

IBM HTTP Server for WebSphere Application Server,
Version 8.0

*IBM HTTP Server for WebSphere
Application Server*



Note

Before using this information, be sure to read the general information under “Notices” on page 219.

Compilation date: June 7, 2011

© Copyright IBM Corporation 2011.

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

Contents

How to send your comments	vii
Changes to serve you more quickly	ix
Chapter 1. Administering and configuring IBM HTTP Server	1
Performing required z/OS system configurations	1
Starting and stopping IBM HTTP Server	4
Using the administrative console to start IBM HTTP Server	5
Using apachectl commands to start IBM HTTP Server	6
Using Windows services to start IBM HTTP Server	7
Using JCL procedures to start IBM HTTP Server on z/OS	7
Configuring IBM HTTP Server	9
Apache modules (containing directives) supported by IBM HTTP Server	9
Apache programs supported by IBM HTTP Server	14
Apache APR and APR-util libraries supported by IBM HTTP Server	15
Apache MPM and addressing modes supported by IBM HTTP Server	16
IPv4 and IPv6 configuration for Windows operating systems	16
Serving static content faster with Fast Response Cache Accelerator	16
Customizing Fast Response Cache Accelerator logging	17
Restrictions on cached content	18
Fast Response Cache Accelerator operational restrictions	19
Servlets and JavaServer Pages files caching	19
AIX considerations for Fast Response Cache Accelerator (FRCA)	20
AFPA directives	20
Enabling IBM HTTP Server for FastCGI applications	23
Learn about FastCGI	24
FastCGI directives	24
Managing IBM HTTP Server remotely with the WebSphere Application Server administrative console	33
Extending IBM HTTP Server functionality with third-party plug-in modules	34
Viable compilers for Apache and third-party plug-in modules	34
Build method options for dynamic modules	35
Considerations for building dynamic modules on Windows platforms	35
Chapter 2. Administering and configuring the administration server	37
Starting and stopping the IBM HTTP Server administration server	37
Protecting access to the IBM HTTP Server administration server	38
Enabling access to the administration server using the htpasswd utility	38
Running the setupadm command for the administration server	38
Setting permissions manually for the administration server	40
Chapter 3. Migrating and installing IBM HTTP Server	43
Installing, updating, rolling back, and uninstalling IBM HTTP Server	43
Installing IBM HTTP Server using the GUI	45
Running multiple instances of IBM HTTP Server from a single install	53
Updating IBM HTTP Server	55
Installing IBM HTTP Server silently	55
Uninstalling IBM HTTP Server using the GUI	65
Rolling back IBM HTTP Server	70
Migrating and installing IBM HTTP Server on z/OS systems	71
Installing IBM HTTP Server for WebSphere Application Server for z/OS	72
Configuring an instance of IBM HTTP Server on the z/OS system	78
Uninstalling IBM HTTP Server for z/OS	82

Chapter 4. Product overview and quick start	85
What is new in this release	85
Key differences from the Apache HTTP Server	86
Chapter 5. Securing IBM HTTP Server	87
Securing IBM HTTP Server	87
Configure SSL between the IBM HTTP Server Administration Server and the deployment manager	87
Securing with SSL communications	90
Secure Sockets Layer (SSL) protocol	93
SSL directive considerations	97
Authentication	98
Encryption	99
Secure Sockets Layer environment variables	99
SSL directives	103
Setting advanced SSL options	123
Choosing the level of client authentication	123
Choosing the type of client authentication protection	124
Defining SSL for multiple-IP virtual hosts	129
Setting up a reverse proxy configuration with SSL	129
IBM HTTP Server certificate management	129
Managing keys with the IKEYMAN graphical interface (Distributed systems)	132
Starting the Key Management utility user interface	133
Working with key databases	133
Changing the database password	134
Creating a new key pair and certificate request	135
Importing and exporting keys	136
Listing certificate authorities	137
Certificate expiration dates	138
Creating a self-signed certificate	138
Receiving a signed certificate from a certificate authority	139
Displaying default keys and certificate authorities	140
Storing a certificate authority certificate	140
Storing the encrypted database password in a stash file	141
Managing keys with the gskcmd command line interface (Distributed systems)	141
Using the gskcmd command	142
Key Management Utility command-line interface (gskcmd) syntax	143
Creating a new key database using the command-line interface	145
Managing the database password using the command line	146
Creating a new key pair and certificate request	147
Importing and exporting keys using the command line	148
Creating a self-signed certificate	150
Receiving a signed certificate from a certificate authority	151
Displaying default keys and certificate authorities	152
Storing a certificate authority certificate	153
Storing the encrypted database password in a stash file	153
Managing keys with the native key database gskkyman (z/OS systems)	153
Getting started with the cryptographic hardware for SSL (Distributed systems)	154
Cryptographic hardware for Secure Sockets Layer	154
Initializing IBM 4758 and IBM e-business Cryptographic Accelerator on AIX systems	156
Initializing IBM 4758 Cryptographic Accelerator on Windows systems	156
Using IKEYMAN to store keys on a PKCS11 device	156
Configuring IBM HTTP Server to use nCipher and Rainbow accelerator devices and PKCS11 devices	158
Authenticating with LDAP on IBM HTTP Server using mod_ibm_ldap (Distributed systems)	159
Lightweight Directory Access Protocol	162
Querying the Lightweight Directory Access Protocol server	163

Secure Sockets Layer and the Lightweight Directory Access Protocol module	163
SSL certificate revocation list	164
LDAP directives	165
Converting your directives from mod_ibm_ldap to mod_ldap	175
Authenticating with LDAP on IBM HTTP Server using mod_ldap	185
Authenticating with SAF on IBM HTTP Server (z/OS systems)	187
SAF directives	188
Chapter 6. Troubleshooting and support: IBM HTTP Server	195
Troubleshooting IBM HTTP Server.	195
Known problems on Windows platforms	195
Known problems on z/OS platforms	196
Known problems with hardware cryptographic support on AIX.	198
Symptