper Poll Interval**.
b. Click **Edit**.
c. In the Session Reaper Poll Interval window, define the value in seconds.
d. Click **OK**.

Set the keystore that is used by the runtime server

This parameter defines the key database that contains the runtime server's private key.

a. On the **Runtime Tuning Parameters** panel, select **Keystore**.
b. Click **Edit**.
c. In the Keystore window, select the key database from the list.
d. Click **OK**.

Set the truststore that is used by the runtime server

This parameter defines the key database that contains keys that are trusted by the runtime server

a. On the **Runtime Tuning Parameters** panel, select **Truststore**.
b. Click **Edit**.
c. In the Truststore window, select the key database from the list.
d. Click **OK**.

Configure an outbound HTTP proxy

You must specify values for the properties for the HTTP proxy. You might also need to import the root CA certificate from the proxy. See the instructions that follow.

Table 4. HTTP proxy properties

Name	Sample Value	Description
`http.proxyHost`	`http.proxy.ibm.com`	The hostname or IP address of the HTTP proxy
`http.proxyPort`	`3128`	The port of the HTTP proxy

Table 4. HTTP proxy properties (continued)

Name	Sample Value	Description
`https.proxyHost`	`https.proxy.ibm.com`	The hostname or IP address of the HTTPS proxy
`https.proxyPort`	`3128`	The port of the HTTPS proxy

a. For each property in the table above:
   1) On the **Runtime Tuning Parameters** panel, select the property.
   2) Click **Edit**.
   3) In the property window, enter the value. See the sample values in the table.
   4) Click **OK**.

b. When all properties are set, follow the prompt to deploy the pending changes.

Certain functions, such as the OpenID connect single sign-on flow, require the root CA certificate of the outbound HTTP proxy to be imported to the Security Access Manager runtime keystore.

Complete the following steps:

a. Go to your HTTP Proxy application and obtain the necessary certificate for exchange. The exact steps to take are specific to the proxy application. Place the certificate on the local file system where it can be accessed by the appliance.

b. On the Security Access Manager system, log in to the local management interface and select **Manage System Settings** > **Secure Settings** > **SSL Certificates**

c. Select the `rt_profile_keys` keystore.

d. Select **Manage** > **Edit SSL Certificate Database**.

e. Select **Manage** > **Import**.

f. On the Signer Certificate panel, browse to locate the **Certificate File**. Enter a **Certificate Label**. Click **Import**.

g. Deploy the changes.

Delete the value of a parameter

Use this button to delete the existing value of a parameter.

a. Select the parameter to reset the value for.

b. Click **Delete**. The value of the parameter is then changed to `Unset`.

Manage the application interface on which the runtime listens

a. On the **Runtime Tuning Parameters** panel, under **Runtime Listening Interfaces**, you can add, edit, or delete a listening interface.

   To add a listening interface

   1) Click **Add**.
   2) In the Runtime Listening Interfaces window, select the listening interface from the list.
   3) Specify the listening port.
   4) Select the **SSL** check box if security is required.
   5) Click **OK**.

   To modify a listening interface

1) Select the listening interface to edit.
2) Click **Edit**.
3) In the Runtime Listening Interfaces window, edit the values as needed.
4) Click **OK** to save the changes.

To delete a listening interface

1) Select the listening interface to delete.
2) Select **Delete**.
3) Confirm the deletion.

Manage tracing specification

a. Select the **Runtime Tracing** link from the top of this page. You can also access this panel from the top menu by selecting **Monitor Analysis and Diagnostics** > **Logs** > **Runtime Tracing**.

b. 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 5. 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.
`WARNING`	Potential error or impending error. This level can also indicate a progressive failure. For example: the potential leaking of resources

Table 5. Valid trace levels (continued). The following table contains the valid trace levels.

Level	Significance
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.isam.*=INFO
  ```

c. Click **Save**.

3. When you make changes, the appliance displays a message that there are undeployed changes. If you have finished making changes, deploy them.

Monitoring log files in the command-line interface

Use the **monitor** menu in the command-line interface so that you can easily monitor the status of the log files.

Procedure

1. Access the command-line interface of the appliance by using either an ssh session or the console.
2. Enter `isam`.
3. Enter `logs`.
4. Enter `monitor`.
5. Select the log file that you want to monitor.

Known issues and solutions

Use the solutions to troubleshoot issues that you might encounter.

The known issues are:

- "Fix for the Sweet32 Birthday vulnerability causes GSKit to disable the DES ciphers" on page 29
- "Enabling Compatibility View in Internet Explorer 9 returns the Browser not Supported message" on page 29
- "Help page content does not display" on page 30
- "On Windows operating systems, you cannot use basic authentication for WebSEAL from IBM Security Access Manager for Web" on page 30
- "Database rolls back with an error when you attempt to remove a large quantity of records from a DB2 runtime database" on page 30
- "A cluster configuration update fails to deploy and generates a timeout error message" on page 30
- "Database Maintenance panel returns a retrieval error" on page 31
- "Cannot export and import a template page file in the same session" on page 31

- “Kerberos Configuration is reset to the default values on the appliance” on page 31

Fix for the Sweet32 Birthday vulnerability causes GSKit to disable the DES ciphers

In order to address the vulnerabilities exploited by theSWEET32 Birthday attack (CVE-2016-2183), ISAM GSKIT implemented a limit on the amount of data which can be downloaded from a single connection (32GB). The support for the DES ciphers is NOT affected. This means that if you have a scan tool which only checks to see if these ciphers are enabled, you could end up with a false positive for this issue.

If you want to disable the ciphers, you can Disable the DES ciphers as per: http://www-01.ibm.com/support/docview.wss?uid=swg21698249

After the 32GB limit is reached, the connection will be broken and an error will be logged such as:

```
2017-05-05-12:10:01.647+10:00I----- 0x38AD5425 webseald ERROR wiv
socket WsSslListener.cpp 1867 0x7f88fc161700 -- DPWIV1061E
Could not write to socket (445)

445 is GSK_ERROR_BYTECOUNT_EXHAUSTED
```

If requests/applications are negatively effected by the GSK_ERROR_BYTECOUNT_EXHAUSTED (445) error, the best solution would be to disable the Sweet32 ciphers either in the application logging the error (recommended if possible in your environment) or if this is not possible then disabling them in the other application sharing this connection.

For example if WebSEAL is logging the error: http://www-01.ibm.com/support/docview.wss?uid=swg21698249

Enabling Compatibility View in Internet Explorer 9 returns the Browser not Supported message

The IBM Security Access Manager appliance does not support the browser operating in this mode. The following message is displayed in the local management interface:

```
Browser not Supported. The IBM Security Access Manager appliance does
not support this browser.

The following browsers are currently supported:

- Internet Explorer 9 or later
- Firefox 17.0 or later
- Google Chrome 27.0 or later
```

Solution:

The appliance does not support Internet Explorer if the Compatibility View is turned on. Ensure that the Compatibility View in Internet Explorer is turned off. The **Compatibility View** option is under the **Tools** menu in the Internet Explorer browser.

Help page content does not display

When you click the **Help** link from the appliance user interface while using Microsoft Internet Explorer version 9.0 or later, the topic content might not display.

Solution:

Ensure that the Compatibility View in Internet Explorer is turned on. The **Compatibility View** option is under the **Tools** menu in the Internet Explorer browser. The Help System supports compatibility mode.

On Windows operating systems, you cannot use basic authentication for WebSEAL from IBM Security Access Manager for Web

WebSEAL does not start properly if your configuration meets all of these conditions:

- Windows operating system
- WebSEAL from IBM Security Access Manager for Web 7.0.0.1
- Basic authentication configured in the WebSEAL configuration file with `basic-auth-user` and `basic-auth-passwd` entries in the `[rtss-cluster:cluster1]` stanza

Solution:

To work around this issue, configure certificate authentication for WebSEAL. See the WebSEAL administration information in the Knowledge Center.

If you are using the **`isamcfg`** tool to configure WebSEAL, be sure to select certificate authentication for the authentication method response.

Database rolls back with an error when you attempt to remove a large quantity of records from a DB2 runtime database

When you try to delete many device fingerprints or user session data records from an external DB2 runtime database, the following error might occur:

```
Error occurred. The database was rolled back to the previous version.
The transaction log for the database is full. SQLCODE=-964, SQLSTATE=57011
```

Solution:

Increase the log capacity by completing the following actions:

- Increase the number of primary and secondary transaction log files.
- Increase the size of each transaction log.

For information about the available transaction log configuration parameters, see the DB2 documentation.

A cluster configuration update fails to deploy and generates a timeout error message

An update to the cluster configuration, such as the External Reference Entity IP address or First Port value, might fail to deploy in the allotted time. The following error message is printed in the event log:

```
WGASY0007E The pending changes failed to deploy within the allotted time.
```

Solution:

Increase the **wga.cmd.timeout** value. In the local management interface, select **Manage System Settings** > **System Settings** > **Advanced Tuning Parameters**. Add a parameter that is called **wga.cmd.timeout** and set the timeout value in seconds. The default value is 300 seconds.

Database Maintenance panel returns a retrieval error

The following error message returns in the **Database Maintenance** panel after the location of the runtime database is changed:

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

This error message returns after the location of the runtime database is changed from **Local to the cluster** to **Remote to the cluster** under the **Database** tab in the **Cluster Configuration** panel.

Solution: Complete the following steps to restart the local management interface:

1. Use an ssh session to access the local management interface.
2. Log in as the administrator.
3. Type `lmi`, and press Enter.
4. Type `restart`, and press Enter.
5. Type `exit`, and press Enter.

Cannot export and import a template page file in the same session

If you export a template page file and immediately try to import a file, no action occurs, and the file is not imported.

Solution: After you export a file, refresh the browser before you try to import a file.

Kerberos Configuration is reset to the default values on the appliance

After saving and deploying the Kerberos Configuration settings, you might find that the values of these Kerberos settings are reset to the defaults. For example `default_realm` is set back to `tbd`.

In a clustered Security Access Manager environment, when you configure the Kerberos settings on a node appliance, the values you modified are set back to the default values. Or, the settings are overwritten by the values set on the primary master of the cluster. This happens after some time.

Specifically, propagation of the values happens once one of the following conditions is met:

- The **Replicate with Cluster** option is toggled **on** from the local management interface runtime component panel. And, the settings are changed on the master.
- Sufficient time has elapsed since the last propagation (approximately 15 minutes).

The resetting of these values is evident on the **Defaults**, **Realms**, **Domains**, and **CA Paths** tabs, but not on the **Keyfiles** tab.

In general, when you choose to replicate the runtime environment with the cluster, these values are synchronized from the primary master. Therefore, if you change these values on a node other than the primary master, they are overwritten with the values from the primary master during the next synchronization operation.

Solution: When your configuration is set to replicate the runtime environment with the cluster, always update the Kerberos Configuration settings on the primary master of the cluster, instead of on a node. You must also import the keytab file on each node that is using the Kerberos Configuration settings.

Known limitations

Consider these known limitations when you are configuring the IBM Security Access Manager environment.

Local management interface (LMI) session timeouts

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, you are automatically logged out and any unsaved data on the current page is lost.

Save your configuration updates in the LMI regularly to avoid data loss in the event of a session timeout.

Chapter 5. Verifying the deployment

The installation of Security Access Manager involves the installation and configuration of a number of prerequisites.

These prerequisites can include Security Access Manager components. Operational failures can result from the failure to install or correctly configure a prerequisite.

If you have a failure in the following instances, always examine your installed software:

- A new Security Access Manager installation fails to work properly
- An existing Security Access Manager installation fails to work properly after you update a prerequisite or Security Access Manager components

For more information, see the product requirements link on the Welcome page in the Knowledge Center.

System types in a deployment

A typical deployment of Security Access Manager involves the installation and configuration of a number of systems.

All Security Access Manager deployments include several types of Security Access Manager systems that are set up in a secure domain, including:

- Base systems
- Web Security systems

Each system type has software prerequisites that include Security Access Manager components. You must successfully install and configure both the prerequisites and Security Access Manager components to avoid operational problems in your environment.

For example, the Security Access Manager policy server requires installation of the following components:

- IBM Global Security Kit (GSKit)
- IBM Security Directory Server client (depending on the registry used)
- IBM Security Utilities
- Security Access Manager License
- Security Access Manager Runtime
- Security Access Manager Policy Server

Base systems

Table 6 on page 34 lists the types of Security Access Manager base systems that you can set up in your secure domain.

Note:

You must install the Security Directory Server client on each system that runs Security Access Manager, with the following exceptions:

- The Security Access Manager system is on a supported Windows operating system that is either the Active Directory domain or is joined to the Active Directory domain where the policy server is to be configured.
- You are setting up an attribute retrieval service or IBM Security Access Manager Runtime for Java,

Table 6. Required components for the Security Access Manager base systems

System type	Required software components
C development	• Global Security Kit • Security Directory Server client (depending on the registry used) • Security Access Manager Runtime • Security Access Manager Application Development Kit
Security Directory Server	If you plan to install the Security Directory Server as your Security Access Manager registry, the following components are required: • Global Security Kit • Security Directory Server client • DB2 Enterprise Server Edition **Note:** Refer to the Security Directory Server documentation for information about which versions of the server are supported.
Runtime for Java	• IBM Security Access Manager Runtime for Java
Runtime	• Global Security Kit • Security Directory Server client (depending on the registry used) • Security Access Manager Runtime

Checking installed software

Security Access Manager provides commands and utilities to help you determine whether the correct software, and software level, is installed on any operating system.

Use the **pdversion** utility to list Security Access Manager components that are installed on the system along with their version number. This utility does not list prerequisite software, such as IBM Global Security Kit (GSKit).

You can also determine the presence or absence of prerequisite software with operating system utilities. The utilities, by operating system, are as follows:

AIX The **lslpp –l** command

Linux The **rpm –qa** command

Solaris
The **pkginfo –l** command

Windows
The Add/Remove Programs facility from the Control Panel

Use these tools with the product requirements that are listed in the Clearinghouse reports you can obtain from the product requirement link on the Welcome page in the Knowledge Center. Ensure that all the required software is installed on each system in your Security Access Manager deployment.

Verifying Global Security Kit

Ensure that the Global Security Kit is installed and working properly for your environment.

The two most common problems with the IBM Global Security Kit (GSKit) include:

- GSKit was not installed
- The wrong version of GSKit is installed, left over from a previous installation

GSKit on Windows operating systems does not add an entry to the Add/Remove Program list. On Windows operating systems, GSKit is typically installed in the following directory:

```
C:\Program Files\IBM\gsk8
```

Use the following command on Windows operating systems to display the GSKit version:

```
C:\Program Files\IBM\gsk8\bin\gsk8ver_64
```

Verifying user registries

This section provides information about verifying the different Security Access Manager user registries.

Security Directory Server

This section provides information about verifying Security Directory Server when it is used as the Security Access Manager user registry.

Verifying the server

Communication between the Security Directory Server client and the LDAP server can be tested by using the **ldapsearch** command. This command also reveals the version of the LDAP server software.

This command (entered as one line) can be run from any machine with Security Directory Server client installed. The structure of the **ldapsearch** command different when you use SSL.

Without SSL

The following sample command is appropriate when the LDAP server is configured for non-SSL communication:

```
ldapsearch -h ldapserver-hostname -p 389 -D "ldapadminDN" \
-w ldapadmin-password -b "" -s base objectclass=*
```

Note: If this command fails on a Windows operating system, check that the **ldapsearch** command is the one that is provided by Security Directory Server client.

The output of this command varies depending on which supported LDAP server you are using.

With SSL

The following sample command is appropriate when the LDAP server is configured for SSL communication:

```
ldapsearch -h ldapserver-hostname -p 636 -D "ldapadminDN" \
-w ldapadmin-password -Z -K client-keyfile \
-P key-password -b "" -s base objectclass=*
```

The output of this command varies depending on which supported LDAP server you are using.

Verifying the client

The previous verification procedure for the LDAP server used the Security Directory Server client that was installed on the LDAP server system itself. However, unless you are using Active Directory server for your Security Access Manager user registry, each Security Access Manager system requires the installation of the Security Directory Server client, not just the machine with the LDAP server installed.

The previous verification procedure for the LDAP server is also appropriate for verifying the functions of the Security Directory Server client on each Security Access Manager system.

Microsoft Active Directory

This section provides information about verifying Microsoft Active Directory when it is used as the Security Access Manager user registry.

Verifying the configuration

Verify that the configuration of the server and client completed successfully.

Server Start the MMC for Active Directory by selecting **Start** → **Program** → **Administrative Tools** → **Active Directory Users and Computers**.

If the Active Directory management console started and you can browse all the objects in Active Directory, the Active Directory server completed its configuration correctly. Otherwise, you can complete the procedures to unconfigure and reconfigure Active Directory as described in the ADK Installation topics in the Knowledge Center.

Client Configure the client into an existing Active Directory domain to complete Security Access Manager configuration. To ensure that the client system is part of the Active Directory domain, you can use System Properties to ensure the correct configuration of the client machine by selecting **Start** → **Settings** → **Control Panel** → **System** → **System Properties**.

In the Control Panel, double-click on System icon. The System Properties window is displayed. On the System Properties windows, click the Network Identification menu. If the Domain on the Network Identification contains the correct Active Directory domain, it indicates that the client machine is properly configured into the Active Directory domain.

Verifying version numbers

Server Active Directory is included with Windows Advanced Server installation. Therefore, only one version of this software is possible. After successfully configuring the Active Directory server, it is started automatically during the reboot process.

Client Active Directory is included with Windows Advanced Server installation. Therefore, only one version of this software is possible.

Confirming connectivity

Install Windows 2008 Support Tools from the `\support\tools` directory on the Windows 2008 operating system DVD. From that directory, run **`setup.exe`** and follow the installation guide to complete the installation.

Activate the ADSI Edit MMC window by selecting **Start** → **Programs** → **Windows 2008 Support Tools** → **Tools** → **ADSI Edit**.

The user can set up the server connection by selecting **Action** → **Settings**.

Additionally, click **Advanced** on the Connection window to input the administrator ID and password to connect to the remote Active Directory server. If the connection is successful, the client machine will be able to communicate with the Active Directory server using ADSI.

Verifying base systems

At a high level, you can determine which Security Access Manager servers are configured and which are running.

Verifying the policy server

The **pdadmin** command can be used to verify the correct operation of the policy server.

About this task

Enter the following command to log in as a Security Access Manager administrator:

```
pdadmin –a sec_master –p password
```

Assuming that WebSEAL is configured on the machine, at the **pdadmin** prompt, complete the following steps:

1. List the servers with the **server list** command. For this purpose, this command has the following syntax:
   ```
   pdadmin> server list webseald-machinename
   ```
2. List the objects with the **object list** command without options, as follows:
   ```
   pdadmin> object list
   ```
3. List ACLs with the **acl list** command without options, as follows:
   ```
   pdadmin> acl list
   ```
4. List users with the **user list** command. For this purpose, this command has the following syntax:
   ```
   pdadmin> user list name count
   ```

Verifying the authorization server

The Security Access Manager application development toolkit (ADK) includes the `authzn_demo` demonstration program. You can use this program, in remote mode, to validate the correct operation of the authorization server.

See the `README` file that accompanies this demonstration program for setup and execution instructions. The `README` file is in the following directory:

```
authzn-adk-install-dir/example/auth_demo/cpp
```

Verifying the runtime

The Security Access Manager Runtime can be installed on a system with only GSKit and Security Directory Server client.

In this case, the verification procedure for the Security Access Manager Runtime is the same as that described in “Verifying the policy server” on page 37.

Verifying Web security systems

You can verify whether Web security systems are operating properly by connecting from your browser to a URL.

Verifying WebSEAL

You can use a browser to verify that WebSEAL is operating properly.

To verify, enter the following URL into your browser:

```
https://webseal-machinename
```

Because a port number is not specified, it is assumed that WebSEAL is listening on port 443 (HTTPS).

Your browser might give you the following warnings:

1. The certificate received from this Web server is issued by a company that you have not yet chosen to trust
2. The name within the certificate received from WebSEAL does not match the name of the system from which it was received

If these warnings occur, they indicate that you did not yet purchase your own server certificate for your WebSEAL server. Your browser is complaining that it received a default server certificate from WebSEAL which contains default names for the issuing certificate authority and the name of the Web server.

Next, the browser prompts you to specify a Security Access Manager user name and password. Enter `sec_master` for the user name and the password that you configured for **sec_master** during installation. If authentication is successful, an image that is labeled Security Access Manager for WebSEAL displays.

Chapter 6. Collecting events to diagnose or audit server operations

You can collect events for diagnostic and auditing purposes of the servers.

Events for diagnostics and auditing are for operations of the Security Access Manager servers.

These events are not for installation of the servers.

To enable diagnostics and auditing, you define which types of events to capture. When events are captured, they can be written to:

- Log file
- Standard output (STDOUT) device
- Standard error (STDERR) device
- Combination of these destinations.

Beyond these destinations, when events are captured, they can be redirected to a remote server or redirected to an application for processing with log agents.

Diagnostic events

For collecting diagnostic information, define which *message events* and which *trace events* to capture. These events can help you troubleshoot problems.

A *message event* is a record of a noteworthy event that occurred, such as a failure to connect error message.

A *trace event* is a capture of information about the current operating environment at the time that a component or application failed to operate as intended.

Trace event logs provide IBM Support with information that relates to the condition of the system at the time a problem occurred. Trace logging enablement can cause large amounts of data to collect in a short amount of time. Therefore, enable trace events only at the direction of IBM Support.

To configure message or trace events, define *statements* in the routing file or Java properties file for the server:

- For message events, define the statements by severity level.
- For trace events, define the statements by trace level and, optionally, by component.

Auditing events

For auditing purposes, define audit, statistic, or other type of events for capture in a log file.

Auditing provides tracking and archiving of auditable events. These events create records of various server activities. Configure audit event logs by using r the base Security Access Manager.

Configure audit events by defining stanza entries in configuration files. Define stanza entries in configuration files:

- For base Security Access Manager auditing: define `logcfg` entries in the `[pdaudit-filter]` stanza.

For WebSEAL, you can also log HTTP events by using the `[logging]` stanza in the WebSEAL configuration file.

For more information about audit events, see the Auditing topics in the Knowledge Center.

Chapter 7. Customize logging events with tracing configuration files

Tracing configuration files are ASCII files that customize the logging of message and trace events for C language-based servers, daemons, and other C-language programs and applications.

To customize logging of message and trace events for Java applications, see Chapter 8, "Java properties files," on page 45.

At application startup, the application-specific tracing configuration file is read. With the tracing configuration file content, you can control the following aspects of logging event activity:

- Whether to enable logging for specific event classes.
- Where to direct the output for each event class.
- How many log files to use for each event class.
- How large each log file can be for each event class.

Location of tracing configuration files

You can locate tracing configuration files to control the logging of both message events and trace events.

Tracing configuration files can be edited through the appliance Web interface.

- Runtime Component

  From the main menu select **Secure Web Settings -> Manage -> Runtime Component**. In the Runtime Component panel select **Manage -> Configuration Files -> Tracing Configuration File**.
- WebSEAL servers

  From the main menu select **Secure Web Settings -> Manage -> Reverse Proxy**. In the Reverse Proxy panel select a reverse proxy instance, then select **Manage -> Configuration -> Edit Tracing Configuration File**.
- Authorization Server

  From the main menu select **Secure Web Settings -> Manage -> Authorization Server**. In the Authorization Server panel select an authorization server instance, then select **Manage -> Configuration -> Edit Tracing Configuration File**.

Tracing configuration file entries

Each tracing configuration file contains entries that control the logging of message events and trace events. However, the format of these entries differs by event type.

Use one of the following formats (entered on a single line without spaces) when you define entries in tracing configuration files:

Message events

```
severity:destination:location [[;destination:location]...]
[;GOESTO:{other_severity | other_component}]
```

Trace events

```
component:subcomponent.level[[,subcomponent.level]...]
```

```
:destination:location [[;destination:location]...]
[;GOESTO:{other_severity | other_component}]
```

Where:

***component*:*subcomponent.level*[[,*subcomponent.level*]...]**

Specifies the component, subcomponents, and reporting levels of trace events to log. For trace events only.

For the component portion, you can specify an asterisk (*) to log trace data for all components.

For the subcomponent portion, you can specify an asterisk (*) to log trace data for all subcomponents of the specified component.

For the level portion, specify the reporting level to log. This value is a number 1 - 9. A level of `1` indicates the least number of details, and a level of `9` indicates the greatest number of details.

destination

Specifies where to log the events. For each destination, you need to specify a location. When you specify multiple destination-location pairs, separate each pair with a semicolon (;). The following destinations are valid:

DISCARD

Discards the events.

FILE

Writes the events as ASCII text in the current code page and locale to the specified location. When you use this destination, you must specify a location for the file. Optionally, you can follow the `FILE` destination by a period and two numbers that are separated by a period (for example, `FILE.10.100`). The first value indicates the number of files to use. The second value indicates the number of events each file can contain. If you do not specify these values, there is only one log file that grows without limit.

The average size of an ASCII event is 200 bytes. Because the maximum size of a log file is 2 GB, limit the maximum number of events to approximately 10,000,000 events.

STDERR

Writes the events as ASCII text in the current code page and locale to the standard error device.

STDOUT

Writes the events as ASCII text in the current code page and locale to the standard output device.

TEXTFILE

Same as `FILE`.

UTF8FILE

Writes the events as UTF-8 text to the specified location. When you use this destination, you must specify a location for the file. Optionally, you can follow the `UTF8FILE` destination by a period and two numbers that are separated by a period (for example, `UTF8FILE.10.100`). The first value indicates the number of files to use. The second value indicates the number of events each file can contain. If you do not specify these values, there is only one log file that grows without limit.

The average size of a UTF-8 event is 200 bytes. Because the maximum size of a log file is 2 GB, limit the maximum number of events to approximately 10,000,000 events.

Note: When the operating system does not use a UTF-8 code page, the conversion to UTF-8 can result in data loss. When data loss occurs, the log file contains a series of question mark (?) characters at the location where the data conversion was problematic.

XMLFILE

Writes events to the specified location in the XML log format. When you use this destination, you must specify a location for the file. Optionally, you can follow the `XMLFILE` destination by a period and two numbers that are separated by a period (for example, `XMLFILE.10.100`). The first value indicates the number of files to use. The second value indicates the number of events each file can contain. If you do not specify these values, there is only one log file that grows without limit.

The average size of an XML message event is 650 bytes, and the average size of an XML trace event is 500 bytes. Because the maximum size of a log file is 2 GB, limit the maximum number of events to approximately 3,000,000 message events or 4,000,000 trace events.

XMLSTDERR

Writes events to the standard error device in the XML log format.

XMLSTDOUT

Writes events to the standard output device in the XML log format.

GOESTO:{*other_severity* | *other_component*}]

Specifies to route events to the same destination and location as either message events of the specified severity or trace events of the specified component.

location

Specifies the name and location of the log file. When the destination is `TEXT`, `TEXTFILE`, `UTF8FILE` or `XMLFILE`, you must specify a location. When the destination is `DISCARD`, `STDERR`, `STDOUT`, `XMLSTDERR` or `XMLSTDOUT`, you must specify a hyphen (`-`).

When you specify a fully qualified file name, you can use the `%ld` character string to insert the process ID into the file name.

When the number of files is specified as part of the destination, a period and the file number are appended to the specified log file.

Note: On Windows operating systems, the file name must not end with a period. If the file name ends with a period, when the file number is appended, the file name contains two consecutive periods. File names with two consecutive periods are not valid.

On AIX, Linux, and Solaris operating systems, the file name must be followed by file permissions, the user who owns the file, and the group that owns the file. Use the following format:

```
location:permissions:owner:group
```

To specify the location for message events from the policy server and write them to the default UTF-8 log file, you can specify the following location:

```
/var/PolicyDirector/log/msg__pdmgrd_utf8.log:644:ivmgr:ivmgr
```

severity

Specifies the severity of the message events to log. For message events only.

The following message severities are valid:

- `FATAL`
- `ERROR`
- `WARNING`
- `NOTICE`
- `NOTICE_VERBOSE`

You can specify an asterisk (*) to log messages regardless of severity.

For complete details about the severity of message events, see "Severity of message events" on page 51.

Chapter 8. Java properties files

Java properties files are ASCII files that are used to customize event logs for Java based Security Access Manager servers, daemons, and other Java-language programs and applications. Beyond customizing logs, these properties files are used to configure other aspects of the application.

The contents of the properties file enables the user to control the following aspects of message logs:

- Whether event logs are enabled
- Where to direct the output
- If the output is to a file, the number of files to use and the size of each file

Lines in the file that start with a number sign (#) are comments and do not affect logging.

The application name (*app_name*) is part of each logging property for a Java application. The application name is specified when you use the `com.tivoli.pd.jcfg.SvrSslCfg` command.

Location of Java properties files

The default locations for the Java properties files for Security Access Manager components are shown in Table 7.

Table 7. Location of Java properties files

Component	Default file name
Java application that is configured by using the `com.tivoli.pd.jcfg.SvrSslCfg` class.	The output application configuration file as specified in the `com.tivoli.pd.jcfg.SvrSslCfg` class.
Java based Security Access Manager commands, such as the **`pdjrtecfg`** command and `com.tivoli.pd.jcfg.SvrSslCfg` or applications not explicitly configured.	`$JAVA_HOME/PolicyDirector/PDJLog.properties`

Note: If the `com.tivoli.pd.jcfg.SvrSslCfg` command was not run, no application-specific configuration file exists. If there is no configuration file, the `PDJLog.properties` file is used.

Application-specific logging of Java applications

Configuration of message and trace logs for the IBM Security Access Manager Runtime for Java (AMJRTE) component is completed on a per-application basis. This configuration removes the file contention and ownership problems that are encountered in previous versions of Security Access Manager.

In addition to existing configuration properties, the application properties file created by the `com.tivoli.pd.jcfg.SvrSslCfg` command contains logging properties that are associated with the application-logging properties for the application.

The names of the logging objects and the log files in this configuration file contain the application server name that is supplied by the **appSvr** parameter of `com.tivoli.pd.jcfg.SvrSslCfg`. Thus, each application has a unique set of objects and log files. If the configuration file for an application is not being used (for instance, when the **pdjrtecfg** command is used), message log and tracing properties are taken from the existing `PDJLog.properties` file.

In addition, the size and the number of files that are used for messages and trace entries are now configurable.

Configuring message events with the Java properties file

To capture message events for Java applications, you need to configure the Java properties file.

Message loggers and file handlers

Each properties file contains properties for one or more message loggers. The `isLogging` property specifies whether message logs are enabled.

To turn on logging for a specific message logger, use:

```
baseGroup.PDJapp_nameMessageLogger.isLogging=true
```

To disable logging for a specific message logger, use:

```
baseGroup.PDJapp_nameMessageLogger.isLogging=false
```

Associated with each message logger is at least one file handler. A file handler specifies the destination for messages. After message logs are enabled by the message logger, the file handler properties are examined to determine whether to log messages, and if so, how and where. The properties that are associated with a file handler are:

```
baseGroup.PDJapp_nameFileHandler.fileName=
baseGroup.PDJapp_nameFileHandler.maxFileSize=
baseGroup.PDJapp_nameFileHandler.maxFiles=
```

where:

fileName
: Specifies the fully qualified file name to be used as the base name for message log files. The file can be in any location accessible by the Java application.

maxFileSize
: Specifies the maximum size, in kilobytes, of each message log file. Default is `512`.

maxFiles
: Specifies the maximum number of files to be used for message logs. Default is 3.

To specify what classes of messages to log, use the `MessageAllMaskFilter.mask` property as illustrated in Figure 1.

```
baseGroup.PDJapp_nameMessageAllMaskFilter.mask=FATAL
 | ERROR | WARNING | NOTICE | NOTICE_VERBOSE
```

Figure 1. Specifying what messages to log in a properties file

When PDJLog.properties is used

The $JAVA_HOME/PolicyDirector/PDJLog.properties file is used to define message and trace log properties in the following cases:

- For non-application-related Java commands, such as **pdjrtecfg** and com.tivoli.pd.jcfg.SvrSslCfg.
- If a Java application was not explicitly configured with the com.tivoli.pd.jcfg.SvrSslCfg command.
- If the application-specific properties file is inaccessible or does not exist.
- If a required property in the application-specific properties file is not found.

When you use the default PDJLog.properties file, message logs are enabled only for FATAL, ERROR, and WARNING messages. This is shown in the portion of the PDJLog.properties file in Figure 2. Logging can be enabled for NOTICE and NOTICE_VERBOSE messages by changing the isLogging property to true for the last two properties that are shown in Figure 2.

```
baseGroup.PDJMessageLogger.isLogging=true
baseGroup.PDJFatalFileHandler.isLogging=true
baseGroup.PDJErrorFileHandler.isLogging=true
baseGroup.PDJWarningFileHandler.isLogging=true
baseGroup.PDJNoticeFileHandler.isLogging=false
baseGroup.PDJNoticeVerboseFileHandler.isLogging=false
```

Figure 2. Portion of the default PDJLog.properties file

On a AIX, Linux, or Solaris operating system, to enable both NOTICE and NOTICE_VERBOSE messages, and to change the destination properties of NOTICE_VERBOSE messages, the following changes can be made, as indicated in **bold**:

```
baseGroup.PDJNoticeFileHandler.isLogging=true

baseGroup.PDJNoticeVerboseFileHandler.fileName=
/tmp/logs/msg__amjrte_verbose.log
baseGroup.PDJNoticeVerboseFileHandler.maxFileSize=1024
baseGroup.PDJNoticeVerboseFileHandler.maxFiles=4
baseGroup.
PDJNoticeVerboseFileHandler.isLogging=true
```

After you make these changes, NOTICE_VERBOSE messages are written to the /tmp/logs/msg__amjrte_verbose.log1 file. After that file reaches 1024 KB, the file is renamed /tmp/logs/msg__amjrte_verbose.log2 and logging continues with a new /tmp/logs/msg__amjrte_verbose.log1 log file. A maximum of four message log files is used.

(The procedure would be the same on a Windows operating system. The file name just needs to be changed to reflect a fully qualified file name on Windows operating systems.)

Console handler and console message logging

A console handler and a message console handler also are configured in the $JAVA_HOME/PolicyDirector/PDJLog.properties file. Both are disabled by default.

To send messages to the console, set both isLogging properties to true:

```
baseGroup.PDJConsoleHandler.isLogging=true
baseGroup.PDJMessageConsoleHandler.isLogging=true
```

Configuring trace events with the Java properties file

To capture trace events for Java applications, you need to configure the Java properties file.

Trace loggers and file handlers

Each properties file contains properties for one or more trace loggers. The `isLogging` property specifies whether trace logs are enabled.

To turn on tracing for a specific trace logger, use:

```
baseGroup.trace_logger_name.isLogging=true
```

To disable logging for a specific trace logger, use:

```
baseGroup.PDJapp_nameTraceLogger.isLogging=false
```

Associated with each trace logger is at least one file handler. A file handler specifies the destination for a specific class, or severity, of messages. After trace logs are enabled by the trace logger, the file handler properties are examined to determine whether to log traces, and if so, how and where. The properties for a file handler are:

```
baseGroup.PDJapp_nameTraceFileHandler.fileName=
baseGroup.PDJapp_nameTraceFileHandler.maxFileSize=
baseGroup.PDJapp_nameTraceFileHandler.maxFiles=
```

where:

fileName
Specifies the fully qualified file name to be used as the base name for trace log files. The file can be in any location accessible by the Java application.

maxFileSize
Specifies the maximum size, in KB, of each trace log file. Default is `512`.

maxFiles
Specifies the maximum number of files to be used for trace logs. Default is `3`.

Enabling trace in a Runtime for Java environment

Trace logging for components that use the Security Access Manager Runtime for Java environment is controlled through the application-specific properties file or, for applications that are not explicitly configured with the `com.tivoli.pd.jcfg.SvrSslCfg` command, the `$JAVA_HOME/PolicyDirector/PDJLog.properties` file.

To enable tracing:

```
baseGroup.PDJapp_nameTraceLogger.isLogging=true
```

For each trace logger, the properties file defines a mask attribute, `baseGroup.PDJapp_nameTraceAllMaskFilter.mask`, that determines what levels of tracing are enabled. Valid mask values are 1 through 9. The precise meaning of any specified mask value is unimportant. The general intention is that ascending from a lower mask value to a higher mask value (for example, 1 to 2) increases the level of information detail that is traced.

Setting the mask value to a particular level means that all tracing levels up to and including the specified level are enabled. For example, if the mask value is 4, then tracing levels 1, 2, 3, and 4 are traced.

Chapter 9. Message event logging

The contents of log files can be useful sources of information when you monitor or troubleshoot Security Access Manager servers.

You can use log files to capture any Security Access Manager message.

- Message logging for the C-language portions of Security Access Manager is controlled through *routing files*.
- Message logging for the Java language portions is controlled through *Java properties files*.

When relevant, the distinctions between these methods are mentioned.

Use the statements within routing files to control which messages to log, the location of the log files, and format of the messages. Use the information in this chapter to learn the configuration syntax that is used in the routing files and defines the default file name and location of the message log files. The directory location for message log files can be different, depending on whether Tivoli Common Directory is configured.

Severity of message events

In the message log file, each message event has an associated severity level.

The following message severities are valid:

- `FATAL`
- `ERROR`
- `WARNING`
- `NOTICE`
- `NOTICE_VERBOSE`

FATAL messages

An unrecoverable error occurred, such as a database corruption. The process that encounters the error usually terminates and might produce a core file.

This error might require manual intervention to recover or require special recovery actions. Depending on the nature of the failure, IBM Software Support might need to be consulted.

The identifier for these messages uses the error (`E`) message severity.

ERROR messages

An unexpected or nonterminal event, such as a timeout, or correctable event that requires manually intervention occurred.

The product continues to function, but some services or functions might not be available. This severity also indicates that a particular request or action was not completed. Administrative action might be required.

The identifier for these messages uses the error (`E`) message severity.

WARNING messages

An event occurred that is possibly not the wanted or requested result, such as a configuration file not being found and a default value used instead.

The program continues to function normally. This severity also indicates a condition that might be an error if the effects are unwanted or indicates a condition, which if not corrected, can result in an error. For example, a low memory or disk space condition.

The identifier for these messages uses the warning (`W`) message severity.

NOTICE messages

An event took place that does not directly require action, such as starting a server. The event conveys general information about running state or normal actions.

The identifier for these messages uses the information (`I`) message severity.

NOTICE_VERBOSE messages

This event is similar to `NOTICE`, but the events that are logged might contain more detailed information.

The identifier for these messages uses the information (`I`) message severity.

Names of message logs

The names of the message log files depend on whether the log is for the Runtime component or a server component.

The messages for the Runtime component are separated into severity-specific files, while the messages for the server components are written to the same file.

Names of runtime logs

IBM Security Access Manager runtime messages are messages that are produced by applications, commands, and utilities that use the Security Access Manager Runtime component. The sources include the C language-based utilities, such as the **pdadmin** commands and the **svrsslcfg** utility.

Table 8 lists the names of the default message log files for both C and Java language applications.

Table 8. Message severity levels and associated message logs

Message severity	Default log name
FATAL	**C runtime log name** msg__fatal.log **Java runtime log name** msg__*app_nameN*.log **WebSEAL log name** Written to the standard error file (STDERR)

Table 8. Message severity levels and associated message logs (continued)

Message severity	Default log name
ERROR	**C runtime log name** msg__error.log **Java runtime log name** msg__*app_nameN*.log **WebSEAL log name** Written to the standard error file (STDERR)
WARNING	**C runtime log name** msg__warning.log **Java runtime log name** msg__*app_nameN*.log **WebSEAL log name** Written to the standard error file (STDERR)
NOTICE	**C runtime log name** msg__notice.log **Java runtime log name** msg__*app_nameN*.log **WebSEAL log name** msg__notice_*PID*.log **Note:** Logging is not enabled by default.
NOTICE_VERBOSE	**C runtime log name** msg__verbose.log **Note:** Logging is not enabled by default. **Java runtime log name** msg__*app_nameN*.log **WebSEAL log name** msg__verbose_*PID*.log **Note:** Logging is not enabled by default.

Notes:

- When WebSEAL is running as a background process, FATAL, ERROR, and WARNING messages are redirected to the server message log file for that WebSEAL instance (msg__webseald–*instance_name*.log).
- If an application-specific configuration file does not exist for a Java application, message logs are controlled by the $JAVA_HOME/PolicyDirector/PDJLog.properties file. In these cases, messages are written to the following files:

  FATAL
  msg__amj_fatal*N*.log

  ERROR
  msg__amj_error*N*.log

  WARNING
  msg__amj_warning*N*.log

  NOTICE
  msg__amj_notice*N*.log

  NOTICE_VERBOSE
  msg__amj_verbose*N*.log

Note: By default, logging of NOTICE and NOTICE_VERBOSE messages is not enabled.

Based on the severity level, runtime messages from C-language applications are written to different log files. For example, WARNING messages are written to the msg__warning.log file and FATAL messages are written to msg__fatal.log file. Error messages from WebSEAL are written to STDERR, unless WebSEAL is running in the background. In this case, the messages are written to the WebSEAL server log file.

Runtime message log files that are associated with C-language applications are allowed to grow without bound. Periodically check the available disk space and adjust as necessary, perhaps by archiving or pruning the log files. You can change the name, location, and put size constraints on the runtime message log files, as explained in "Tracing configuration file entries" on page 41.

Runtime message log files for Java language applications can grow to a maximum size of 512 KB. A maximum of three message files can exist, with the most recent messages always being in the file that ends in "1". When the file reaches its maximum size, the files are renamed. For example, when the msg__*appname*1.log file reaches 512 KB, the following process occurs:

1. The msg__*appname*3.log file is deleted, if it exists
2. The msg__*appname*2.log file, if it exists, is renamed to msg__*appname*3.log
3. The msg__*appname*1.log file is renamed msg__*appname*2.log
4. A new msg__*appname*1.log file is created

The names, location, number, and size of the Java runtime logs can be changed, as explained in Table 7 on page 45.

Names of server logs

Server messages are messages that are generated by the daemons and servers that are associated with Security Access Manager.

Unlike runtime messages from C applications, which are written to different log files based on severity, server messages are always written to the message log for that particular server. Thus, all FATAL, WARNING, ERROR, NOTICE, and NOTICE_VERBOSE messages for the policy server are written to the msg__pdmgrd_utf8.log file. Similarly, WebSEAL messages are written to the msg__webseald-*instance_name*.log file.

Table 9 lists the default names for the server message log files.

Table 9. Message log files that are associated with servers

Server	Default message log file
Security Access Manager policy server	msg__pdmgrd_utf8.log
Security Access Manager authorization server	msg__pdacld_utf8.log
Security Access Manager WebSEAL	msg__webseald-*instance_name*.log
Security Access Manager Attribute Retrieval Service	msg__amwebars_exceptions.log

By default, the Security Access Manager server message log files (the ones that start with msg__) are allowed to grow without bound. Be sure to periodically check the available disk space and adjust as necessary. You might want to archive or prune the log files on a periodic basis as well.

You can change the name, location, and put size constraints on the server message log files as explained in “Tracing configuration file entries” on page 41.

Format of messages in logs

Messages in logs are recorded in a specific format.

To configure Security Access Manager message logs and trace logs to produce output in an XML log format, see “Sending messages to multiple places in different formats” on page 64 and “Trace logging in XML log format” on page 69.

Messages in text format

This log entry is a sample of a server message in text format.

Figure 3 shows an example of a message log entry.

```
2005-10-26-20:09:10.984-06:00I----- 0x1354A41E pdmgrd ERROR
ivc socket e:\am600\src\mts\mtsclient.cpp 1832 0x000001c4
HPDCO1054E  Could not connect to the server acld2 on port 7137.
```

Figure 3. Sample message log entry in text format

The following list explains the log entry fields that are shown in Figure 3:

2005-10-26-20:09:10.984-06:00I
Indicates the timestamp of the message entry. The timestamp is in the following format:

```
YYYY–MM–DD–hh:mm:ss.fff[+|–]hh:mmI
```

where:

YYYY-MM-DD
Specifies the date in year, month, and day.

hh:mm:ss.fff
Specifies the time in hours, minutes, seconds, and fractional seconds.

hh:mmI
Specifies the time inaccuracy factor.

0x1354A41E
Indicates the 32-bit message number in hexadecimal.

pdmgrd
Indicates the name of the process that created the entry.

ERROR
Indicates the severity of the message.

ivc Indicates the component for the process that generated the entry.

socket Indicates the subcomponent for the process that generated the entry.

e:\am600\src\mts\mtsclient.cpp
Indicates the name of the source file that generated the entry.

1832 Indicates the exact line number in the source file.

0x000001c4
Indicates the 32–bit thread ID in hexadecimal.

HPDCO1054E
Indicates the message ID.

Could not connect to the server acld2 on port 7137.
Indicates the message text.

Messages in XML log format

Figure 4 uses the message from Figure 3 on page 55, but shows the message in the XML log format.

```
<Message Id="HPDCO1054E" Severity="ERROR">
<Time Millis="1067220550984">2005-10-26-20:09:10.984</Time>
<Component>ivc/socket</Component>
<LogAttribs><KeyName><![CDATA[Message Number]]></KeyName>
<Value><![CDATA[0x1354A41E]]></Value>
</LogAttribs>
<Source
FileName="e:\am600\src\mts\mtsclient.cpp"
Method="unknown" Line="1832">
</Source>
<Process>pdmgrd</Process>
<Thread>0x000001c4</Thread>
<TranslationInfo
Type="XPG4" Catalog="pdbivc.cat" SetId="1" MsgKey="1354a41e">
<Param><![CDATA[acld2]]></Param>
<Param><![CDATA[7137]]></Param>
</TranslationInfo>
<LogText>
<![CDATA[HPDCO1054E Could not connect to the server acld2 on port 7137.]]
</LogText>
</Message>
```

Figure 4. Sample message entry in the XML log format

Environment variables

The message log behavior that is specified by a routing file can be changed by using environment variables.

The PD_SVC_ROUTING_FILE environment variable can specify a fully qualified file name for a routing file to replace the one currently in use. If the file is not accessible, or does not exist, no change in logging messages is made.

Routing for messages of a specific severity can be manipulated by using environment variables as well. Set the appropriate message log entry format string to the wanted environment variable:

- SVC_FATAL
- SVC_ERROR
- SVC_WARNING
- SVC_NOTICE
- SVC_NOTICE_VERBOSE

For example, on Windows operating systems, the following command overrides the setting in the corresponding routing file and directs `WARNING` messages to the standard error device and a file:

```
SET SVC_WARNING="STDERR:-;FILE:D:\MSGS\MSG__WARNING.LOG"
```

See "Tracing configuration file entries" on page 41 for a description of message log entry format strings.

Displaying and not displaying environment variables in the log

On the computer on which the Security Access Manager server is running, you can configure environment variables to display or not display in the server log.

Use the PD_SVC_DISPLAY_ENV_VARS or PD_SVC_DONT_DISPLAY_ENV_VARS environment variable to display or not display environment variables as required. Separate the environment variable keys by the "|" character.

PD_SVC_DISPLAY_ENV_VARS

Displays only the specified environment variables.

The following example shows how to set this environment variable on an AIX, Linux, or Solaris system:

```
export PD_SVC_DISPLAY_ENV_VARS="PATH|HOME|PWD"
```

The log displays only the PATH, HOME, and PWD environment variables display when the Security Access Manager server starts.

PD_SVC_DONT_DISPLAY_ENV_VARS

Displays all available environment variables except the specified environment variables.

The following example shows how to use this environment variable on an AIX, Linux, or Solaris system:

```
export PD_SVC_DONT_DISPLAY_ENV_VARS="LANG|PATH|PWD"
```

The log displays all environment variables except LANG, PATH, and PWD when the Security Access Manager server starts.

Notes:

- The PD_SVC_DISPLAY_ENV_VARS option takes precedence over the PD_SVC_DONT_DISPLAY_ENV_VARS option.
- If the value of one of the specified environment variables keys does not exist, Security Access Manager ignores that key.

Routing files for message events

Logging of message events is controlled by a routing file. Entries in the routing file for a server determine which message events to log.

The server configuration files pick up the information from the routing files. If for any reason a routing file is deleted, the log-file stanza entry for the appropriate server is used instead.

Within routing files, you can disable or enable any type of message logs by adding or removing the comment character (#) at the beginning of the line in the routing file.

C runtime routing file

AIX, Linux, or Solaris: default routing file

```
FATAL:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__fatal.log
:644:ivmgr:ivmgr
ERROR:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__error.log
:644:ivmgr:ivmgr
WARNING:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__warning.log
:644:ivmgr:ivmgr
NOTICE:FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__notice.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__verbose.log
:644:ivmgr:ivmgr
```

Windows: default routing file

```
FATAL:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__fatal.log
ERROR:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__error.log
WARNING:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__warning.log
NOTICE:FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__notice.log
#NOTICE_VERBOSE:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__verbose.log
```

Policy server pdmgrd_routing file

AIX, Linux, or Solaris: default routing file

```
FATAL:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdmgrd_utf8.log
:644:ivmgr:ivmgr
ERROR:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdmgrd_utf8.log
:644:ivmgr:ivmgr
WARNING:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdmgrd_utf8.log
:644:ivmgr:ivmgr
NOTICE:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdmgrd_utf8.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:STDOUT:-;UTF8FILE:/pd_install_dir/log/
msg__pdmgrd_utf8.log:644:ivmgr:ivmgr
```

where *pd_install_dir* is either `/var/PolicyDirector` or the Tivoli Common Directory location.

Windows: default routing file

```
FATAL:STDERR:-;FILE:pd_install_dir\log\msg__pdmgrd_utf8.log
ERROR:STDERR:-;FILE:pd_install_dir\log\msg__pdmgrd_utf8.log
WARNING:STDERR:-;FILE:pd_install_dir\mog\msg__pdmgrd_utf8.log
NOTICE:FILE:pd_install_dir\log\msg__pdmgrd_utf8.log
#NOTICE_VERBOSE:STDOUT:-;UTF8FILE:pd_install_dir\log\
msg__pdmgrd_utf8.log
```

where *pd_install_dir* is either the directory where Security Access Manager is installed or the Tivoli Common Directory.

Authorization server pdacld_routing file

AIX, Linux, or Solaris: default routing file

```
FATAL:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdacld_utf8.log
:644:ivmgr:ivmgr
ERROR:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdacld_utf8.log
:644:ivmgr:ivmgr
WARNING:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdacld_utf8.log
:644:ivmgr:ivmgr
```

```
NOTICE:STDOUT:-;UTF8FILE:/pd_install_dir/log/msg__pdacld_utf8.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:STDOUT:-;UTF8FILE:/pd_install_dir/log/
msg__pdacld_utf8.log:644:ivmgr:ivmgr
```

where *pd_install_dir* is either `/var/PolicyDirector` or the Tivoli Common Directory location.

Note: A log file name such as `msg__pdacld_utf8.log` applies to a default instance of the authorization server. If multiple instances of the authorization server exist, each log file contains the specified instance name. For example, if an authorization server instance name is `instance1`, the log file is called `msg__instance1-pdacld_utf8.log`.

Windows: default routing file

```
FATAL:STDOUT:-;UTF8FILE:pd_install_dir\log\msg__pdacld_utf8.log
ERROR:STDOUT:-;UTF8FILE:pd_install_dir\log\msg__pdacld_utf8.log
WARNING:STDOUT:-;UTF8FILE:pd_install_dir\log\msg__pdacld_utf8.log
NOTICE:STDOUT:-;UTF8FILE:pd_install_dir\log\msg__pdacld_utf8.log
#NOTICE_VERBOSE:STDOUT:-;UTF8FILE:pd_install_dir\log\
msg__pdacld_utf8.log
```

where *pd_install_dir* is either the directory where Security Access Manager is installed or the Tivoli Common Directory.

Note: A log file name such as `msg__pdacld_utf8.log` applies to a default instance of the authorization server. If multiple instances of the authorization server exist, each log file contains the specified instance name. For example, if an authorization server instance name is `instance1`, the log file is called `msg__instance1-pdacld_utf8.log`.

WebSEAL routing file

By default WebSEAL writes all messages to the standard error device.

When WebSEAL is running in the background, standard error is redirected to the `msg__webseald-`*`instance_name`*`.log` file.

AIX, Linux, or Solaris: default routing file

```
FATAL:STDERR:-
ERROR:STDERR:-
WARNING:STDERR:-
#NOTICE:FILE.10.100:pdweb_install_dir/log/msg__notice_%ld.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:FILE.10.100:pdweb_install_dir/log/msg__verbose_%ld.log
:644:ivmgr:ivmgr
```

where *pdweb_install_dir* is either `/var/pdweb` or the Tivoli Common Directory location.

Windows: default routing file

```
FATAL:STDERR:-
ERROR:STDERR:-
WARNING:STDERR:-
#NOTICE:FILE.10.100:pdweb_install_dir/log/msg__notice_%ld.log
#NOTICE_VERBOSE:FILE.10.100:pdweb_install_dir/log/msg__verbose_%ld.log
```

where *pdweb_install_dir* is the value of the PD_WEB environment variable. The PD_WEB environment variable is set the WebSEAL installation directory during the initialization of the WebSEAL runtime environment.

Message routing files

Logging of message events is controlled by a routing file. Entries in the routing file for a server determine which message events to log.

The server configuration files pick up the information from the routing files. If for any reason a routing file is deleted, the `log-file` stanza entry for the appropriate server is used instead.

Within routing files, you can disable or enable any type of message logs by adding or removing the comment character (#) at the beginning of the line in the routing file.

C runtime routing file on Windows

This topic describes the C runtime routing files on Windows.

The default routing file for the C runtime on a Windows operating system, `%PD_HOME%\etc\routing`, contains message log specifications similar to the specifications in Figure 5.

```
FATAL:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__fatal.log
ERROR:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__error.log
WARNING:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__warning.log
NOTICE:FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__notice.log
#NOTICE_VERBOSE:STDERR:-;FILE:C:/PROGRA~1/Tivoli/POLICY~1/log/msg__verbose.log
```

Figure 5. Sample C runtime routing file

Using this routing file, messages are logged in the following manner:

- `FATAL` messages are sent to the standard error device as ASCII text in the current code page and locale and also are written to the `C:/PROGRA~1/Tivoli/POLICY~1/log/msg__fatal.log` file in the same format.
- `ERROR` messages are sent to the standard error device as ASCII text in the current code page and locale and also are written to the `C:/PROGRA~1/Tivoli/POLICY~1/log/msg__error.log` file in the same format.
- `WARNING` messages are sent to the standard error device as ASCII text in the current code page and locale and also are written to the `C:/PROGRA~1/Tivoli/POLICY~1/log/msg__warning.log` file in the same format.
- `NOTICE` messages are written to the `C:/PROGRA~1/Tivoli/POLICY~1/log/msg__notice.log` file only.
- `NOTICE_VERBOSE` messages are not captured, because this line starts with the number sign (#).

Policy server routing file on AIX, Linux, or Solaris

This topic describes the policy server routing files on AIX, Linux, or Solaris.

The default routing file for the policy server on a AIX, Linux, or Solaris operating system, `/opt/PolicyDirector/etc/pdmgrd_routing`, contains message log specifications similar to the specifications in Figure 6 on page 61.

```
FATAL:STDOUT:-;UTF8FILE:/var/PolicyDirector/log/msg__pdmgrd_utf8.log:
644:ivmgr:ivmgr
ERROR:STDOUT:-;UTF8FILE:/var/PolicyDirector/log/msg__pdmgrd_utf8.log:
644:ivmgr:ivmgr
WARNING:STDOUT:-;UTF8FILE:/var/PolicyDirector/log/msg__pdmgrd_utf8.lo
g:644:ivmgr:ivmgr
NOTICE:STDOUT:-;UTF8FILE:/var/PolicyDirector/log/msg__pdmgrd_utf8.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:STDOUT:-;UTF8FILE:/var/PolicyDirector/log/msg__pdmgrd
_utf8.log:644:ivmgr:ivmgr
```

Figure 6. Sample policy server routing file

With this routing file, messages are logged as follows:

- FATAL messages are written to the standard output device as ASCII text in the current code page and locale and written as UTF-8 text to the /var/PolicyDirector/log/msg__pdmgrd_utf8.log file. When the file is initially created, user ivmgr is the owner, ivmgr is the group, and 644 is the file permission.
- ERROR messages are written to the standard output device as ASCII text in the current code page and locale and written as UTF-8 text to the /var/PolicyDirector/log/msg__pdmgrd_utf8.log file. When the file is initially created, user ivmgr is the owner, ivmgr is the group, and 644 is the file permission.
- WARNING messages are written to the standard output device as ASCII text in the current code page and locale and written as UTF-8 text to the /var/PolicyDirector/log/msg__pdmgrd_utf8.log file. When the file is initially created, user ivmgr is the owner, ivmgr is the group, and 644 is the file permission.
- NOTICE messages are written to the standard output device as ASCII text in the current code page and locale and written as UTF-8 text to the /var/PolicyDirector/log/msg__pdmgrd_utf8.log file. When the file is initially created, user ivmgr is the owner, ivmgr is the group, and 644 is the file permission.
- NOTICE_VERBOSE messages are not written, because this line starts with the number sign (#).

WebSEAL routing file

This topic describes the details and specifications of a WebSEAL routing file.

The default routing file for a WebSEAL server contains message log specifications similar to the specifications in Figure 7.

```
FATAL:STDERR:-
ERROR:STDERR:-
WARNING:STDERR:-
#NOTICE:FILE.10.100:/var/pdweb/log/msg__notice_%ld.log
:644:ivmgr:ivmgr
#NOTICE_VERBOSE:FILE.10.100:/var/pdweb/log/msg__verbose_%ld.log
:644:ivmgr:ivmgr
```

Figure 7. Sample WebSEAL routing file

By removing the number sign (#) from the `NOTICE` specification and stopping and then restarting the WebSEAL server, `NOTICE` messages are written to a set of 10 files. Assuming that the process ID of the WebSEAL server is 1017, the names of the 10 files would be:

```
/var/pdweb/log/msg__notice_1017.log.1
/var/pdweb/log/msg__notice_1017.log.2
/var/pdweb/log/msg__notice_1017.log.3
/var/pdweb/log/msg__notice_1017.log.4
/var/pdweb/log/msg__notice_1017.log.5
/var/pdweb/log/msg__notice_1017.log.6
/var/pdweb/log/msg__notice_1017.log.7
/var/pdweb/log/msg__notice_1017.log.8
/var/pdweb/log/msg__notice_1017.log.9
/var/pdweb/log/msg__notice_1017.log.10
```

Message logging starts with the first file, `/var/pdweb/log/msg__notice_1017.log.1`. After 100 `NOTICE` messages are logged to that file, messages are written to the next file, `/var/pdweb/log/msg__notice_1017.log.2`. Message logging continues in this manner until 100 messages are written to the `/var/pdweb/log/msg__notice_1017.log.10` file. At that point, the messages in the first file are deleted, and logging resumes again to the `/var/pdweb/log/msg__notice_1017.log.1` file.

Note: By default WebSEAL writes all messages to the standard error device. When WebSEAL is running in the background, standard error is redirected to the `msg__webseald-`*`instance_name`*`.log` file.

Limiting the size of message logs

By default, message logs grow without limit. Limiting message log size requires that the directories and file systems that contain message log files must be checked on a periodic basis to ensure that enough space is available, and to prune the log files or make more space available as necessary.

The routing files can be modified to limit the amount of disk space that is used for message logs.

Consider the routing specification that is shown in Figure 8.

```
FATAL:STDOUT:-;UTF8FILE.10.100:/var/PolicyDirector/log/msg__pd
mgrd_fatal_utf8.log:644:ivmgr:ivmgr
ERROR:STDOUT:-;UTF8FILE.10.100:/var/PolicyDirector/log/msg__pd
mgrd__error_utf8.log:644:ivmgr:ivmgr
WARNING:STDOUT:-;UTF8FILE.5.1000:/var/PolicyDirector/log/msg__
pdmgrd_warning_utf8.log:644:ivmgr:ivmgr
NOTICE:STDOUT:-;UTF8FILE.5.1000:/var/PolicyDirector/log/msg__p
dmgrd_notice_utf8.log:644:ivmgr:ivmgr
NOTICE_VERBOSE:STDOUT:-;XMLFILE.10.500:/var/PolicyDirector/log
/msg__pdmgrd_verbose_utf8.xml:644:ivmgr:ivmgr
```

Figure 8. Multiple log files

All message files that are produced by this message log specification are of a determinate size, thus the maximum disk space that can be used by all of the files can be calculated.

Estimating the size of message logs

Each entry in a message log file with a destination of `FILE`, `TEXTFILE`, or `UTF8FILE` is an average of 200 bytes in size. The maximum size of any log file is 2 GB.

To estimate the size of all the log files, in bytes, use the following equation:

```
200 × (Number of log files) × (Number of entries per log file)
```

For example, given the following specification:

```
NOTICE:UTF8FILE.10.1000:E:\LOGS\PDPROXYMGRD.LOG
```

The maximum size for the `PDPROXYMGRD.LOG` file would be approximately (`200 × 10 × 1000`) or 2,000,000 bytes.

Each entry in a message log file with a destination of `XMLFILE` is an average of 650 bytes in size. Therefore, the maximum size of a log file, in bytes, can be estimated using the following equation:

```
650 × (Number of log files) × (Number of entries per log file)
```

For example, given the following specification:

```
NOTICE:XMLFILE.10.500:E:\LOGS\MSG__NOTICE.XML
```

The maximum size for the `MSG__NOTICE.XML` file would be approximately (650 × 10 × 500) or 3,250,000 bytes.

Logging all messages the same way

You can log messages to a single file.

To send all runtime messages to a single file regardless of severity, the routing specification that is shown in Figure 9 can be used in the `/opt/PolicyDirector/etc/routing` file.

```
*:UTF8FILE:/tmp/msg__amrte_utf8.log:666:ivmgr:ivmgr
```

Figure 9. Sending all messages to one location

Changing the message format in log files

You can change the message format for a server-specific message log file.

About this task

To change the message format for a server-specific message log file, complete the following steps:

Procedure

1. Go to the directory where the routing file is located. The default location for the Security Access Manager servers is one of the following operating system-specific locations:

   AIX, Linux, and Solaris operating systems

   `/opt/PolicyDirector/etc/`

Windows operating systems
  `%PD_HOME%\etc\`

The default location for a WebSEAL server is one of the following operating system-specific locations:

AIX, Linux, and Solaris operating systems
  `/opt/pdweb/etc/`

Windows operating systems
  `%PD_WEB%\etc\`

2. Edit the appropriate server-related routing file. The following list contains the names of the routing files:

   `pdmgrd_routing`
   The routing file for the policy server

   `pdacld_routing`
   The routing file for the authorization server

   `routing`
   The routing file that is used for the C runtime

   `routing`
   The routing file for the WebSEAL server

   Note: Although the routing file for the C runtime and the WebSEAL server have the same file name, these files are in different directories.

3. Find the statement that defines how to log messages of a specific severity. For example, find the `ERROR` statement in the routing file to change the logging of error messages. The `ERROR` statement might be similar the following statement:

   ```
   ERROR:STDOUT:-;UTF8FILE:pd_install_dir/log/msg__pdmgrd_utf8.log
   :644:ivmgr:ivmgr
   ```

   where *pd_install_dir* is either `/var/PolicyDirector` or the directory that is defined by the `tivoli_common_dir` stanza entry of the `pd.conf` configuration file (`/var/ibm/tivoli/common`).

   For example, to log error messages in the XML format and to log these messages to the `msg__error.log` file, make the following changes:

   a. Change `UTF8FILE` to `XMLFILE`
   b. Change `msg__pdmgrd_utf8.log` to `msg__error.log`

   The following statement is the result of these changes:

   ```
   ERROR:STDOUT:-;XMLFILE:pd_install_dir/log/msg__error.log
   :644:ivmgr:ivmgr
   ```

4. Save and exit the routing file.

Results

After this statement is changed, `ERROR` messages are written to the standard output device in ASCII text and written to the `msg__error.log` file in the XML format.

Sending messages to multiple places in different formats

You can send messages to multiple locations with different message formats.

To send `FATAL`, `ERROR`, and `WARNING` messages from the authorization server to the standard output device as ASCII text, to a file as UTF-8 text, and to another file in XML log format, the specification that is shown in Figure 10 on page 65 can be used in the %PD_HOME%\etc\msg__pdacld.routing file.

```
FATAL:STDOUT:-;UTF8FILE:C:\LOGS\MSG__PDACLD_%LD.LOG;XMLFILE:C:\XMLLO
GS\MSG__PDACLD_%LD.XML
ERROR:STDOUT:-;UTF8FILE:C:\LOGS\MSG__PDACLD_%LD.LOG;XMLFILE:C:\XMLLO
GS\MSG__PDACLD_%LD.XML
WARNING:STDOUT:-;UTF8FILE:C:\LOGS\MSG__PDACLD_%LD.LOG;XMLFILE:C:\XML
LOGS\MSG__PDACLD_%LD.XML
NOTICE:UTF8FILE:C:\LOGS\MSG__PDACLD_%LD.LOG;XMLFILE:C:\XMLLOGS\MSG__
PDACLD_%LD.XML
NOTICE_VERBOSE:UTF8FILE:C:\LOGS\MSG__PDACLD_%LD.LOG;XMLFILE:C:\XMLLO
GS\MSG__PDACLD_%LD.XML
```

Figure 10. Sending messages to standard output to file in UTF-8, and to file in XML format

After you stop and restart the authorization server to pick up the changes to the routing file and assuming that the process ID of the authorization server is 1253, the following files are created to contain messages:

- C:\PDACLD_LOGS\MSG__PDACLD_1253.LOG
- C:\PDACLD_XMLLOGS\MSG__PDACLD_1253.XML

Chapter 10. Trace event logs

Security Access Manager provides configurable tracing capabilities that can aid in problem determination. Unlike message logs, trace logs (or tracing) are *not* enabled by default.

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 a Security Access Manager environment.

Attention: 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.

Trace logs are best suited to situations where a problem is easily reproduced, is short-lived in duration, and can be produced without significant trace generation from other system activity. Enabling tracing can adversely affect the performance of Security Access Manager and its associated products and applications.

Tracing can be activated when servers, daemons, and applications start by using routing files and Java properties files. In some cases, tracing can be activated dynamically by using the **`server task trace`** command with the **set** option.

Note: Tracing from the C-language portions of Security Access Manager is controlled through routing files. Similarly, tracing from the Java language portions of Security Access Manager is controlled through Java properties files. When relevant, the distinction between these two methods of trace handling is mentioned.

Mechanisms for controlling trace logs

There are different trace log techniques that are used, depending on which component is being logged.

Tracing can be controlled by using the following mechanisms:

routing file

A routing file can be used to control tracing of the policy server, authorization server, WebSEAL, and runtime components. An affected component must be stopped and restarted for modifications to the routing file to take effect.

Java properties file

A Java properties file can be used to control tracing of the IBM Security Access Manager Runtime for Java components.

trace command

The **`server task trace`** command can be used to dynamically control trace operations for the authorization server and WebSEAL components. This command can also be used to control trace operations for custom C applications that were developed by using the Security Access Manager authorization C APIs.

Managing Trace

Tracing can be activated either through a routing file or through a Security Access Manager server task administrative command.

The following information outlines details of the second method. For details on how to activate tracing through a routing file, see Chapter 7, “Customize logging events with tracing configuration files,” on page 41.

The **server task trace** command can be used to dynamically control trace operations. As with other Security Access Manager administrative functions, the trace command can be issued through either the **pdadmin** utility or programmatically through the Security Access Manager Administrative API.

Different **pdadmin** commands are available to:
- List all of the available trace points.
- Change the level and destination for specific trace points.
- Retrieve the trace point level for specific trace points.

Listing all trace commands

To list all of the trace components that are offered by a server, issue the trace list command:

```
server task server name trace list
```

Where *server name* specifies the name of the server on which you want to collect trace information.

```
pdadmin> server task PDWebPI-webpi.gc.au.ibm.com trace list
pdwebpi.request
pdwebpi.response
...
```

Adjusting the trace level of a component

To change the level and destination for a specific trace point, use the following command:

```
server task <server name> trace set <component> \
<level> [file path=file|other-log-agent-config]
```

Where:

server name
Specifies the name of the server on which you want to collect trace information.

component
Specifies the name of trace component as shown by the list command.

level Controls the amount of detail to be gathered, in the range of 1 to 9, with 1 collecting the least number of traces, and 9 collecting the most number of traces.

file path
The optional **file path** parameter specifies the location for trace output. If this parameter is not supplied, the trace output is sent to the stdout stream of the server.

The following example sets the trace level to 9 for the **pdwebpi.request** component. Any output that is generated is sent to the /tmp/log.txt file on the WebPI server.

```
pdadmin> server task PDWebPI-wpi.com trace set pdwebpi.request 9
file path=/tmp/log.txt
```

Retrieving the current trace level of a component

To show the names and levels for all enabled trace components, use the following command:

```
server task server-name trace show [component]
```

If the optional **component** parameter is omitted, the output lists the name and level of all of the enabled trace components.

```
pdadmin>server task PDWebPI-wpi.ibm.com trace show pdwebpi.request 9
```

Tracing configuration file examples

To illustrate features available with the tracing configuration files, several examples are provided.

Trace logging in XML log format

To send trace output for the authorization server to a single file in XML log format, the *install_dir*/etc/msg__pdacld.routing file can be modified as follows:

```
*:*.9:XMLFILE:E:\PDACLD_XMLTRACE\TRACE__PDACLD_%LD.XML
```

After you stop and restart the authorization server to pick up the routing file change, and assuming that the process ID of the authorization server is 412, trace output is written to the following file:

```
E:\PDACLD_XMLTRACE\TRACE__PDACLD_412.XML
```

Trace logging to multiple files

By default, trace logs grow without limit. This requires that the directories and file systems that contain trace log files be checked on a periodic basis to ensure that enough space is available, and to prune the log files or make more space available as necessary.

The routing files can be modified to limit the amount of disk space that is used for trace logs.

To send runtime trace output, from reporting levels 1 through 5, to 10 different files, each containing a maximum of 10000 trace entries, the routing specification that is shown in Figure 11 can be used in the /opt/PolicyDirector/etc/routing file.

```
*:*.5:UTF8FILE.10.10000:/tmp/trace__am_utf8.log:666:ivmgr:ivmgr
```

Figure 11. Sending trace output to multiple files

Tracing starts with the first file, /tmp/trace__am_utf8.log.1. After 10000 trace entries are logged to that file, trace entries are written to the second file, /tmp/trace__am_utf8.log.2. Tracing continues in this manner until 10000 trace entries are written to the /tmp/trace__am_utf8.log.10 file. At that point, the trace entries in the first file are deleted, and tracing resumes again to the /tmp/trace__am_utf8.log.1 file.

Tracing a particular component

At the direction of IBM Software Support, you might be asked to enable tracing for a particular component of Security Access Manager.

For example, if asked to trace the `mgr` component of the policy server, the `/opt/PolicyDirector/etc/pdmgrd_routing` file can be modified as shown in Figure 12.

```
mgr:*.9:UTF8FILE:/tmp/trace__pdmgrd_mgr_9_utf8.log:644:ivmgr:ivmgr
```

Figure 12. Tracing a component

Determining maximum size of a trace log

Each entry that is made to a trace log file created that uses a destination of `FILE`, `TEXTFILE`, or `UTF8FILE` is an average of 200 bytes in size.

The maximum size of a log file, in bytes, can be estimated as follows:

```
200 × (Number of log files) × (Number of entries per log file)
```

For example, given a specification of:

```
*:*.9:TEXTFILE.10.10000:C:/PROGRA~1/Tivoli/POLICY~1/log/trace__%ld.log
```

The maximum size would be approximately (`200 × 10 × 10000`) or 20,000,000 bytes.

Trace entries that are written in XML log format are an average of 500 bytes in size, thus for a specification of:

```
*:*.9:XMLFILE.10.10000:/var/dbug20031028A/trace__%ld.xml
```

The maximum size would be approximately (`500 × 10 × 10000`) or 50,000,000 bytes.

Enabling trace

You can enable tracing during startup and be able to view trace records.

About this task

To enable tracing during startup and be able to view trace records, complete the following steps:

Procedure

1. Edit the appropriate routing file for the server. The following list contains the names of the routing files

   `pdmgrd_routing`
   The routing file for the Security Access Manager policy server

   `pdacld_routing`
   The routing file for the Security Access Manager authorization server
2. Add a line similar to the following to the routing file:
   ```
   *:*.9:TEXTFILE:pd_install_dir/log/trace_%ld.log
   ```

where, on a Windows operating system, *pd_install_dir* is either the directory where Security Access Manager is installed or the Tivoli Common Directory location.

Or, remove the number sign (#) at the beginning of this line, if it exists in the routing file, to allow viewing of trace records.

3. Change this line, if you want to log in this debug trace data XML log format. For example, you can change the line to send the output to an XML file instead of a text file:

   ```
   *:*.9:XMLFILE.10.1000:pd_install_dir/log/trace__%ld.log;XMLSTDERR:-
   ```

   where, on a AIX, Linux, or Solaris operating system, *pd_install_dir* is either `/var/PolicyDirector` or the Tivoli Common Directory location.

Using the trace commands

Use the **server tasks trace** command that is provided as by the **pdadmin** utility to manage trace components.

You can use the **trace** command to complete the following operations:

trace list
List all available trace components

trace set
Enable the trace level and trace message destination for a component and its subordinates

trace show
Show the name and level for all enabled trace components or for the specified component

For more information about the **server task trace** command, see "server task trace" on page 116.

Listing available trace components

You can list the specified component or all components that are available to gather and report trace information.

To list the specified component or all components available to gather and report trace information, run the following command:

```
trace list [component]
```

The trace components themselves are organized in a hierarchical fashion. If trace is activated for a parent trace component, it will automatically be activated for all children trace components. As an example, if you activate trace for the component: *pdweb.snoop*, tracing for the sub-component, *pdweb.snoop.jct* will also be activated.

Enabling trace

Use the **server task trace set** command to enable the gathering of trace information for the specified component and level.

```
trace set component level [file path=file | log_agent]
```

where:

component
The trace component name. This required argument indicates the component to be enabled. WebSEAL components are prefixed with `pdweb`.

level Reporting level. This required argument must be in the range of 1 to 9. The *level* argument specifies the number of details that are gathered by the **trace** command. Level 1 indicates the least detailed output, and level 9 indicates the most detailed output.

file The fully qualified name of the file to which trace data is written.

log_agent
Optionally specifies a destination for the trace information that is gathered for the specified component. See the event log information in the Administering topics in the IBM Knowledge Center.

Showing enabled trace components

You can list all enabled trace components or a specific enabled component.

If a specified component is not enabled, no output is displayed.

```
trace show [component]
```

Example:

```
pdadmin> server task webseald-instance trace set pdweb.debug 2
pdadmin> server task webseald-instance trace show pdweb.debug 2
```

Changing the name and location of trace files

Trace log file locations and names depends on which Security Access Manager component is being traced.

For the Security Access Manager authorization server, the trace log file can be explicitly specified by the user in the following command:

```
pdadmin> server task server-name trace set component level [file path=file]
```

where *server-name* is the name of the authorization server that is displayed by the **server list** command, and *file* is the fully qualified trace file name.

Alternatively, if the Security Access Manager routing file is being used to enable and disable tracing, then the location and name of this trace log file can be defined within the routing file.

For WebSEAL, the trace log file can be explicitly specified by the user in the following command:

```
pdadmin> server task server-name trace set component level [file path=file]
```

where *server-name* is the name of the WebSEAL server that is displayed by the **server list** command, and *file* is the fully qualified trace file name.

Alternatively, if the WebSEAL routing file is being used to enable and disable tracing, then the location and name of this trace log file can be defined within the routing file.

Format of trace entry in logs

Trace entries are recorded in a specific format.

Figure 13 shows an example of a trace entry that is taken from a Security Access Manager trace log file.

```
2005-10-29-18:01:06.984-06:00I—–- pdmgrd DEBUG8 mgr general
e:\am600\src\ivmgrd\pdmgrapi\MrMgmtDomainMan.cpp 736 0x000007d0
CII ENTRY: MrMgmtDomainMan::setCurrentDomainName
```

Figure 13. Sample trace log entry in text format

The following list describes the trace entry that is shown in Figure 13.

2005-10-29-18:01:06.984-06:00I
Indicates the timestamp of the trace entry. Timestamp is in the following format: `YYYY–MM–DD–hh:mm:ss.fff[+|–]hh:mmI`

where:

YYYY-MM-DD
Specifies the date in year, month, and day.

hh:mm:ss.fff
Specifies the time in hours, minutes, seconds, and fractional seconds.

hh:mmI
Specifies the time inaccuracy factor

pdmgrd
Indicates the name of the process which created the entry.

DEBUG8
Indicates the reporting level of the trace entry.

mgr Indicates the component for the process that generated the entry.

general
Indicates the subcomponent for the process that generated the entry.

e:\am600\src\ivmgrd\pdmgrapi\MrMgmtDomainMan.cpp
Indicates the name of the product source file that generated the entry.

736 Indicates the exact line number in source file.

0x000007d0
Indicates the thread ID in hexadecimal.

CII ENTRY: MrMgmtDomainMan:setCurrentDomainName
Indicates the text of the trace entry.

Trace logging for WebSEAL

This section describes some of the options available for trace logs for WebSEAL.

WebSEAL provides the following components to trace HTTP requests:

- pdweb.debug
- pdweb.snoop

Note: The amount of data that is produced by the trace options, especially by the snoop trace command, can be large.

pd.ivc.ira trace for LDAP server interaction

The **pd.ivc.ira** component traces Security Access Manager interaction with the LDAP server.

The trace helps with determining problems that occur during authentication. This trace can show the following information:

- The LDAP search path that is used during the search for a user
- Whether authentication succeeded for the user
- Whether any policy (for example, password, time-of-day) took effect

This information helps to identify Security Access Manager problems that originate from the LDAP server. The trace also shows the general interaction with the local user registry cache.

If the trace level is set to 7, approximately 30 lines of trace are produced for every transaction. This trace level mostly shows the authentication process. It can be used to determine whether a DN for the user was successfully located, and whether authentication for the user succeeded.

If the trace level is set to 8, approximately 170 lines of trace are produced for every transaction. In addition to the authentication process, this trace level logs the steps that are involved in validating the user policy. It also shows some interaction with the local user registry cache.

The following sample output is an extract of the trace that is produced during a standard authentication for a trace level of 8. The output shows that the user, `scotte`, was successfully authenticated and that the DN of the user is `cn=Scott Exton,o=ibm,c=au`.

Example pd.ivc.ira output (extract)

```
...
2007-03-09-14:31:00.329+10:00I----- thread(2) trace.pd.ivc.ira:8 /project/am610/build/am61
0/src/ivrgy/ira_auth.c:1221: CII ENTRY: ira_get_dn_utf8() parm: scotte
...
2007-03-09-14:31:00.329+10:00I----- thread(2) trace.pd.ivc.ira:7 /project/am610/build/am61
0/src/ivrgy/ira_entry.c:2879: ira_ldap_search_ext_s() base: SECAUTHORITY=DEFAULT scope: 2
filter: (secDomainId=Default%scotte)
2007-03-09-14:31:00.329+10:00I----- thread(2) trace.pd.ivc.ira:7 /project/am610/build/am61
0/src/ivrgy/ira_ldap.c:3009: ira_ldap_search_ext_s(): No timeout - calling ldap_search_ext
_s
2007-03-09-14:31:00.331+10:00I----- thread(2) trace.pd.ivc.ira:7 /project/am610/build/am61
0/src/ivrgy/ira_ldap.c:3029: ira_ldap_search_ext_s: Returning LDAP rc x0
...
2007-03-09-14:31:00.332+10:00I----- thread(2) trace.pd.ivc.ira:8 /project/am610/build/am61
0/src/ivrgy/ira_auth.c:1738: CII ENTRY: ira_authenticate_user3() parm: cn=Scott Exton,o=ib
m,c=au
...
2007-03-09-14:31:00.334+10:00I----- thread(2) trace.pd.ivc.ira:8 /project/am610/build/am61
0/src/ivrgy/ira_auth.c:1596: CII EXIT ira_auth_passwd_compare() with status:  0x00000000
...
2007-03-09-14:31:00.334+10:00I----- thread(2) trace.pd.ivc.ira:8 /project/am610/build/am61
0/src/ivrgy/ira_cache.c:1588: CII EXIT ira_cache_user_get_account_state() with status:  0x
00000000
...
2007-03-09-14:31:00.340+10:00I----- thread(2) trace.pd.ivc.ira:8 /project/am610/build/am61
0/src/ivrgy/ira_auth.c:2160: CII EXIT ira_authenticate_user3() with status:  0x00000000
```

pdweb.debug trace of HTTP header requests and responses

The pdweb.debug component traces the HTTP headers for requests and responses.

The pdweb.debug component operates at level 2 only. To log the message body, see "pdweb.snoop trace of HTTP traffic with WebSEAL" on page 75.

The following example command starts the trace utility for the pdweb.debug component at level 2 and directs the output to a file:

```
pdadmin> server task webseald-instance trace set pdweb.debug 2 \
file path=debug.log
```

Sample output of this command as it displays in the `debug.log` file:

```
2012-08-10-23:42:19.725+00:00I----- thread(7) trace.pdweb.debug:2
/amweb600/src/wand/wand/log.c:278: ------------- Browser ===> PD -------------
Thread_ID:27
GET /junction/footer.gif HTTP/1.1
Accept: */*
Referer: https://bevan/junction/
Accept-Language: en-us
Accept-Encoding: gzip, deflate
If-Modified-Since: Wed, 11 Jul 2009 21:11:14 GMT
If-None-Match: "abe09-3c8-3b4cc0f2"
User-Agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; .NET CLR 1.0.3705)
Host: bevan
Connection: Keep-Alive
---------------------------------------------------

2012-08-10-23:42:19.736+00:00I----- thread(7) trace.pdweb.debug:2
/amweb600/src/wand/wand/log.c:278: ------------- PD ===> BackEnd -------------
Thread_ID:27
GET /footer.gif HTTP/1.1
via: HTTP/1.1 bevan:443
host: blade.cruz.ibm.com:444
user-agent: Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 5.1; .NET CLR 1.0.3705)
accept: */*
accept-language: en-us
accept-encoding: gzip, deflate
if-none-match: "abe09-3c8-3b4cc0f2"
referer: https://bevan/blade/
if-modified-since: Wed, 11 Jul 2009 21:11:14 GMT
connection: close
---------------------------------------------------

2012-08-10-23:42:19.739+00:00I----- thread(7) trace.pdweb.debug:2
/amweb600/src/wand/wand/log.c:278: ------------- PD <=== BackEnd -------------
Thread_ID:27
HTTP/1.1 304 Not Modified
date: Wed, 10 Aug 2012 23:34:17 GMT
etag: "abe09-3c8-3b4cc0f2"
server: IBM_HTTP_SERVER/1.3.19.1Apache/1.3.20 (Unix)
connection: close
---------------------------------------------------

2012-08-10-23:42:19.740+00:00I----- thread(7) trace.pdweb.debug:2
/amweb600/src/wand/wand/log.c:278: ------------- Browser <=== PD ------------
Thread_ID:27
HTTP/1.1 304 Not Modified
date: Wed, 10 Aug 2012 23:34:17 GMT
etag: "abe09-3c8-3b4cc0f2"
server: IBM_HTTP_SERVER/1.3.19.1Apache/1.3.20 (Unix)
---------------------------------------------------
```

pdweb.snoop trace of HTTP traffic with WebSEAL

The pdweb.snoop component traces HTTP traffic. This component logs the HTTP headers and the message body for requests and responses.

Note: The snoop component traces the entire request and response as it is read off the socket. This trace might contain sensitive information.

The pdweb.snoop component has the following subcomponents:

pdweb.snoop.client
   Traces data that is sent between WebSEAL and clients.

pdweb.snoop.jct

   Traces data that is sent between WebSEAL and junctions.

To trace only the message headers, see "pdweb.debug trace of HTTP header requests and responses" on page 74.

The following example command starts the trace utility for the pdweb.snoop component at level 9 and directs the output to a file:

```
pdadmin> server task webseald-instance trace set pdweb.snoop 9 \
file path=snoop.out
```

The following sample output shows the WebSEAL server that is sending 2137 bytes of data to a client at IP address 10.4.5.12:

```
----------------------------------------
2012-08-10-19:35:18.541+00:00I----- thread(5) trace.pdweb.snoop.client:1
/home/amweb600/src/pdwebrte/webcore/amw_snoop.cpp:159:
----------------------------------------
Thread 2828; fd 22; local 10.4.5.10:443; remote 10.4.5.12:1250
Sending 2137 bytes
0x0000   4854 5450 2f31 2e31 2034 3033 2046 6f72        HTTP/1.1.403.For
0x0010   6269 6464 656e 0d0a 6461 7465 3a20 5475        bidden..date:.We
0x0020   652c 2032 3820 4f63 7420 3230 3033 2031        d,.10.Aug.2005.1
0x0030   393a 3335 3a31 3820 474d 540d 0a73 6572        9:35:18.GMT..ser
0x0040   7665 723a 2057 6562 5345 414c 2f35 2e31        ver:.WebSEAL/5.1
0x0050   2e30 2e30 2028 4275 696c 6420 3033 3130        .0.0.(Build.0310
0x0060   3230 290d 0a63 6163 6865 2d63 6f6e 7472        20)..cache-contr
0x0070   6f6c 3a20 6e6f 2d63 6163 6865 0d0a 7072        ol:.no-cache..pr
0x0080   6167 6d61 3a20 6e6f 2d63 6163 6865 0d0a        agma:.no-cache..
0x0090   636f 6e74 656e 742d 6c65 6e67 7468 3a20        content-length:.
0x00a0   3139 3038 0d0a 7033 703a 2043 503d 224e        1908..p3p:.CP="N
0x00b0   4f4e 2043 5552 204f 5450 6920 4f55 5220        ON.CUR.OTPi.OUR.
0x00c0   4e4f 5220 554e 4922 0d0a 636f 6e74 656e        NOR.UNI"..conten
...
```

pdweb.wan.azn trace for transaction authorization decisions

The **pdweb.wan.azn** component traces authorization decisions for all transactions.

The trace information includes:

- Credential details upon which the authorization decision is made.
- The resource that is accessed.
- The result of the authorization decision.

The following sample output is an extract of the trace which is produced from a single authorization decision.

Example pdweb.wan.azn output (extract)

```
...
2007-03-09-16:12:35.195+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:95: [10.251.140.1] Dumping attrlist: creds
2007-03-09-16:12:35.195+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AUTHENTICATI
ON_LEVEL , Value: 1

2007-03-09-16:12:35.195+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_AUT
HNMECH_INFO , Value: LDAP Registry
```

```
2007-03-09-16:12:35.195+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_AUT
HZN_ID , Value: cn=Scott Exton,o=ibm,c=au

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_AUT
H_METHOD , Value: password

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_BRO
WSER_INFO , Value: curl/7.12.1 (i386-redhat-linux-gnu) libcurl/7.12.1 OpenSSL/0.9.7a zlib/
1.2.1.2 libidn/0.5.6

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_IP_
FAMILY , Value: AF_INET

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_MEC
H_ID , Value: IV_LDAP_V3.0

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_NET
WORK_ADDRESS_BIN , Value: 0x0afb8c01

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_NET
WORK_ADDRESS_STR , Value: 10.251.140.1

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_PRI
NCIPAL_DOMAIN , Value: Default

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_PRI
NCIPAL_NAME , Value: scotte

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_PRI
NCIPAL_UUID , Value: ad987b08-cdf4-11db-a51a-000c29e9c358

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_QOP
_INFO , Value: SSK: TLSV1: 35

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_REG
ISTRY_ID , Value: cn=Scott Exton,o=ibm,c=au

2007-03-09-16:12:35.196+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_USE
R_INFO , Value:

2007-03-09-16:12:35.197+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:AZN_CRED_VER
SION , Value: 0x00000600

2007-03-09-16:12:35.197+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:tagvalue_log
in_user_name , Value: scotte

2007-03-09-16:12:35.197+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:117: [10.251.140.1] Attr Name:tagvalue_ses
sion_index , Value: 2c5bfdba-ce05-11db-bba0-000c29e9c358

...

2007-03-09-16:12:35.197+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:233: [10.251.140.1] INPUTS - protected_res
ource=/WebSEAL/webpi.vwasp.gc.au.ibm.com-default/index.html, operation=r

...

2007-03-09-16:12:35.198+10:00I----- thread(5) trace.pdweb.wan.azn:9 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:254: [10.251.140.1] OUTPUT - permission=0

...
```

```
2007-03-09-16:12:35.198+10:00I----- thread(5) trace.pdweb.wan.azn:8 /project/amwebrte610/b
uild/amwebrte610/src/pdwebrte/webcore/amw_azn.c:261: [10.251.140.1] CII EXIT amw_azn_decis
ion_access_allowed_ext with status=0x00000000
```

Setting trace for Security Access Manager SPNEGO issues

Enable trace for SPNEGO so that you can retrieve more details about WebSEAL issues.

About this task

When directed by IBM Support, collect the SPNEGO diagnostic data when WebSEAL does not start.

Procedure

1. Log in to the local management interface.
2. From the top menu, select **Secure Web Settings** > **Manage** > **Reverse Proxy**.
3. Select your reverse proxy.
4. Select **Manage** > **Configuration** > **Edit Tracing Configuration File**.
5. Add an entry to the file. Example entry that directs output to the `spnegotrace.log` file:

   ```
   bst:*.9:TEXTFILE:spnegotrace.log
   ```
6. Turn on per-process trace to diagnose WebSEAL start issues:
   a. Uncomment the last line by removing the # symbol.

   ```
   #Route to a per-process text file
   *:*.9:TEXTFILE.10.1000:trace.log
   ```
   b. Stop and restart WebSEAL.
   c. Check the output in `trace.log`.

   Note: If you start WebSEAL with the **`pdweb_start`** command, there are two traces with different `pids`.

Available trace components

The following table contains all trace components that are common to all Security Access Manager servers:

Table 10. Common trace components

Component	Description
pd.bst.general	Used to trace the Kerberos authentication process.
pd.acl.general	The general trace for the authorization API.
pd.acl.client	Used to trace the plug-in services for the authorization server.
pd.acl.authzn	Used to trace the authorization decision.
pd.acl.adminsvc	Used to trace the interface into the administration service plug-in.
pd.acl.remsvc	Used to trace the authorization decision during remote mode operation.
pd.acl.aznapi	Used to trace the usage of the Security Access Manager authorization API.
pd.acl.aznsvc	Used to trace the plug-in services that are provided by the authorization server.

Table 10. Common trace components (continued)

Component	Description
pd.idb.database	Used to trace access to the Security Access Manager policy database.
pd.ivc.ira	The IRA is the Security Access Manager interface into the LDAP server. This trace component is used to trace the Security Access Manager communication with the LDAP server.
pd.mgr.general	Used to trace the Security Access Managerr administration commands in the Policy Server.
pd.mgr.svrmgmt	Used to trace the management of the authorization servers within the policy server.
pd.ias.general	User to trace the Security Access Manager supplied authentication mechanisms, otherwise known as CDASs.
pd.ras.exception.trace	Used to trace any exceptions that might be caught by the server.

The following table contains all available pdadmin trace components:

Table 11. The pdadmin trace components

Component	Description
pdweb.bca.general	Used to trace the client side of the Security Access Manager authorization API.
pdweb.bca.user	Used to trace the client side of **user** pdadmin command.
pdweb.bca.group	Used to trace the client side of **group** pdadmin command.
pdweb.bca.acl	Used to trace the client side of **acl** pdadmin command.
pdweb.bca.protobj	Used to trace the client side of **object** pdadmin command.
pdweb.bca.protobjspace	Used to trace the client side of **objectspace** pdadmin command.
pdweb.bca.appsvrcfg	Used to trace the client side of **user** config command.
pdweb.bca.ssoresource	Used to trace the client side of **user** rsrc command.
pdweb.bca.ssoresourcegroup	Used to trace the client side of **rsrcgroup** pdadmin command.
pdweb.bca.ssocred	Used to trace the client side of **rscrcred** pdadmin command.
pdweb.bca.action	Used to trace the client side of **action** pdadmin command.
pdweb.bca.server	Used to trace the client side of **server** pdadmin command.
pdweb.bca.pop	Used to trace the client side of **pop** pdadmin command.
pdweb.bca.domain	Used to trace the client side of **domain** pdadmin command.
pdweb.bca.authzrule	Used to trace the client side of **authzrule** pdadmin command

The following table contains all available WebSEAL trace components:

Table 12. The WebSEAL trace components

Component	Description
pdweb.wan.ssl	Used to trace the SSL connection between WebSEAL and junctioned web servers.
pdweb.wns.session	Used to trace the WebSEAL sessions, as they are stored within the session cache and retrieved or removed from the session cache.
pdweb.wns.authn	Used to trace the authentication processing. **Note:** This trace component includes the header information that WebSEAL uses for header-based authentication. This header might contain sensitive information. For example, a BA header.
pdweb.adm.config	Used to trace the configuration for e-community SSO.
pdweb.wan.bool	Used to trace the WebSEAL processing of Security Access Manager authorization rules. Additional trace for Security Access Manager authorization rules can be enabled with the **pd.acl.authzn** trace component.
pdweb.wns.compress	Used to trace the WebSEAL compression of HTTP messages.
pdweb.cas.general	Used to trace the interface between WebSEAL and a custom-written CDAS shared library.
pdweb.wco.azn	Used to trace the entitlements service, which manages the maximum concurrent web session policy. The policy is used with SMS to limit the number of times a particular user can create a session concurrently.
pdweb.debug	Used to trace the HTTP headers sent between WebSEAL and the client. **Note:** The pdweb.debug trace could contain sensitive information.
pdweb.snoop.client	Used to trace the HTTP packets that are transmitted between WebSEAL and the client. **Note:** This component traces each request and response in its entirety as it is read off the socket. This trace might contain sensitive information.
pdweb.snoop.jct	Used to trace the HTTP packets that are transmitted between WebSEAL and the junctioned back-end web server. **Note:** This component traces each request and response in its entirety as it is read off the socket. This trace might contain sensitive information.
pdweb.url	Used to trace the creation and parsing of the URL.
pdweb.wan.azn	Used to trace the WebSEAL authorization decision.
pdweb.wan.ltpa	Used to trace the management of LTPA cookies.
pdweb.oauth	Used to trace OAuth EAS authorization decisions. **Note:** This component traces the data that passes into the EAS, which is governed by the **[azn-decision-info]** stanza. This trace might contain sensitive information.
pdweb.http.transformation	Used to trace HTTP transformation processing. **Note:** This component traces the header information in the request, which might contain sensitive information. For example, a Basic Authentication header.
pdweb.http2	Used to trace HTTP/2 client connections.

Table 12. The WebSEAL trace components (continued)

Component	Description
pdweb.http2jct	Used to trace HTTP/2 junction server and proxy connections.

Chapter 11. Common Security Access Manager problems

Check the following information for issues with any of the following Security Access Manager base components:

- Security Access Manager policy server
- Security Access Manager authorization server
- Security Access Manager Runtime
- Security Access Manager Runtime for Java

Environment information messages in the server log file at startup

In Security Access Manager, the severity level of the startup environment dump information added to the server process log is **WARNING**.

When a Security Access Manager server starts, a series of **WARNING** messages display information about the environment and AIX, Linux, or Solaris `ulimit`. These messages are informational. The support team uses these messages to diagnose problems.

The following example shows a warning:

```
2009-08-27-03:54:43.017+00:00I----- 0x1354A0CD pdmgrd WARNING ivc
general azn_maint.cpp 4977 0x00000001 HPDCO0205W  -------------------------
2009-08-27-03:54:43.017+00:00I----- 0x1354A0CC pdmgrd WARNING ivc
general azn_maint.cpp 4998 0x00000001 HPDCO0204W  Informational Message -
The environment variable for the running process :
_=/opt/PolicyDirector/bin/pdmgrd
...
TCD_PRODNAME=HPD
MAILMSG=[You have new mail]
PDCONFOBF=/opt/PolicyDirector/etc/pd.conf.obf
PWD=/
TZ=CST6CDT
PD_SVC_ROUTING_FILE=/opt/PolicyDirector/etc/pdmgrd_routing
....
2009-08-27-03:54:43.018+00:00I-----
0x1354A0CD pdmgrd WARNING ivc general azn_maint.cpp 505
0 0x00000001 HPDCO0205W  getrlimit():
RLIMIT_DATA (rlim_cur: 134217728 ; rlim_max: 2147483647)
2009-08-27-03:54:43.019+00:00I----- 0x1354A0CD
pdmgrd WARNING ivc general azn_maint.cpp 505
0 0x00000001 HPDCO0205W  getrlimit(): RLIMIT_STACK
(rlim_cur: 33554432 ; rlim_max: 2147483646)
2009-08-27-03:54:43.019+00:00I----- 0x1354A0CD
pdmgrd WARNING ivc general azn_maint.cpp 505
0 0x00000001 HPDCO0205W  getrlimit():
RLIMIT_AS (rlim_cur: 2147483647 ; rlim_max: 2147483647)
2009-08-27-03:54:43.019+00:00I-----
0x1354A0CD pdmgrd WARNING ivc general azn_maint.cpp 503
1 0x00000001 HPDCO0205W  --------------------
----------------------------
```

Unable to create new user

One of the most common error messages that you might see when you create a user is as follows:

```
Could not perform the administration request.
Error: Password rejected due to the Minimum Non-Alphabetic Characters policy
(status 0x13212131)
```

This error indicates, for example, that the password "abc" that was specified when you attempt to create the user does not comply with one of the user password policies that is defined. To view the help text for the Security Access Manager policy commands, enter the following command:

```
pdadmin> help policy
```

The previous password policy error can be solved by using one of the following solutions:

- Determine the minimum non-alphabetic character policy with the following command:

  ```
  pdadmin> policy get min-password-non-alphas
  ```

  Using this value, create the user with a password that contains the required minimum number of non-alphabetic characters.
- Modify the Security Access Manager non-alphabetic character policy before creating the user with the following command:

  ```
  pdadmin> policy set min-password-non-alphas number
  ```

Unable to authenticate user

After you create a user, this user cannot authenticate immediately with the new Security Access Manager user identity until the account is modified.

About this task

Security Access Manager user definitions are initially created with the account disabled (`Account valid = no`).

This condition is frequently the cause of authentication failures. To modify the account, complete the following steps:

1. Use the **user modify** command to enable the account:

   ```
   pdadmin> user modify user-name account-valid yes
   ```
2. Use the **user show** command to verify this change:

   ```
   pdadmin> user show user-name
   ```

Unexpected access to resources

Accesses to a protected system resource are either being unexpectedly granted or denied. It is always wise to first validate that the Security Access Manager processes are started and running normally.

Also check to ensure that the Security Access Manager message log files do not flag any operational problems. If Security Access Manager seems operationally sound, the problem is likely due to the policies that have been defined and applied to that system resource.

There are three Security Access Manager policy mechanisms that can be used to control access to your protected resources: ACLs, POPs, and authorization rules. Use the **pdadmin** commands to learn which ACL in your protected object space hierarchy has control over the access to the protected resource.

ACL commands

Security Access Manager access control depends on the following conditions:

- The ACL that controls the requested object must contain appropriate access permissions for the requesting user.
- The requested object must be accessible to the requesting user.

  Accessibility to protected objects is controlled by the traverse (**T**) permission. The traverse permission is only applied to container objects in the protected object space. The traverse permission specifies that a user, group, any-other, or unauthenticated user, that is identified in the ACL entry, has permission to pass through this container object to gain access to a protected resource object that is below it in the hierarchy.

If an ACL is directly attached to the protected object, this ACL defines the ACL policy for that object. If an ACL is not directly attached to the protected object, the controlling ACL is the nearest one that is above it in the protected object hierarchy.

Listing ACLs

Lists all ACLs that are defined in Security Access Manager:

```
padmin> acl list
```

Finding ACLs

Displays where each of those ACLs is attached within the protected object space hierarchy:

```
pdadmin> acl find acl_name
```

Showing ACLs

Examines the controlling ACL to check that it is correct for the type of enforcement wanted:

```
pdadmin> acl show acl_name
```

Correct the ACL definition if needed.

POP commands

A protected object is accessible to a requester if the requester possesses the traverse permission on each ACL attached to container objects above the requested resource on the path towards root and including root.

Additionally, use the **pdadmin** command to learn which POP (if any) within your protected object space hierarchy controls access to the protected resource in question.

If a POP is directly attached to the protected object in question, this POP defines the POP policy for that object. If a POP is not directly attached to the protected object in question, the controlling POP is the nearest one that is above it in the protected object hierarchy.

Listing POPs

The following command lists all of the POPs which are defined for Security Access Manager:

```
padmin> pop list
```

Finding POPs

The following command enables you to learn where a particular POP is attached within the protected object space hierarchy:

```
pdadmin> pop find pop_name
```

Showing POPs

Examine the controlling POP with the following command to ensure that it is correct for the type of enforcement desired:

```
pdadmin> pop show pop_name
```

Correct the POP definition if needed.

Authorization rule commands

If an authorization rule is directly attached to the protected object in question, this authorization rule defines the rule policy for that object. If an authorization rule is not directly attached to the protected object in question, the controlling rule is the nearest one that is above it in the protected object hierarchy.

Listing rules

The following command lists all of the authorization rules defined for Security Access Manager:

```
padmin> authzrule list
```

Finding rules

The following command enables you to learn where a particular authorization rule is attached within the protected object space hierarchy:

```
pdadmin> authzrule find authznrule_name
```

Showing rules

Use the following command to examine the controlling authorization rule and to ensure that it is correct for the type of enforcement required:

```
pdadmin> authzrule show authznrule_name
```

Correct the rule definition if needed.

Password change does not work in a multidomain environment

Specific configuration conditions for policy server, subdomains, and WebSEAL can cause password changes to fail.

About this task

A WebSEAL instance cannot change user passwords under all the following conditions because of the absence of ACL settings that are required to search domain locations:

- You configured the policy server in a nondefault location that is a location other than secAuthority=Default.
- You create Security Access Manager subdomains under the new location.
- You configured a WebSEAL instance in any of the new subdomains.

Complete the following steps to set the correct ACL with the following assumptions:

- The management domain name is Default.
- The Default domain is in an LDAP suffix that is called O=IBM,C=US.
- The subdomain names are Domain1, Domain2, and so on.

1. Place the following in a file called aclEntry.ldif:
   ```
   ##------ START: Do not include this line -----##
   dn: secAuthority=Default,o=ibm,c=us
   changetype: modifyI
   add: aclentry
   ```

```
aclentry:group:cn=SecurityGroup,SecAuthority=Domain1,cn=SubDomains
,SecAuthority=Default,O=IBM,C=US,O=IBM,C=US:object:ad:normal
:rwsc:sensitive:rwsc:critical:rwsc:system:rsc
aclentry:group:cn=SecurityGroup,SecAuthority=Domain2,cn=SubDomains,
SecAuthority=Default,O=IBM,C=US,O=IBM,C=US:object:ad
:normal:rwsc:sensitive:rwsc:critical:rwsc:system:rsc
##------ END: Do not include this line -------##
```

You must replace the management domain name `Default`, suffix `O=IBM,C=US`, and subdomains `Domain1`, `Domain2`, and so on, with the corresponding name of the current installation.

2. Update the ACL by running the following command:`ldapmodify -h host -p port -D cn=root -w pwd -i aclEntry.ldif`

Product pages might not display in browser if website is not trusted

Product pages might not display in a browser if the website is not trusted. For example, if the administrative console pages for Security Access Manager are empty even though the product is working, the issue might be caused by an untrusted website.

To resolve this issue, add the product website to the list of trusted websites in the security settings of your browser.

For example, if your URL to access Web Portal Manager is `http://wpm14.example.com:9060/ibm/console`, add the website to your list of trusted websites in the security settings of your browser.

Errors occur with in pdjrte startup when Java 2 security is enabled

Errors might occur at application startup when Java 2 security is enabled.

On startup, Security Access Manager for Java ensures that the JVM is properly configured for refreshing certificate expirations. If Java 2 security is enabled, invoke some security methods with privileged security enabled.

A workaround is update either the JVM `java.policy` file or the `was.policy` file for a WebSphere Application Server with the following entries:

```
permission java.security.SecurityPermission "insertProvider.IBMJCE";
permission java.security.SecurityPermission "putProviderProperty.IBMJCE";
```

Previously configured virtual host junctions are missing after upgrading to Security Access Manager 9

After upgrading to Security Access Manager 9, you might find that previously configured virtual host junctions cannot be found or used.

Along with the virtual host junctions missing, the following entries can be found on the WebSEAL message log:

```
DPWWA0308W Function ThirdPartyJunction - ltpa_read_key_file failed with errno 1
DPWWA1230E Error building junction / from file /var/pdweb/<webseal_instance>/server-root
/jct/junction.xml: HPDBA0521I Successful completion
DPWWA1211E Could not load junction database (/var/pdweb/<webseal_instance>/server-root
/jct,0x38cf07ce)
```

To resolve this issue, add the LTPA keys to the cluster Primary Master.

Starting from Security Access Manager 9, the LTPA keys are synchronized using the appliance clustering feature. All LTPA keys must be located on the Primary Master. Any keys created on other nodes with the same name will be overwritten when the cluster synchronizes. Any keys created on other nodes that do not match a name on the Primary Master will be deleted.

Chapter 12. Common user registry problems

This chapter details common user registry or directory server problems that you might encounter when you use Security Access Manager.

Security Directory Server common problems

This section details common problems that you might encounter when you use IBM Security Directory Server as the user registry.

Location of error logs

When a problem occurs that seems to be related to Security Directory Server, check for error messages that are related to that product.

You can find locations of Security Directory Server log files in the Knowledge Center for your version of the product.

Security Directory Server error log warnings

Security Directory Server error log indicates several "does not exist" warnings.

When the policy server is configured, the policy server is first unconfigured as part of this configuration process to ensure that it was cleaned up completely.

The unconfiguration step attempts to remove entries in the directory server.

When the policy server configuration is not yet completed, these entries are not yet created and might not exist in the LDAP registry. The Security Directory Server logs these entry removal attempts as warnings in its error log. These warnings are therefore normal and can be safely ignored.

Security Directory Server Instance Administration tool does not display instance on Red Hat Enterprise Linux 6

After you install Security Directory Server with the installation wizard typical installation path, the default instance is created. However, on Red Hat Linux, version 6, the Instance Administration tool started at the end of installation does not display the instance.

To verify that the default instance is listed in the configuration, use the **idsilist** command. By default, this command is in the /opt/ibm/ldap/V6.3/sbin/ directory. For details about the command, see the IBM Security Directory Server, product documentation.

Although the instance does not display in the UI, you can follow the command-line tool configuration steps as documented in the ADK Installation topics in the Knowledge Center.

To use an instance other than the default instance that is created by the Security Directory Server installer, use command-line tools to create and configure a non-default instance. See the Security Directory Server documentation for more information.

Setting up SSL

For information about configuring IBM Directory Server to use SSL communication, see the IBM Knowledge Center.

Go to the following URL: http://www-01.ibm.com/support/knowledgecenter/SSVJJU/welcome

LDAP common problems

This section details common problems that you might encounter when you use an LDAP-based user registry, such as Security Directory Server.

About this task

For common problems that are specific to Security Directory Server, see "Security Directory Server common problems" on page 89.

LDAP does not start after suffix is created

After you create the `secAuthority=Default` suffix Security Directory Server does not start.

The following steps are required to prepare an LDAP server for use with Security Access Manager. These steps must be completed before you configure Security Access Manager:

1. Create the `secAuthority=Default` suffix.
2. Stop and restart the Security Directory Server to enable the server to recognize the newly created suffix

When command-line installation is used, these steps must be completed manually.

If the user attempts to create the `secAuthority=Default` suffix and restart Security Directory Server before you apply the schema modifications that are required by Security Access Manager, Security Directory Server fails to restart.

When the server fails to restart it logs an error message to the `slapd.errors` file. This message indicates that the **secAuthority** attribute is not defined. The `slapd.errors` file is in the `/tmp` directory on AIX, Linux, and Solaris operating systems and in the *`ldap_install_dir`*`\tmp` directory on Windows operating systems, where *ldap_install_dir* is the directory where Security Directory Server was installed.

Insufficient privileges to perform operations

You are receiving the following error:

```
Insufficient LDAP access privileges to perform operation.
```

This message indicates that a supplied LDAP suffix does not have the correct ACLs attached. The following reasons are possible causes for this problem:

- During configuration, Security Access Manager was unable to attach ACLs to the existing suffix because the directory entries necessary to instantiate the suffix were not created.
- The suffix was added after the initial configuration of the Security Access Manager management server and the required ACLs were not added manually.

The Web Administration Tool provided with the Security Directory Server client can be used to check the suffix and to add the appropriate ACLs manually. For information about how to accomplish this, see the Administering topics in the Knowledge Center.

Correct the ACLs on the suffix, and run the command again.

Active Directory common problems

This section details common problems that you might encounter when you use Active Directory as the user registry.

Receiving HPDRG0100E for Active Directory operations

During Security Access Manager configuration, you might receive the following HPDRG0100E error message:

```
HPDRG0100E The operation in the Active Directory registry for operation_id
failed with return error nnnnnnnn.
```

This message is issued when an Active Directory error cannot be resolved during configuration.

Use the following problem determination suggestions to resolve this error before you restart configuration.

- An error message HPDRG0100E that is similar to the following content refers to a schema write failure:

  ```
  HPDRG0100E The operation in the Active Directory registry for
  adschema_update.exe: ADSCHEMA_SET_SCHEMA_WRITE failed with return error
  35.
  ```

  For this case, ensure that the Remote Registry Windows service is running on the root Active Directory domain controller system. During configuration, the Active Directory schema is updated, which requires the Remote Registry Windows service to be running. If the service is not running, start the service and complete configuration. You can stop the service after the configuration is finished.
- If the HPDRG0100E error message contains an 8-digit return code that begins with 8007, use the Microsoft **net helpmsg** command to display relevant help text. Convert the last four digits (digits that are *nnnn* of the return code of the form 8007*nnnn*) from hexadecimal format to decimal format and issue the command by using decimal format for the last four digits of the return code:

  ```
  net helpmsg nnnn
  ```
- Pursue the problem with Microsoft support by using the Active Directory return code value provided in the HPDRG0100E error message.

After you resolve the cause of this error, you can restart the configuration operation.

Receiving HPDRG0101E The user password violates the Active Directory user password policies

The Security Access Manager HPDRG0101E error message can occur during Security Access Manager configuration or during password change operations. This error can occur even if the Active Directory Domain account password policy "Password must meet complexity requirements" is set to "Disabled."

To resolve this error, ensure that the password meets the requirements of the Microsoft account password complexity policy. To learn more about Microsoft account password complexity policy, search the Microsoft knowledge base.

Chapter 13. Single sign-on Issues: Windows Desktop single sign-on, Kerberos, and SPNEGO

Use the following information to troubleshoot and resolve single sign-on issues that involve Windows Desktop single sign-on, Kerberos, and SPNEGO.

Windows Desktop single sign-on, on the client end, uses the Simple and Protected GSS-API Negotiation (SPNEGO) authentication protocol over HTTP to authenticate with WebSEAL.

SPNEGO authentication works by wrapping a Kerberos authentication token, obtained by the windows Desktop browser, and sending it in an HTTP header to the target web server without the need for user action. The user signs on to their Windows Desktop, and the browser can use the sign on to send the Kerberos token by means of SPNEGO to the web server for single sign-on, assuming the Web Server can handle SPNEGO or Kerberos.

WebSEAL on AIX, Linux, or Solaris use Kerberos to validate SPNEGO authentication data.

Problems with SPNEGO

Use the following troubleshooting tips for issues that involve SPNEGO authentication.

Web security server not starting

The following information describes debugging SPNEGO configuration problems with WebSEAL configuration where one of the Web security servers does not start. If the WebSEAL server fails to start, the server log file for that server contains messages that describe the problem.

Collecting data for Security Access Manager: WebSEAL (SPNEGO issues)

When WebSEAL does not start because of a SPNEGO issue, you might need to collect data for problem determination.

About this task

When directed by IBM Support, collect the SPNEGO diagnostic data when WebSEAL does not start.

Procedure

1. Turn on trace for each process by removing the # on the last line of the `/opt/pdweb/etc/routing` file. The last three lines of the routing file are shown:

   ```
   #
   # Route to a per-process text file
   #*:*.9:TEXTFILE.10.1000:/var/pdweb/log/trace__%ld.trace.log:644:ivmgr:ivmgr
   This will create a file in '/var/pdweb/log/trace __%ld.trace.log'
   ```

   Note: Ensure that enough disk space is available in the `/var` directory. If WebSEAL is started with the **`pdweb_start`** command, there are two traces with different `pids`.

2. Start WebSEAL to recreate the issue.
3. Turn off trace for each process by replacing the # at the beginning of the last line of the `/opt/pdweb/etc/routing` file.
4. Collect the following files:
   - `Webseald-`*`instance_name`*`.conf`
   - `msg__webseald-`*`instance_name`*`.log`
   - `trace_`*`pid`*`.trace.log`
   - The `krb5.conf` file if WebSEAL is on AIX, Linux, or Solaris
   - The Keytab file if WebSEAL is on AIX, Linux, or Solaris
   - `ldap.conf` for WebSEAL
   - `Activedir_ldap.conf` if Active Directory is the user registry
5. Collect the following information:
   - The output of the **pdversion** command on the WebSEAL server system
   - If WebSEAL is on AIX, Linux, or Solaris: **kinit** output when you use the keytab file
   - The **ktpass** command that is issued to create the keytab file
   - Active Directory Server version
6. Archive the data and send to support as directed by IBM Support.

No match to principal in key table

The server did not start, and the log file contains the following error:

```
HPDST0130E The security service function gss_import_name returned
the error 'No principal in keytab matches desired name'
(code 0x1cff2901/486484225)
```

The principal name for the SPNEGO service that is defined in the Security Access Manager server configuration file does not have a matching key in the SPNEGO key table. This error can occur for various reasons.

The algorithm to map the service principal name to the key in the SPNEGO key table completes the following processes:

1. Completes forward and reverse name resolution for the host name that is defined in the `spnego-krb-service-name` entry of the `[spnego]` stanza to discover the canonical host name.
2. Compares canonical host name to the realms defined in the `[domain_realm]` stanza of the `krb5.conf` configuration file.
3. Validates the principal key in the SPNEGO key table.

For details about these processes, see "Algorithm to resolve host names" on page 95.

The server configuration file for WebSEAL contains the `[spnego]` stanza. This stanza contains the following entries to examine:

spnego-krb-service-name

Defines the service principal name in the following format:

```
HTTP@hostname
```

The following example shows a definition of this entry in the configuration file:

```
HTTP@diamond.subnet2.ibm.com
```

spnego-krb-keytab-file

Defines the SPNEGO key table. This file contains principal keys in the following format:

```
HTTP/canonical_hostname@realm
```

The following example shows a key in the key table:

```
HTTP/diamond.subnet2.ibm.com@IBM.COM
```

The Kerberos `krb5.conf` configuration file contains the `[domain_realm]` stanza. This stanza contains entries that define the supported Kerberos realms. For details about this configuration file, see your Kerberos documentation.

Algorithm to resolve host names

The following process is used to map a service principal name to a key in the SPNEGO key table:

1. Resolve the host name to an IP address. The mapping process depends on your host name resolution configuration. Typically, the `/etc/hosts` file is checked first followed by the DNS server that is configured in the `resolv.conf` file.

   If the resolution succeeds, the process continues with step 2.

   If the resolution fails, the canonical name is assumed to be the same as the host name. The process continues with step 3.
2. Resolve the IP address to the canonical name. The mapping process depends on your host name resolution configuration. Typically, the `/etc/hosts` file is checked first followed by the DNS server that is configured in the `resolv.conf` file.

   If the IP address is found in the `/etc/hosts` file, the canonical name is set to the first host name that is listed.

   If the IP address is not found in the `/etc/hosts` file, the DNS server is queried to complete a reverse lookup on the IP address. If the DNS server returns a host name for this IP address, this host name becomes the canonical name.

   If the IP address is not found in the `/etc/hosts` file and if the DNS server does not return a host name for this IP address, the canonical name is assumed to be the same as the host name.

   Common error

   The `/etc/hosts` file lists the short name of the host before the fully qualified host name, the format of the `/etc/hosts` file is incorrect. Entries in the `/etc/hosts` file are in the following format:

   ```
   IP_address fully_qualified_hostname short_name
   ```

   When the format is incorrect, host name resolution might return the short name. The canonical name is then set to this short name. When this issue occurs, the Web server searches for the wrong key in the key table. The canonical name must be set to match the host name that clients use to contact the Web server.

   Resolution

   Contact your AIX, Linux, or Solaris system administrator on how to change entries in the following files:

   - `/etc/hosts`
   - `resolv.conf`
3. Map the canonical name from step 1 or step 2 to the realm name by checking the `[domain_realm]` stanza of the `/opt/PolicyDirector/etc/krb5.conf` file. Each entry in this stanza maps a host name or domain name to a realm name.

The canonical host name if checked against each of the host entries. If a matching host entry is found, the realm name becomes the realm that is specified for the host. If no matching host entry is found, the domain entries are checked. If a matching domain entry is found, the realm name becomes the realm that is specified for that domain.

If no matching domain entry is found, the realm name becomes the value of the `[libdefaults] default_realm` entry in the `/opt/PolicyDirector/etc/krb5.conf` file.

Common error
The entries in the `[domain_realm]` stanza of the `/opt/PolicyDirector/etc/krb5.conf` file are incorrect.

Resolution
Verify that the realm name specified in the `[domain_realm]` stanza is correct, and verify that the canonical name matches a host or domain entry in this stanza.

4. Verify that the key table contains this entry.

Common error
The key table does not contain a matching entry.

Resolution
Use the **am_klist** command or the **am_ktutil** program to check the SPNEGO key table for an entry in the following format:

```
HTTP/canonical_name@realm_name
```

For details about using the **am_ktutil** program, see "Validating keys in key tables" on page 99.

Problems with Kerberos

The following information can assist you when troubleshooting issues that concern Kerberos authentication.

Kerberos initialization failing

The following information describes debugging problems with the Kerberos **am_kinit** command. If the **am_kinit** command completes, it generates no output, but you can use the Kerberos **am_klist** command to view the principals in the ticket cache. If the **am_kinit** command fails, the error message provides the following details:

- The primary cause for the failure
- A hexadecimal status code
- The specific reason for the failure

The most common primary causes for the **am_kinit** command to fail are the following reasons:

- Failure to initialize the Kerberos libraries
- Unable to obtain initial credentials

Kerberos configuration

Problem: **am_kinit** crashes when running `am_kinit -k -t`

Solution: Some versions of **am_kinit** do not deal properly with problems when an entry is not found in a keytab file. Double-check that the keytab file has the exact same entry you are passing to **am_kinit**.

Unable to initialize Kerberos libraries

When you run the **am_kinit** command, you might receive an error message that states the following primary reason:

```
Initializing kerberos libraries failed.
```

The most common causes that the initial credentials cannot be obtained are the following reasons:

- Cannot open configuration file
- Improper format of configuration file

Cannot open configuration file: When you use the **am_kinit** command, you receive the following error:

```
Initializing kerberos libraries failed. Status 0x96c73a87
Cannot open or find the Network Authentication Service
configuration file.
```

The /opt/PolicyDirector/etc/krb5.conf file does not exist or cannot be opened. Verify that file exists and is readable by all users.

Improper format of configuration file: When you use the **am_kinit** command, you receive the following error:

```
Initializing kerberos libraries failed. Status 0x96c73a88
Improper format of Network Authentication Service
configuration file.
```

The /opt/PolicyDirector/etc/krb5.conf file contains a syntax error. No details about the syntax error are available. Edit the configuration file to identify and correct the syntax error.

Other configuration items to check when problems occur:

Problems can result because of configuration errors. Check the following configuration items when problems occur.

Other configuration items to check when problems occur:

- Check that the file permissions and ownership of the keytab file allow access by the plug-in authorization server.
- Check that the keytab file contains valid data and keys for the correct principal name by using the **am_ktutil** utility to display information that is contained in the keytab file.
- Check that the DNS configuration for the entire domain (domain controller and clients) is correct and that names resolve correctly and match the values in the service principal name configuration items in various locations (such as keytab file, and plug-in configuration file).
- Check that system clocks are synchronized and that a distributed time service is maintaining clock synchronization on all systems in the domain.
- Check that the network configuration is correct and that there are no issues such as congestion, incorrect routing, or name collision. Ensure that the latency is within tolerable limits. Be sure that firewalls, NAT, and other network security services do not interfere with the operation of the domain.

Unable to obtain initial credentials

When you run the **am_kinit** command, you might receive an error message that states the following primary reason:

```
Unable to obtain initial credentials.
```

The most common causes that the initial credentials cannot be obtained are the following reasons:

- Cannot resolve the network address of the key distribution center (KDC)
- Cannot contact the KDC
- Clocks are not synchronized
- Pre-authentication failure
- Client not found in authentication database, or client that is locked out

Cannot resolve address of key distribution center: When you use the **am_kinit** command, you receive the following error:

```
Unable to obtain initial credentials. Status 0x96c73adc
Cannot resolve network address for KDC in requested realm.
```

The host name for the key distribution center (KDC) that is defined in the /opt/PolicyDirector/etc/krb5.conf file is not valid. Edit this configuration file to correct the host name for the KDC.

Cannot contact the key distribution center: When you use the **am_kinit** command, you receive the following error:

```
Unable to obtain initial credentials. Status 0x96c73a9c
Cannot contact any KDC in requested realm.
```

The host name for the key distribution center (KDC) that is defined in the /opt/PolicyDirector/etc/krb5.conf file is valid, but the KDC cannot be contacted. Verify the following conditions:

- Ensure that the krb5.conf configuration file defines the correct host name and port for the KDC.
- Ensure that the KDC is running.
- Ensure that there is network connectivity between the client and the KDC.

Clocks are not synchronized: When you use the **am_kinit** command to test an AIX, Linux, or Solaris server key table, you receive the following error:

```
Unable to obtain initial credentials. Status 0x96c73a25
Clock skew too great.
```

To resolve this condition, keep system clocks synchronized. For a permanent solution, deploy a time synchronization service on your systems. For a temporary solution, adjust the clocks on the systems so they are within one minute of each other.

Pre-authentication failure: When you use the **am_kinit** command to test an AIX, Linux, or Solaris server key table, you receive the following error:

```
Unable to obtain initial credentials. Status 0x96c73a18
Preauthentication failed.
```

The key in the key table is incorrect. A common reason is that the password for the principal in the Active Directory server was changed. In this case, regenerate the key table. If the password was not changed, make sure that you generated the key table correctly by using the correct principal name, Active Directory user name, and path.

Client not found or locked out: When you use the **am_kinit** command to test an AIX, Linux, or Solaris server key table, you receive the following error:

```
Unable to obtain initial credentials. Status 0x96c73a06
Client not found in Network Authentication Service database
or client locked out.
```

The key table does not have a key for the specified principal. Check whether an error was made when the principal was typed, or whether the key table was generated incorrectly. You can use the Kerberos **am_ktutil** commands to check which keys are in the key table. For details about this procedure, see "Validating keys in key tables."

Useful Kerberos procedures

Use the following procedures to help troubleshoot a SPNEGO or Kerberos problem.

Validating keys in key tables

About this task

There is a test function available in the appliance Web interface to test Kerberos keyfiles.

From the main menu select **Secure Web Settings -> Global Settings -> Kerberos Configuration**. In the Kerberos Configuration panel select the Keyfiles tab and click on the **Test** button.

Unable to authenticate

If a user attempts to authenticate to WebSEAL by using SPNEGO authentication and the authentication fails, an HTML error page is displayed and a message is added to the log.

Ticket not yet valid

A user attempts to access WebSEAL and receives an HTML page with the following error:

```
HPDIA0100E An internal error has occurred.
```

The trace log file contains the following one of the following messages:

- `HPDST0130E The security service function gss_accept_sec_context returned the error 'Ticket not yet valid' (code 0x96c73a21/-1765328351).`
- `HPDST0130E The security service function gss_accept_sec_context returned the error 'Clock skew too great' (code 0x96c73a25/-1765328347).`

The system clock on the client system is not synchronized with the system clock on the Active Directory server. When you use Kerberos, these clocks must be synchronized. For a permanent solution, deploy a time synchronization service on your system. For a temporary solution, adjust the clocks on the system so that they are within one minute of each other.

Cannot acquire credentials

A user attempts to access WebSEAL and receives an HTML page with the following error:

```
HPDIA0114E Could not acquire a client credential.
```

This same message is written to the log file.

The user exists in the Active Directory user registry and presented valid SPNEGO authentication data, but the user does not exist in the Security Access Manager user registry.

SPNEGO authentication requires that the user exists in both the Active Directory and the Security Access Manager user registries. If you believe that the user exists in both user registries, verify that the user ID produced by SPNEGO authentication matches what you expect.

To see the user ID, complete the following steps:

1. Enable the `pd.ias` authentication trace by using the following **pdadmin** command:
   ```
   pdadmin sec_master> server task serverName trace set  \
      pd.ias 9 file path=/tmp/ias.log
   ```
2. Have the same user attempt to access the Web server again. After this user receives the HPDIA0114E message, disable the authentication trace by using the following **pdadmin** command:
   ```
   pdadmin sec_master> server task server trace set pd.ias 0
   ```
3. Examine the `/tmp/ias.log` file for a message that is similar to the following message:
   ```
   Mapped name user@realm to am_user
   ```
4. Ensure that the *am_user* user is a defined user in the Security Access Manager user registry.

Wrong principal in request

A user attempts to access WebSEAL and receives an HTML page with the following error:

```
HPDIA0100E An internal error has occurred.
```

The trace log file contains the following message:

```
HPDST0130E The security service function gss_accept_sec_context returned
the error 'Wrong principal in request' (code 0x96c73a90/-1765328240).
```

The server principal name (SPN) supplied by the client in the SPNEGO authentication header does not match the SPN being used by the Web security server. This error can be caused in the following situations:

- The user did not specify the fully qualified host name (FQHN) when you contact the Web security server. Clients must use the FQHN so that the Active Directory server can provide the client with an appropriate Kerberos authentication ticket.
- The Web security server is configured to use the wrong SPN. The host name portion of the principal in the Kerberos key table must match the host name that is being used by the client to contact the Web security server. If the principal name in the key table is incorrect, the key table must regenerate on the key distribution center (KDC) by using the **ktpass** command with the **–princ** option. The value that is specified for the **–princ** option must be the same host name that client uses to contact the Web security server.

  For example, for clients to contact the Web security server at `https://diamond.example.ibm.com` and the Web security server is in the `IBM.COM` Kerberos realm, specify the following value for the **–princ** option:

  ```
  HTTP/diamond.example.ibm.com@IBM.COM
  ```

You can use the **am_ktutil** program to examine the contents of the Kerberos key table.

Encryption type not permitted

A user attempts to access WebSEAL and receives an HTML page with the following error:

```
HPDIA0100E An internal error has occurred.
```

The trace log file contains the following message:

```
HPDST0130E The security service function gss_accept_sec_context returned
the error 'Encryption type not permitted' (code 0x96c73ae9/-1765328151).
```

The encryption type in the SPNEGO authentication header does not match any of the encryption types that the Kerberos libraries are configured to accept. To resolve the issue, ensure that the /opt/PolicyDirector/etc/krb5.conf configuration file defines the following entries in the [libdefaults] stanza:

```
default_tkt_enctypes = rc4-hmac des-cbc-md5 des-cbc-crc aes256-cts aes128-cts
default_tgs_enctypes = rc4-hmac des-cbc-md5 des-cbc-crc aes256-cts aes128-cts
```

After you save these changes, restart the Web security server

Key version is incorrect

A user attempts to access WebSEAL and receives an HTML page with the following error:

```
HPDIA0100E An internal error has occurred.
```

The trace log file contains the following message:

```
HPDST0130E The security service function gss_accept_sec_context returned
the error 'Key version number for principal in key table is incorrect'
(code 0x96c73ae3/-1765328157).
```

The key version in the Kerberos authentication header does not match the key version in the SPNEGO key table. This error might occur after the password for the Kerberos principal for the Web security server is changed in the Active Directory server. After you change this password, complete the following steps:

1. Regenerate the SPNEGO key table.
2. Replace this key table on the Web security server.
3. Restart the Web security server

For more information, see item 870987 in the Microsoft knowledge base.

Cannot authenticate by using NTLM

When you attempt to access a Web security server, you receive the following error messages:

```
DPWWA2403E Your browser supplied NTLM authentication data.
NTLM is not supported by WebSEAL. Ensure that your browser
is configured to use Integrated Windows Authentication.
```

WebSEAL does not support NT LAN Manager (NTLM) authentication. Some browsers support NTLM authentication only or are configured to send NTLM authentication tokens instead of SPNEGO tokens. A browser that supports SPNEGO might be sending NTLM tokens for the following reasons:

- Microsoft Internet Explorer is not configured with the WebSEAL server in the "Trusted sites" or "Local intranet" zone.
- Microsoft Internet Explorer is not configured for Integrated Windows Authentication.
- The client workstation and the WebSEAL server might be a member of different Active Directory domains (Kerberos realms).
- The client workstation is not logged in to the Active Directory domain.
- The client workstation is not specifying the correct host name to access the WebSEAL server. The value that is specified for the **–princ** option of the **`ktpass`** command must be the same host name that client uses to contact the Web security server.

  For example, for clients to contact the Web security server at `https://diamond.subnet2.ibm.com` and the Web security server is in the `IBM.COM` Kerberos realm, specify the following value for the **–princ** option:

  ```
  HTTP/diamond.subnet2.ibm.com@IBM.COM
  ```

Under certain circumstances, clients cannot be prevented from sending NTLM authentication tokens. Under these circumstances, you might not be able to directly use SPNEGO authentication with the WebSEAL server. Instead, you can configure the Web Server Plug-in for IIS to serve as an e-community SSO (ECSSO) master authentication server (MAS). In this configuration, the Web server plug-in must be configured to support both NTLM and SPNEGO tokens. The WebSEAL server can now receive ECSSO tokens from the MAS.

Cannot complete authentication

When you attempt to access a security Web server, you receive the following error messages:

```
HPDIA0220I Authentication requires continuation before
completion status can be determined.
```

This error occurs when the password used to encrypt the SPNEGO authentication data is not synchronized with the password that is used by the Web security server to decrypt the SPNEGO authentication data.

On AIX, Linux, and Solaris operating systems, this error can be caused for the following reasons:

- The password for the Web security principal in Active Directory changed. When this password is changed, you must regenerate the SPNEGO key table.
- The SPNEGO key table was updated, but the client is still presenting an old authentication token. To clear the cached copy of the authentication token because SPNEGO authentication tokens are cached, log the client out of the workstation and log back in or restart the workstation.

On Windows operating systems, this error can be caused for the following reasons:

- The Web security server might be running in the foreground instead of as a service. The Web security server must be running as a service to have access to the correct password needed to decrypt the security token.
- The password for the Web security principal in Active Directory changed. If the Windows service for the Web security server is configured to log on as the Local System account, you might need to restart the system. If the Windows service for the Web security server is configured to log on using a particular Active Directory domain account, you might need to update the Windows service "Logon as" configuration to specify the correct account and password.

- Both the password and the Windows service configuration were updated, but the client is still presenting an old authentication token. To clear the cached copy of the authentication token because SPNEGO authentication tokens are cached, log the client out of the workstation and log back in, or restart the workstation.

Chapter 14. Common problems with WebSEAL servers

The following information details the common problems with WebSEAL servers.

For more information about Security Access Manager WebSEAL, see the Administrating topics in the Knowledge Center.

Cannot customize basic authentication response

When you use basic authentication (`BA-auth`) with Security Access Manager, you cannot customize the `acct_locked.html` file to contain more images. Although you can embed images in the file, subsequent requests to access the embedded images fail.

To customize user authentication for your environment, use other authentication methods. For example, you can use forms authentication to bypass the authentication check when images from the error page are requested.

WebSEAL not responding on ports 80 or 443

WebSEAL does not respond on either port 80 or port 443.

Determine whether you have another web server that is installed on the WebSEAL system. For example, the IBM HTTP Server is commonly installed during the installation of Security Directory Server. If another Web server is on the same system as WebSEAL, reconfigure it to listen on ports other than the ports that are used by WebSEAL.

Servers fail to start because of exceeded LDAP replica server limit

Security Access Manager supports a maximum of one host and nine LDAP replica servers, which are in the `ldap.conf` file.

If more than nine LDAP replica entries are listed in the `ldap.conf` file, the Security Access Manager servers cannot start. The following errors indicate this problem:

- **For the policy server:**

  ```
  HPDC00190E  Unable to configure LDAP replica
  "ldap10.ibm.com,3899,readonly,5" \
  into server, errorcode=0xe9.
  ```

  where "`ldap10.ibm.com,3899,readonly,5`" indicates the 10th replica entry in the `ldap.conf` file.
- **For WebSEAL:**

  ```
  DPWWA0314E Initialization of authorization API failed. Major status =0x1, \
  minor status = 0x1005b3a3
  ```
- **For the authorization server:**

  ```
  HPDAC0180E The Security Access Manager authorization server could not
  be started (0x1005b3a3).
  Please consult "Error Message Reference Guide" for explanation of
  minor error status 0x1005b3a3.
  ```

To resolve the error, specify no more than nine replica LDAP servers in the `replica` entry of the `[ldap]` stanza in the `ldap.conf` file.

For information about the `replica` entry of the `[ldap]` stanza in the `ldap.conf` file, see the Configuring topics in the Knowledge Center.

Multiple logins with e-community

WebSEAL e-community users are prompted to log in more than one time.

This problem can occur if two WebSEAL servers are configured in the same domain. In these cases, one WebSEAL server is configured as the Master Authentication Server (MAS), and the other WebSEAL server is configured to use the MAS for authentication. Attempts to access the latter might require the user to log in more than one time if the difference in system times between the two WebSEAL servers is too great.

To address this concern, synchronize the system time on each WebSEAL server that participates in an e-community.

e-Community SSO Master Authentication Server configured with EAI

You can configure e-Community single sign-on (SSO) so that the master authentication server (MAS) uses an external authentication service (EAI) to create a token for a consumer server.

To use this configuration:

1. Specify the MAS server as the default server.
2. Enable EAI in the WebSEAL default configuration file.
3. If you use virtual hosts, configure the WebSEAL virtual hosts configuration file (`webseal-vh.conf`) to use e-Community SSO.

This configuration works for any virtual host junction that is defined in the virtual hosts configuration that is in the same domain that is defined for e-Community SSO.

Verifying junctioned, third-party Web server

Verifying the correct operation of a junctioned, third-party Web application server is similar to the procedure for WebSEAL. Enter the following URL into your browser to verify that the third-party Web server is functioning properly:

```
http://junctioned-webserver-machinename
```

Do not specify a port number so that you can determine if the server is listening on port 80 (HTTP). If successful, the `index.html` page of the third-party Web server is displayed.

The WebSEAL junction for the Web server is created with a **`pdadmin`** command. In the following example, the junction points to the third-party Web server in the WebSEAL `/myjunction` file space:

```
pdadmin> server task webseald-webseal-machinename create -t tcp \
-p junctioned-server-port -h junctioned-webserver-machinename \
-c iv_user,iv_groups /myjunction
```

Try to access the `index.html` page on the junctioned Web server through WebSEAL with the following URL:

```
webseal-machinename/myjunction
```

If you configured the junction to use secure communication (**–t** ssl), your browser might issue warnings about the WebSEAL server certificate and prompt you for a user name and password. Enter sec_master for the user name and the appropriate password. If successful, the index.html page of the third-party Web server is displayed.

WebSEAL performance is degraded during file downloads

In Security Access Manager, you can use the io-buffer-size parameter in the **[junction]** stanza to configure the buffer size for reading and writing data to-and-from the junction. This value limits the amount of data that can be written from the socket to a junctioned server.

The optimum value for this io-buffer-size parameter is 8191 bytes (one byte less than the typical TCP buffer size of 8 KB). Severe performance degradation might occur if the value of the io-buffer-size is larger than 8191.

Similarly, you can use the io-buffer-size parameter in the **[server]** stanza to control the buffer size to read and write data to-and-from the client. The amount of data that can be written from the socket to an HTTP browser depends on the value of this parameter.

For either of these io-buffer-size parameters, a small value (for instance, 10 bytes) can hurt performance by causing frequent calls to the low-level read/write APIs. Up to a certain point, larger values improve performance. However, if the io-buffer-size exceeds the size of low-level I/O functions, there is no longer any improvement in performance. Using an io-buffer-size value that is too high degrades performance.

DPWWA0305E inconsistent message severity

The WebSEAL message log might include an entry such as the following that is recorded as a WARNING but that the code indicates is an ERROR and the error reference reports as FATAL.

```
2006-09-04-07:17:35.188+02:00I----- 0x38CF0131 webseald WARNING wwaserver
s:\amweb600\src\pdweb\webseald\http\server\WsTcpListener.cpp
3930x000007c8 DPWWA0305E     The 'pd_tcp_write' routine failed for
'WsTcpConnector::write', errno = 10054
```

This issue is a common connection reset error. The error message reports that the other end of the connection terminated the connection. The correct severity level is WARNING.

For more information about severity levels, see “Severity of message events” on page 51.

Error when you create an LTPA junction

When you create a lightweight third-party authentication (LTPA) junction, you might receive an error that WebSEAL is unable to parse the LTPA key.

WebSphere® Application Server does not add the "realm" component to a token unless global security is enabled, and WebSEAL expects this component to be present. To resolve this issue, configure global security for the LDAP registry in WebSphere and then regenerate the LTPA keyfile; WebSEAL can then successfully load the keyfile.

Password change fails in a multi-domain environment

In a multi-domain environment, WebSEAL can fail to change a user password because of insufficient ACL settings.

About this task

WebSEAL does not have correct ACL settings to search the Management Domain information in environments where:

- Security Access Manager Policy Server is configured in a non-default location. That is, a location other than **secAuthority=Default**.
- Security Access Manager subdomains exist.
- The WebSEAL instance is configured in one of the subdomains.

In this situation, WebSEAL cannot successfully change user passwords because of the lack of correct ACL settings.

You must set the correct ACLs so that WebSEAL can search the Management Domain and change user passwords in a multi-domain environment.

The provided procedure is based on the following environment:

- The Management Domain name is **Default**.
- The Management Domain is in an LDAP Suffix that is called **O=IBM,C=US**.
- There are two subdomains that are called **Domain1** and **Domain2**.

Note: You must modify the following steps to use the domain names and locations that match your environment.

Procedure

1. Create a file called `aclEntry.ldif`.
2. Copy the following contents into the file:

   Note: The two entries that start with `aclentry:` must each be entered as one line.

   ```
   ##------ START: Do not include this line -----##
   dn: secAuthority=Default,o=ibm,c=us
   changetype: modify
   add: aclentry
   aclentry:group:cn=SecurityGroup,SecAuthority=Domain1,cn=SubDomains,
     SecAuthority=Default,O=IBM,C=US,O=IBM,C=US:object:ad:normal:
     rwsc:sensitive:rwsc:critical:rwsc:system:rsc
   aclentry:group:cn=SecurityGroup,SecAuthority=Domain2,cn=SubDomains,
     SecAuthority=Default,O=IBM,C=US,O=IBM,C=US:object:ad:normal:
     rwsc:sensitive:rwsc:critical:rwsc:system:rsc
   ##------ END: Do not include this line -------##
   ```
3. Save the file.
4. Run the following command to update the ACL:

   ```
   ldapmodify -h host -p port -D cn=root -w pwd -i aclEntry.ldif
   ```

Results

WebSEAL can now successfully change user passwords.

Firefox does not send JSON POST requests correctly

When using the Firefox browser, JSON POST requests specified in the WebSEAL configuration file are not parsed. Your policies might produce incorrect or unexpected results.

In the WebSEAL configuration file, you can capture JSON data for a custom attribute by using the **`[azn-decision-info]`** stanza. In that stanza, if you have an entry for post-data that specifies JSON POST data, it will not be parsed.

Use an alternate browser when you have custom attributes configured that specify JSON POST requests. For example, use Chrome or Internet Explorer.

Chapter 15. Risk-Based Access External Authorization Service plug-in

The Risk-Based Access (RBA) External Authorization Service (EAS) component provides a runtime XACML EAS plug-in for WebSEAL to enforce a policy decision. WebSEAL becomes the authorization enforcement point to access resources protected by RBA.

The EAS collects context information about the user and the request, creates an XACML over SOAP decision request, and sends the information to the server.

Manage the EAS with entries in the `webseald.conf` file.

For more information about the risk-based EAS, see the Configuring topics in the IBM Knowledge Center. Search for **Runtime security services external authorization service** for details.

For assistance in troubleshooting RBA EAS issues, you can enable tracing, then review the logs for information about any issue that might be occurring.

Enabling External Authorization Service tracing on WebSEAL

To enable tracing and logging for the XACML EAS plug-in, issue the following **pdadmin** command:

```
pdadmin > server task WebSEAL_server_name trace set xacml_eas_comp_name 9
filepath=path_to_log_file
```

where:

webseal_server_name
: Is the name of the WebSEAL server.

xacml_eas_comp_name
: Is the name of the XACML EAS component.

path_to_log_file
: Is the directory where you want to store the trace log file.

For example:

```
pdadmin > server task default-webseald-localhost
trace set pdweb.xacml 9 file path=/tmp/xacml.log
```

Note: Tracing is disabled when you restart WebSEAL.

Chapter 16. Troubleshooting certificate compliance issues

When you enable Security Access Manager applications to implement a security compliance standard, certain settings are required.

The required settings apply to the standards of the following security settings:

- FIPS 140-2
- NIST Special Publications 800-131a (or SP 800-131a) Transition
- NIST SP800-131a Strict
- National Security Agency (NSA) Suite B 128 bit
- NSA Suite B 192 bit

To ensure a successful regeneration of the Security Access Manager side of the certificates, see the Administering topics in the IBM Knowledge Center.

WebSphere Application Server, version 8.0, requires certain settings to properly enable compliance. See

> http://publib.boulder.ibm.com/infocenter/ieduasst/v1r1m0/index.jsp?topic=/com.ibm.iea.was_v8/was/8.0.0.3/Security/WASV8003_SecurityCryptoSignatureAlgorithm/player.html

For support for NIST SP 800-131 and NSA Suite B, you must use IBM WebSphere Application Server, version 8.0.0.3 or later.

Other troubleshooting tips:

- **Check browser configuration**

  Your browser might not support or not be configured to support the TLS protocol.

  TLS 1.2 is not enabled by default. Check your browser documentation for instructions on how to enable TLS version 1.2.

  For example, for Internet Explorer, version 8 on Windows 7 and Windows 2008, go to **Tools** > **Internet Options** > **Advanced (Tab)** > **Security** and select **Use TLS 1.2**.

- **Check user registry configuration**

  Changing an SSL protocol to TLS, version 1.2, can affect communication between WebSphere Application Server and the user registry. If you receive an error message about failed connection, check your user registry configuration.

  The user registry must support TLS, version 1.2, if you use an SSL connection.

Chapter 17. Serviceability commands

The following reference information describes the serviceability and problem determination **pdadmin** commands and utilities.

Reading syntax statements

The reference documentation uses the following special characters to define syntax:

[ ] Identifies optional options. Options that are not enclosed in brackets are required.

... Indicates that you can specify multiple values for the previous option.

| Indicates mutually exclusive information. You can use the option to the left of the separator or the option to the right of the separator. You cannot use both options in a single use of the command.

{ } Delimits a set of mutually exclusive options when one of the options is required. If the options are optional, they are enclosed in brackets ([ ]).

\ Indicates that the command line wraps to the next line. It is a continuation character.

The options for each command or utility are listed alphabetically in the Options section or Parameters section. When the order of the options or parameters must be used in a specific order, this order is shown in the syntax statements.

Serviceability and problem determination commands

Table 13 lists the serviceability and problem determination commands that are available with the **pdadmin** utility.

Table 13. Serviceability and problem determination commands

Command	Description
"server list"	Lists all registered Security Access Manager servers.
"server task trace" on page 116	Enables the gathering of trace information for components of installed Security Access Manager servers or server instances that support debug event tracing.

For information about the command modes for the **pdadmin** utility, see the Command Reference topics in the Knowledge Center.

server list

Lists all registered Security Access Manager servers.

Requires authentication (administrator ID and password) to use this command.

Syntax

```
server list
```

Description

Lists all registered Security Access Manager servers. The name of the server for all server commands must be entered in the exact format as it is displayed in the output of this command. The **server list** command does not have such a requirement.

Options

None.

Return codes

0 The command completed successfully.

1 The command failed. When a command fails, the **pdadmin** command provides a description of the error and an error status code in hexadecimal format (for example, 0x14c012f2). See the Message topics in the Knowledge Center for more information.

Example

The following example lists registered servers:

```
pdadmin> server list
```

The output is as follows:

```
ivmgrd-master
ivacld-server1
ivacld-server2
```

where ivmgrd-master represents the Policy server; ivacld-server2 and ivacld-server1 represent Authorization server instances.

server task trace

Enables the gathering of trace information for components of installed Security Access Manager servers or server instances.

Requires authentication (administrator ID and password) to use this command.

Syntax

server task *server_name–host_name* **trace list** [*component*]

server task *server_name–host_name* **trace set** *component level* [*destination*]

server task *server_name–host_name* **trace show** [*component*]

Description

The **server task trace** command enables the gathering of trace information for components of installed Security Access Manager servers or server instances that support debug event tracing. The content of trace messages is generally undocumented and is intended to be used for debugging purposes only. The format and content of trace messages might vary between product releases.

Options

component
Specifies the component for which to enable (set) tracing.

destination
Specifies where the gathered statistics are written, where *destination* can be one of the following:

file path=*file_name*
Specifies the fully qualified file name.

log_agent
Specifies a destination for the statistics information gathered for the specified component. For more information about event logging, see the Administering topics in the Knowledge Center.

level
Specifies the level of tracing. The supported values for this option are `1` through `9`, with `9` reporting the most detailed level of information in the trace log.

server_name–*host_name*
Specifies the name of the server or server instance. You must specify the server name in the exact format as it is shown in the output of the **`server list`** command.

For example, if the configured name of a single WebSEAL server on host `cruz.dallas.ibm.com` is `default`, the *server_name* would be `default-webseald` and the *host_name* would be `cruz.dallas.ibm.com`. For this example, the name of the server would be `default-webseald-cruz.dallas.ibm.com`.

If there are multiple configured server instances on the same machine, for example, the host `cruz.dallas.ibm.com`, and the configured name of the WebSEAL server instance is `webseal2-webseald`, the *server_name* would be `webseal2-webseald` and the *host_name* would be `cruz.dallas.ibm.com`. For this example, the name of the server instance would be `webseal2-webseald-cruz.dallas.ibm.com`.

trace list [*component*]
Lists all enabled trace components that are available to gather and report trace information. If you specify the *component* option and the component is enabled, the output lists that component; otherwise, no output is displayed. If you do not specify the *component* option, the output lists all enabled components.

trace set *component level* [*destination*]
Sets the trace level and trace message destination for a specific *component* and its subordinates. The value for the *level* option is a single integer from `1` to `9`, with `9` reporting the most detailed level of information. The *destination* option specifies where the gathered trace information is written.

trace show [*component*]
Shows the names and levels for all enabled trace components. If you specify the *component* option, the output lists the name and level for the specified component.

Return codes

0
The command completed successfully.

1
The command failed. When a command fails, the **`pdadmin`** command

provides a description of the error and an error status code in hexadecimal format (for example, 0x14c012f2). For more information, refer to the Messages topics in the Knowledge Center.

Examples

- The following example enables the pdweb.debug trace component to level 2. Then displays the output for all enabled components. Note that WebSEAL–specific components are prefixed with pdweb.

  ```
  pdadmin sec_master> server task webseald-instance_name trace set
  pdweb.debug 2

  pdadmin sec_master> server task webseald-instance_name trace show
  ```

  Output from the **trace show** command is similar to:

  ```
  pdweb.debug 2
  ```

- The following example enables the pdwebpi.module.session-cookie trace component to level 9. Then displays the output for all enabled components. Components that are specific to the Web server plug-ins are prefixed with pdwebpi.

  ```
  pdadmin sec_master> server task pdwpi-tivoli.com trace set
  pdwebpi.module.session-cookie 9

  pdadmin sec_master> server task pdwpi-tivoli.com trace show
  ```

  Output from the **trace show** command is similar to:

  ```
  pdwebpi.module.session-cookie 9
  ```

See also

"server list" on page 115

Serviceability and problem determination utilities

Table 14 lists the serviceability and problem determination utilities.

Table 14. Serviceability and problem determination utilities

Utility	Description
"pdjservicelevel"	Returns the service level of installed Security Access Manager files that use the IBM Security Access Manager Runtime for Java package.
"pdservicelevel" on page 119	Returns the service level of installed Security Access Manager files that use the Security Access Manager Runtime package.
"pdversion" on page 120	Lists the current version of Security Access Manager components that are installed on the system.

pdjservicelevel

Returns the service level of installed Security Access Manager files that use the IBM Security Access Manager Runtime for Java package.

Note: This utility is for use by support personnel.

Syntax

pdjservicelevel *directory*

Description

The **pdjservicelevel** utility recursively scans the specified directory and returns the name and service level for each file to standard output. Only executable programs, shared libraries, archives, and other such files have a service level.

If the service level for a file cannot be determined, the string `"Unknown"` is written to standard output. Generally, ASCII files and other such files do not have service levels.

Note: For this utility to determine the service level of a JAR file, the Java **jar** utility must exist in the system PATH statement. When the **jar** utility cannot be found, the service level that is reported for all JAR files is `"Unknown"`.

Parameters

directory
Specifies the fully qualified name of the directory.

Availability

This utility is installed as part of the IBM Security Access Manager Runtime for Java package. It is in one of the following default installation directories:

- On AIX, Linux, and Solaris operating systems:
  `/opt/PolicyDirector/sbin`
- On Windows operating systems:
  `C:\Program Files\Tivoli\Policy Director\sbin`

When an installation directory other than the default is selected, this utility is in the `/sbin` directory under the installation directory (for example, *installation_directory*/`sbin`).

Return codes

0 The utility completed successfully.

1 The utility failed. When a utility fails, a description of the error and an error status code in hexadecimal format is provided (for example, `0x15c3a00c`). See the Messages topics in the Knowledge Center.

pdservicelevel

Returns the service level of installed Security Access Manager files that use the Security Access Manager Runtime package.

Note: This utility is for use by support personnel.

Syntax

pdservicelevel *directory*

Description

The **pdservicelevel** utility recursively scans the specified directory and returns the name and service level for each file to standard output. Only executable programs, shared libraries, archives, and other such files have a service level.

If the service level for a file cannot be determined, the string "Unknown" is written to standard output. Generally, ASCII files and other such files do not have service levels.

Note: For this utility to determine the service level of a JAR file, the Java **jar** utility must exist in the system PATH statement. When the **jar** utility cannot be found, the service level that is reported for all JAR files is "Unknown".

Parameters

directory
Specifies the fully qualified name of the directory.

Availability

This utility is installed as part of the Security Access Manager Runtime package. It is located in one of the following default installation directories:

- On AIX, Linux, and Solaris operating systems:
  `/opt/PolicyDirector/sbin`
- On Windows operating systems:
  `C:\Program Files\Tivoli\Policy Director\sbin`

When an installation directory other than the default is selected, this utility is located in the `/sbin` directory under the installation directory (for example, *installation_directory*`/sbin`).

Return codes

0 The utility completed successfully.

1 The utility failed. When a utility fails, a description of the error and an error status code in hexadecimal format is provided (for example, `0x15c3a00c`). For more information, see the Messages topics in the Knowledge Center.

pdversion

Lists the current version of Security Access Manager components that are installed on the system.

Syntax

pdversion [–key *key1, key2...keyX*] [–separator *delimiter_character*]

Parameters

–key *key1***,** *key2***...***keyX*
Specifies the component or components of the current version. (Optional) The following are possible values of –key:
- `pdacld` – Security Access Manager Authorization Server
- `pdauthadk` – Security Access Manager Application Developer Kit

- pdjrte – Security Access Manager Runtime for Java
- pdmgr – Security Access Manager Policy Server
- pdrte – Security Access Manager Runtime
- pdsms – Security Access Manager Session Manager Server
- pdweb – Security Access Manager WebSEAL
- pdwebars – Security Access Manager Attribute Retrieval Service
- pdwebadk – Security Access Manager Web Security ADK
- pdwebrte – Security Access Manager Web Security Runtime
- tivsecutl – IBM Tivoli® Security Utilities

The version information for the various blades shows up when the blade packages are installed on the system. The following components are basic components:

- Security Access Manager Runtime
- Security Access Manager Policy Server
- Security Access ManagerWeb Portal Manager
- Security Access Manager Application Developer Kit
- Security Access Manager Authorization Server
- Security Access Manager Runtime for Java

The following components are blades:

- Security Access Manager WebSEAL
- Security Access Manager web Security Runtime
- Security Access Manager web Security ADK
- Security Access Manager Session Manager Server
- Security Access Manager Attribute Retrieval Service

–separator *delimiter_character*
Specifies the separator that is used to delimit the description of the component from its version. (Optional)

Availability

This utility is in one of the following default installation directories:

- On AIX, Linux, and Solaris operating systems:
  /opt/PolicyDirector/bin
- On Windows operating systems:
  C:\Program Files\Tivoli\Policy Director\bin

When an installation directory other than the default is selected, this utility is in the /bin directory under the installation directory (for example, *installation directory*/bin).

Return codes

0 The utility completed successfully.

1 The utility failed. When a utility fails, a description of the error and an error status code in hexadecimal format is provided (for example, 0x15c3a00c). See the Messages topics in the Knowledge Center for more information.

Chapter 18. IPMItool

If you are using the hardware appliance, you can use the IPMItool to manage the Baseboard Management Controller (BMC) module. The IPMItool cannot be used with the virtual appliance.

The IPMItool is a utility to monitor, configure, and manage devices that support the Intelligent Platform Management Interface (IPMI). IPMI is a standardized message-based hardware management interface. A hardware chip that is known as the Baseboard Management Controller (BMC), or Management Controller (MC), implements the core of IPMI.

The BMC provides key interfaces that are needed for monitoring the health of the system hardware. There are interfaces for user channels, monitoring elements (temperature, voltage, fan speed, bus errors, and so on), manually driven recovery (local or remote system resets and power on/off operations), and an interface for logging without operating system intervention for abnormal or 'out-of-range' conditions for later examination and alerting. It is important to note that the BMC is always powered ON. It contains a small processor that runs IPMI even when the main system is OFF, or the operating system has crashed. So the BMC can be configured to look at the status of local hardware from another server for secure remote monitoring and recovery (such as system reset) regardless of the status of the platform.

To access the IPMItool, follow these steps:

1. Log on to the command-line interface (CLI) as a root user.
2. Enter the `Hardware` menu.
3. Enter `ipmitool`.

   Note: The IPMItool can be run only with the hardware appliance. Attempts to enter the IPMItool from a virtual appliance result in errors that are returned.
4. Enter specific commands to manage the BMC module as needed.

   To see a list of IPMItool commands, enter `# ipmitool help`.

   You can also get help for many specific IPMItool commands by adding the word help after the command. For example, `# ipmitool channel help`.

Chapter 19. Support Information

To isolate and resolve problems with your IBM products, you can use the troubleshooting and support information. This information contains instructions for using the problem-determination resources that are provided with your IBM products.

Techniques for troubleshooting problems

Troubleshooting is a systematic approach to solving a problem. The goal of troubleshooting is to determine why something does not work as expected and how to resolve the problem. Certain common techniques can help with the task of troubleshooting.

The first step in the troubleshooting process is to describe the problem completely. Problem descriptions help you and the IBM technical-support representative know where to start to find the cause of the problem. This step includes asking yourself basic questions:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, which can then lead you to a problem resolution.

What are the symptoms of the problem?

When you start to describe a problem, the most obvious question is “What is the problem?” This question might seem straightforward; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, crash, performance degradation, or incorrect result?

Where does the problem occur?

Determining where the problem originates is not always easy, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few of the components to consider when you are investigating problems.

The following questions help you to focus on where the problem occurs to isolate the problem layer:

- Is the problem specific to one platform or operating system, or is it common across multiple platforms or operating systems?
- Is the current environment and configuration supported?

- Do all users have the problem?
- (For multi-site installations.) Do all sites have the problem?

If one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many problems can be traced back to incompatible levels of software that are not intended to run together or were not fully tested together.

When does the problem occur?

Develop a detailed timeline of events that led up to a failure, especially for those cases that are one-time occurrences. You can most easily develop a timeline by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log.

To develop a detailed timeline of events, answer these questions:
- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to these types of questions can give you a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing which systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These questions about your environment can help you to identify the root cause of the problem:
- Does the problem always occur when the same task is being completed?
- Does a certain sequence of events need to happen for the problem to occur?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might occur around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the ideal problem is one that can be reproduced. Typically, when a problem can be reproduced you have a larger set of tools or procedures at your disposal to help you investigate. Problems that you can reproduce are often easier to debug and solve.

However, problems that you can reproduce can have a disadvantage: If the problem is of significant business affect, you do not want it to recur. If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

- Can the problem be re-created on a test system?
- Are multiple users or applications that encounter the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, or a particular application?

Searching knowledge bases

You can often find solutions to problems by searching IBM knowledge bases. You can optimize your results by using available resources, support tools, and search methods.

About this task

You can find useful information by searching the IBM Knowledge Center for this product. However, sometimes you need to look beyond the IBM Knowledge Center to answer your questions or resolve problems.

Procedure

To search knowledge bases for information that you need, use one or more of the following approaches:

- Search for content by using the IBM Support Assistant (ISA).

  ISA is a no-charge software serviceability workbench that helps you answer questions and resolve problems with IBM software products. You can find instructions for downloading and installing ISA on the ISA website.
- Find the content that you need by using the IBM Support Portal.

  The IBM Support Portal is a unified, centralized view of all technical support tools and information for all IBM systems, software, and services. The IBM Support Portal lets you access the IBM electronic support portfolio from one place. You can tailor the pages to focus on the information and resources that you need for problem prevention and faster problem resolution.
- Search for content by using the IBM masthead search. You can use the IBM masthead search by typing your search string into the **Search** field at the top of any ibm.com® page.
- Search for content by using any external search engine, such as Google, Yahoo, or Bing. If you use an external search engine, your results are more likely to include information that is outside the ibm.com domain. However, sometimes you can find useful problem-solving information about IBM products in newsgroups, forums, and blogs that are not on ibm.com.

  Tip: Include “IBM” and the name of the product in your search if you are looking for information about an IBM product.

Getting fixes

A product fix might be available to resolve your problem.

About this task

Procedure

To find and install fixes:

1. Obtain the tools required to get the fix.
2. Determine which fix you need.
3. Download the fix. Open the download document and follow the link in the "Download package" section.
4. Apply the fix. Follow the instructions in the "Installation Instructions" section of the download document.
5. Subscribe to receive weekly email notifications about fixes and other IBM Support information.

Getting fixes from Fix Central

You can use Fix Central to find the fixes for IBM products. With Fix Central, you can search, select, order, and download fixes for your system with a choice of delivery options. A product fix might be available to resolve your problem.

Procedure

To find and install fixes:

1. Obtain the tools that are required to get the fix. If it is not installed, obtain your product update installer. You can download the installer from Fix Central. This site provides download, installation, and configuration instructions for the update installer.
2. Select the product, and select one or more check box that are relevant to the problem that you want to resolve.
3. Identify and select the fix that is required.
4. Download the fix.
   a. Open the download document and follow the link in the "Download Package" section.
   b. When downloading the file, ensure that the name of the maintenance file is not changed. This change might be intentional, or it might be an inadvertent change that is caused by certain web browsers or download utilities.
5. Apply the fix.
   a. Follow the instructions in the "Installation Instructions" section of the download document.
   b. For more information, see the "Installing fixes with the Update Installer" topic in the product documentation.
6. Optional: Subscribe to receive weekly email notifications about fixes and other IBM Support updates.

Contacting IBM Support

IBM Support assists with product defects, answers FAQs, and helps users resolve problems with the product.

Before you begin

After you try to find your answer or solution by using other self-help options such as technotes, you can contact IBM Support. Before you contact IBM Support, your company or organization must have an active IBM software subscription and support contract, and you must be authorized to submit problems to IBM. For information about the types of available support, see the Support portfolio topic in the "*Software Support Handbook*".

Procedure

To contact IBM Support about a problem:

1. Define the problem, gather background information, and determine the severity of the problem. For more information, see the Getting IBM support topic in the *Software Support Handbook*.
2. Gather diagnostic information.
3. Submit the problem to IBM Support in one of the following ways:
   - Using IBM Support Assistant (ISA):
   - Online through the IBM Support Portal: You can open, update, and view all of your service requests from the Service Request portlet on the Service Request page.
   - By telephone for critical, system down, or severity 1 issues: For the phone number to call in your region, see the Directory of worldwide contacts web page.

Results

If the problem that you submit is for a software defect or for missing or inaccurate documentation, IBM Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the IBM Support website daily, so that other users who experience the same problem can benefit from the same resolution.

What to do next

Exchanging information with IBM

To diagnose or identify a problem, you might need to provide IBM Support with data and information from your system. In other cases, IBM Support might provide you with tools or utilities to use for problem determination.

Sending information to IBM Support

To reduce the time that is required to resolve your problem, you can send trace and diagnostic information to IBM Support.

Procedure

To submit diagnostic information to IBM Support:

1. Open a problem management record (PMR).

2. Collect the diagnostic data that you need. Diagnostic data helps reduce the time that it takes to resolve your PMR. You can collect the diagnostic data manually or automatically:
   - Collect the data manually.
   - Collect the data automatically.
3. Compress the files by using the `.zip` or `.tar` file format.
4. Transfer the files to IBM. You can use one of the following methods to transfer the files to IBM:
   - IBM Support Assistant
   - The Service Request tool
   - Standard data upload methods: FTP, HTTP
   - Secure data upload methods: FTPS, SFTP, HTTPS
   - Email

   All of these data exchange methods are explained on the IBM Support website.

Receiving information from IBM Support

Occasionally an IBM technical-support representative might ask you to download diagnostic tools or other files. You can use FTP to download these files.

Before you begin

Ensure that your IBM technical-support representative provided you with the preferred server to use for downloading the files and the exact directory and file names to access.

Procedure

To download files from IBM Support:

1. Use FTP to connect to the site that your IBM technical-support representative provided and log in as `anonymous`. Use your email address as the password.
2. Change to the appropriate directory:
   a. Change to the `/fromibm` directory.
   ```
   cd fromibm
   ```
   b. Change to the directory that your IBM technical-support representative provided.
   ```
   cd nameofdirectory
   ```
3. Enable binary mode for your session.
   ```
   binary
   ```
4. Use the **`get`** command to download the file that your IBM technical-support representative specified.
   ```
   get filename.extension
   ```
5. End your FTP session.
   ```
   quit
   ```

Subscribing to Support updates

To stay informed of important information about the IBM products that you use, you can subscribe to updates.

About this task

By subscribing to receive updates about this product, you can receive important technical information and updates for specific IBM Support tools and resources. You can subscribe to updates by using one of two approaches:

RSS feeds and social media subscriptions

For general information about RSS, including steps for getting started and a list of RSS-enabled IBM web pages, visit the IBM Software Support RSS feeds site.

My Notifications

With My Notifications, you can subscribe to Support updates for any IBM product. (My Notifications replaces My Support, which is a similar tool that you might have used in the past.) With My Notifications, you can specify that you want to receive daily or weekly email announcements. You can specify what type of information you want to receive (such as publications, hints, and tips, product flashes (also known as alerts), downloads, and drivers). My Notifications enable you to customize and categorize the products about which you want to be informed and the delivery methods that best suit your needs.

Procedure

To subscribe to Support updates:

1. Subscribe to the RSS feeds by accessing the IBM Software Support RSS feeds site and subscribe to the product feed.
2. Subscribe to My Notifications by going to the IBM Support Portal and click **My Notifications** in the **Notifications** portlet.
3. Sign in using your IBM ID and password, and click **Submit**.
4. Identify what and how you want to receive updates.
   a. Click the **Subscribe** tab.
   b. Select the appropriate software brand or type of hardware.
   c. Select one or more products by name and click **Continue**.
   d. Select your preferences for how to receive updates, whether by email, online in a designated folder, or as an RSS or Atom feed.
   e. Select the types of documentation updates that you want to receive, for example, new information about product downloads and discussion group comments.
   f. Click **Submit**.

Results

Until you modify your RSS feeds and My Notifications preferences, you receive notifications of updates that you requested. You can modify your preferences when needed (for example, if you stop using one product and begin using another product).

Related information

IBM Software Support RSS feeds

Subscribe to My Notifications support content updates

My Notifications for IBM technical support

My Notifications for IBM technical support overview

Part 2. Advanced Access Control troubleshooting

Chapter 20. Advanced Access Control known issues and solutions

Determine if an issue you are having has a fix or a workaround.

Troubleshooting PIP server connections

To troubleshoot connection errors, ensure that runtime tracing for the PIPs is enabled.

1. In the local management interface, click **Secure Access Control** > **Global Settings** > **Runtime Parameters** > **Runtime Tracing**.
2. In the Tracing Specification field, type the tracing specifications for the PIPs:

   For JNDI (LDAP PIP)

   ```
   com.tivoli.am.rba.pip.LdapPIP=FINE
   javax.naming.*=FINE
   ```

   For JDBC (Database PIP)

   ```
   com.tivoli.am.rba.pip.JdbcPIP=FINE
   java.sql.*=FINE
   ```

   Note: A trace level of FINE provides general trace, method entry, exit, and return values. For more detailed information, consider using FINER or FINEST as the trace level.
3. Click **Save**.

Removal of a server connection defined for a PIP causes a problem

A problem with a policy information point (PIP) can be caused by:

- Deleting a server connection, especially if it is defined by a PIP and being used to return attributes in a policy or risk score.
- Updating the JNDI ID or host name of a server connection after it was already referenced in a PIP definition.
- Updating any of the server connection properties that are not valid after it was already referenced in a PIP definition.

Solution:

Review the list of server connections to determine whether any were deleted, and re-create any deleted server connections.

Ensure that the PIP specifies the correct server connection:

1. In the local management interface, click **Secure Access Control**.
2. Under **Policy**, click **Information Points**.
3. Select the server connection PIP and click **Modify**.
4. In the Connection tab, locate the **Server Connection** field and determine whether it contains the correct server connection.
5. Update the server connection, if necessary.

WebSEAL sends client certificate by default

The **isamcfg** tool can configure a WebSEAL or Web Reverse Proxy instance to use client certificate authentication to the run time. The run time can be configured to optionally accept client certificates. When configured in this mode, WebSEAL or the Web Reverse Proxy sends a client certificate regardless of whether a certificate label is specified in the configuration.

Solution:

If you require client certificate authentication, ensure that a valid client certificate is specified in the `ssl-keyfile-label` entry of the `[rtss-cluser:<cluster name>]` stanza. If you do not require client certificate authentication, either disable the optional acceptance of client certificates or specify an invalid client certificate label in the `ssl-keyfile-label` entry of the `[rtss-cluser:<cluster name>]` stanza. This method ensures that WebSEAL does not send a client certificate.

A junction error occurs when a duplicate `dn` or certificate label is detected

The **isamcfg** tool imports the Access Manager server certificate to the WebSEAL or Web Reverse Proxy keystore when it uses an SSL connection. If an entry exists in this keystore with the same `dn` or the same certificate label, an error can occur when it creates the junction to the Access Manager server.

Solution:

You must manually export the certificate that is presented by the run time and import it to the WebSEAL key database as a signer certificate. The junction then becomes accessible.

The appliance fails to connect to the authorization service endpoint

When the appliance is configured in FIPS and NIST SP800-131A compliant mode, the RBA EAS configured in the POC appliance fails to connect to the authorization service endpoint. This failure causes the RBA flow to fail. This issue also affects the ping call that is issued regularly. The connection fails because the authorization service EAS uses SSLv2, which is not supported by the appliance when it operates in the NIST SP800-131A strict compliant mode.

Solution:

1. In the appliance local management interface, select **Reverse Proxy Settings** > **<your instance>** > **Manage** > **Configuration** > **Edit configuration file**. The Advanced Configuration File Editor opens.
2. Add the parameter to the existing stanza.

   ```
   [rtss-cluster:cluster1]
   gsk-attr-name = enum:438:1
   ```
3. Click **Save**.
4. Deploy the changes.
5. Restart the instance.

Incorrectly formatted FORM or JSON data in custom attributes causes policy failure

If a policy contains a custom attribute with incorrect FORM or JSON data, the policy fails and the user is not permitted to access the resource. Specifically, all of the following conditions apply:

- The WebSEAL configuration file contains a `post-data` entry in the `[azn-decision-info]` stanza.
- The <*post-data-name*> value for the `post-data` entry is not formatted properly. For example, the date format is not correct.
- The `.datatype` entry in the `[user-attribute-definitions]` stanza for this `azn-decision-info` attribute is a type other than `string`.

Solution:

Ensure that you specify the correct format for the data. For example:

- Valid date format is: *yyyy-mm-ddzzzzzz*

  For example: `2013-05-20–06:00`
- Valid time format is: *hh*:*mm*:*sszzzzzz*

  For example: `13:12:36-06:00`

Advanced Access Control cookies are not removed when a user logs out by using non-standard junction name or cookie names

The Advanced Access Control runtime sets cookies for attribute collection purposes. The **isamcfg** tool configures a WebSEAL or Web Reverse Proxy configuration option to clean these cookies upon session termination. It uses the Advanced Access Control advanced configuration to determine which cookie values to clear.

Solution:

When the junction name `attributeCollection.serviceLocation` or cookie name `attributeCollection.cookieName` in the advanced configuration changes, you must run the **isamcfg** tool so that these changes are picked up.

Access tokens are not cleared in a failed resource owner password credential flow

The access tokens that are generated from resource owner password credential flow are not cleared when resource owner password validation fails. These tokens must be removed so that malicious users cannot use the resource owner password credential flow to fill up the token cache by using a public client.

Solution:

In the pre-mapping rule, make a validate request to the user directory with the user name and password that you want to verify. This method stops the resource owner password credential flow before it generates the access token if the user name and password verification fails.

For more information, see OAuth 2.0 mapping rule methods.

The one-time password value cannot be validated

Policies that permit one-time password obligations might result in an error if the user who made the access request did not set a secret key on the self-care secret key page.

Solution: To ensure that users who encounter this error are directed to create their secret key, edit the `error_could_not_validate_otp.html` file and add a link that opens the self-care secret key page at `/sps/mga/user/mgmt/html/otp/otp.html`. For example, add the following link:

```
<a "href="/sps/mga/user/mgmt/html/otp/otp.html">
Generate your Secret Key..</a>
```

Attributes cause errors when hashed in attributeCollection.attributesHashEnabled

In the Advanced Configuration settings, any attribute that uses a matcher other than the exact matcher causes an error if it is hashed.

Solution: Do not hash the following attributes with the `attributeCollection.attributesHashEnabled` setting of Advanced Configuration:

- `ipAddress`
- `geoLocation`
- `accessTime`

Policies with X.500 names

If your LDAP root DN is `secauthority=default`, you can use the = (equal) operator only in policies that use X.500 names `userDN` and `groupsDN`.

IP addresses are granted access despite their IP reputations

Despite their IP reputation classifications, IP addresses are granted access. When the appliance cannot resolve the domain name for the license server, the IP reputation database cannot update. When the IP reputation database cannot update, it either contains inaccurate IP reputation data or no IP reputation data.

You know that the IP reputation database cannot contact the license server when all IP addresses are granted access despite their IP reputations. If the following conditions are true, the IP reputation policy information point (PIP) might return an incorrect reputation:

- An administrator writes a policy that denies access based on IP reputation.
- The database cannot contact the license server.

Solution: Configure direct access to the license server or indirect access to the license server. See the Administrating topics about the license server for more information.

Error message is returned when IBM solidDB is used to deploy an external runtime database

When you deploy an external runtime database with IBM solidDB, `isam_access_control_soliddb.sql` attempts to create a duplicate index on the `RBA_USER_DEVICE` table. The following error message is returned:

```
SOLID Table Error 13199: Duplicate index definition
```

Solution: Ignore this error message. Investigate any other error messages about the deployment of the run time database.

Authentication alias message in startup log

The following message might be displayed in the runtime server startup log:

```
I               J2CA8050I: An authentication alias should be used instead
of defining a user name               and password on
com.ibm.ws.jdbc.dataSource-config/properties-0
```

Solution: Ignore this message.

Creating access control policies during switch between daylight saving time and standard time

If you create an access control policy when the appliance clock switches from daylight saving time (DST) to standard time (ST), the policy might not work as you expect. For example:

1. Create an access control policy at 1:45 am on the Sunday when DST returns to ST.
2. Modify the policy at 2:15 am DST, which is 1:15 am ST and within 1 hour of the creating the policy. The policy that you created in step 1 is used because it is newer (more recently created) than the modified policy in step 2.

Solution: To correct this issue, modify and republish the policy that you want to use after 1:45 am standard time.

Filter stops working after changing a parameter on the Advanced Configuration panel

You can filter the data displayed on the Advanced Configuration panel. After you change a parameter and click the **Change** button, the filter is no longer applied to the displayed data.

Solution: Click **Enter** next to the filter field to reapply the filter.

Reverse proxy user password page is inaccessible

After authenticating with the authentication mechanisms, the reverse proxy user password change page (/pkmspasswd) becomes inaccessible.

Solution: This page is working as designed. The reverse proxy makes this page inaccessible for users who are authenticated with the External Authentication Interface (EAI). The authentication service relies on EAI to establish the authenticated session.

Database failover capabilities vary during a cluster upgrade

The distributed session cache, runtime database, and configuration database have different failover capabilities during cluster upgrades.

Table 15. Database capabilities

Database	Behaviour
Distributed session cache Runtime database	If the primary master fails, failover goes to secondary master. Changes are done in the secondary master and reconciliation occurs when primary master is restored.
Configuration database	If the primary master fails, there is no failover to the secondary master. No changes are possible on the primary or secondary master until the primary master is back online.

Issue 1

When a high availability cluster is active, a situation exists during the firmware upgrade on the primary node where the configuration database is read-only.

A read-only database prevents the upgrade process from writing to the configuration database when creating new tables, modifying schema of existing tables, and inserting or updating rows on tables.

The reason the database on the primary node becomes read-only is the firmware upgrade requires the appliance to be rebooted. During a reboot, the high availability controller switches the secondary master to be read-write and act as the temporary primary master.

When the primary node reboots and the database starts:

1. It recognizes that the secondary node is in control and starts in read-only mode.
2. The appliance cluster manager includes a background thread which will eventually switch the primary node to resume its role as the primary master database.
3. The database on the primary node becomes writeable.

However, during an upgrade, the database upgrade scripts are executed before the primary database has become writeable.

Issue 2

For the distributed session cache and runtime database, a situation exists where changes to the secondary databases are not reflected in the primary database after the completion of a cluster upgrade.

Solution:

To address Issue 1, make the cluster a single master cluster for the duration of the firmware upgrade. For detailed instructions, see the Use the local management interface for a cluster of appliances section in Upgrading to the current version.

To address Issue 2, stop traffic to the cluster before starting the upgrade.

Error occurs after switching the runtime database in the appliance

After you switch the runtime database from local to remote and deploying the pending changes, SQL-related runtime errors occur in the appliance.

Solution: Restart the runtime.

1. In the local management interface, select **Secure Access Control** > **Global Settings** > **Runtime Parameters** > **Runtime Status**.
2. Click **Restart All Clustered Runtimes**.

Error code does not display for type mismatch of RESTful PIP

The `message.log` contains an error message without an error code when the response type does not match the RESTful web service type. In the message log, you might see messages without an error code from `com.tivoli.am.rba.pip.RestPIP`.

This problem occurs during RESTful PIP configuration if:

- You manually type a response type instead of selecting a predefined type.
- The type that you entered does not match the type received from the RESTful web service.

Solution:

- Correct the RESTful PIP configuration in the local management interface.
- Verify that the response type you specify is correct for the PIP.
- Whenever possible, select a predefined type from the list instead of typing the response type.

Firmware version and Last update information not displayed in Update History and Overview section

Firmware information is not displayed in the Update History section of the user interface. Information on Last update is not displayed in the Overview section of the user interface. This is the case for firmware upgrades and fixpacks.

The firmware version can be determined using the Firmware Settings user interface on the appliance.

To determine the firmware version:

1. Logon to the appliance as the administrator.
2. Click **Manage System Settings** > **Firmware Settings**. The active partition is reported. The Details column includes the firmware version.

Alternatively, you can verify the firmware version from **Manage System Settings** > **Updates and Licensing** > **Overview**.

Information on Last update not displaying is a known limitation.

Default component port number usage

When you assign port numbers during installation or configuration, do not use port numbers already assigned to other components. You might see unpredictable behavior if the same port number is assigned to multiple components.

This is important when there are two Security Access Manager appliances and one of them has Advanced Access Control.

Solution: Consider the following situation:

- If you set up external WebSEAL servers as part of the cluster configuration, do not assign port 9080 to the cluster configuration if the Support internal and external clients option is selected. Port 9080 is the default port assigned to WebSphere Application Server.

The following table lists the default port numbers that you must be aware of when configuring components on a Security Access Manager appliance that also has Advanced Access Control. Do not assign them to other components.

Table 16. Default port numbers for Access Manager components

Component	**Port number**
WebSphere Application Server	9080, 9443
Security Access Manager Policy Server	7135
Security Access Manager Authorization Server	7136, 7137
WebSEAL listening port	7234
Cluster configuration	2020-2050
LDAP server, SSL port	636
Remote syslog	514
Local management interface (LMI)	443
WebSEAL HTTPS	443
LDAP server, non-SSL port	389
WebSEAL HTTP	80

Deploy request fails

A call to the Deploy the pending configuration changes web service (`pending_changes/deploy`) might fail.

A call to the Deploy the pending configuration changes (`pending_changes/deploy`) web service returns an HTTP Internal Server Error.

Solution: Retry the operation.

Runtime server unable to obtain federated directory information from the local management interface

If the runtime server cannot contact the local management interface (LMI) under the following configuration settings, federated directory users cannot authenticate:

- Federated directories are used for username and password authentication.
- Resource owner password credentials (ROPC) are used for API protection.

A user whose information is housed in a federated directory cannot authenticate through the username and password authentication mechanism or access API protection definitions that use resource owner password credentials.

The runtime server attempts to contact the local management interface one time to obtain all federated directory information to authenticate federated directory users. If the runtime server cannot contact the local management interface for any reason, it cannot obtain the required federated directory information to authenticate federated directory users. For example: If the local management interface server is not running, the runtime server cannot obtain the necessary federated directory information to authenticate users in a federated directory.

Solution: Follow these steps to ensure that this issue is resolved:

1. Restart the local management interface:
   a. Use an ssh session to access the local management interface.
   b. Log in as the administrator.
   c. Type `lmi` and press Enter.
   d. Type `restart` and press Enter.
   e. Type `exit` and press Enter.
2. Restart the runtime server:
   a. Log in to the local management interface.
   b. Select **Secure Access Control** > **Global Settings** > **Runtime Parameters** > **Runtime Status**.
   c. Click **Restart Local Runtime**.
3. Ensure that the runtime server can communicate with the local management interface so that it can obtain all of the necessary federated directory information.

Chapter 21. Advanced Access Control known limitations

Consider these known limitations when you are configuring an Advanced Access Control environment on the appliance.

External clients cannot use the session cache

The distributed session cache in the Advanced Access Control does not support external clients.

The **Support internal and external clients** option on the Session Cache tab on the Cluster Configuration management page is not relevant in an Advanced Access Control environment.

Advanced Access Control disregards the **Port**, **Keyfile**, and **Label** fields, which relate to external clients.

Descriptions of default attribute and obligations might not display in the correct language

If you clear your browser cache while logged into an appliance session, you might not see the descriptions of default attributes and obligations in the correct language. This scenario happens when you perform steps similar to these:

1. Log in to the appliance.
2. Change the language of the local management interface.
3. Clear the browser cache.
4. Display obligations or attributes. For example, to display the attributes:
   a. Select **Secure Access Control**.
   b. Under **Policy**, select **Attributes**. Under the name of each default attribute is the description. This description might display in an incorrect language.

Therefore, do not clear the browser cache during an appliance session because you might see an incorrect language displayed in this scenario.

Certain characters in JSON messages are displayed in Unicode

Non-ASCII characters are escaped in the JSON response from the REST API endpoints. This format is specified in RFC 4627.

The non-ASCII character is represented as a six-character sequence: a reverse solidus, followed by the lowercase letter u, and followed by four hexadecimal digits that encode the code point of the character. For example, `\u00e9`. For more information, see RFC 4627.

The Quick Response (QR) Code generator in Advanced Access Control only accepts US-ASCII alphanumeric characters as valid inputs

Advanced Access Control can display the OAuth 2.0 authorization code as a QR code image.

The QR code endpoint creates the QR code image. The endpoint is designed to accept US-ASCII alphanumeric characters only. This is to ensure maximum interoperability with existing QR code scanners.

Ensure that only US-ASCII alphanumeric characters are used to create the QR code image.

Authentication service cannot use the group information in the credential

You can create a custom authentication mechanism by using the authentication mechanism Software Development Kit. Aside from authenticating the user, the authentication mechanism can modify the credential of the current user.

After the user completes the execution of the authentication policy, which contains your custom authentication mechanism, the authentication service logs in the current user to IBM Security Access Manager by using the resulting credential. Advanced Access Control has a limitation that the group information in the resulting credential is not used by the authentication service to log in the current user to IBM Security Access Manager.

Chapter 22. Deploying pending changes

Some configuration and administration changes require an extra deployment step.

About this task

When you use the graphical user interface on the appliance to specify changes, some configuration and administration tasks take effect immediately. Other tasks require a deployment step to take effect. For these tasks, the appliance gives you a choice of deploying immediately or deploying later. When you must make multiple changes, you can wait until all changes are complete, and then deploy all of them at one time.

When a deployment step is required, the user interface presents a message that says that there is an undeployed change. The number of pending changes is displayed in the message, and increments for each change you make.

Note: If any of the changes require the runtime server to be restarted, the restart occurs automatically when you select **Deploy**. The runtime server will then be unavailable for a period of time until the restart completes.

Procedure

1. When you finish making configuration changes, select **Click here to review the changes or apply them to the system**.

   The Deploy Pending Changes window is displayed.
2. Select one of the following options:

Option	Description
Cancel	Do not deploy the changes now. Retain the undeployed configuration changes. The appliance user interface returns to the previous panel.
Roll Back	Abandon configuration changes. A message is displayed, stating that the pending changes were reverted. The appliance user interface returns to the previous panel.
Deploy	Deploy all configuration changes. When you select **Deploy**, a system message is displayed, stating that the changes were deployed. If any of the changes require the runtime server to be restarted, the restart occurs automatically when you select **Deploy**. The runtime server will then be unavailable for a period of time until the restart completes.

Part 3. Appendixes

Index

Printed in USA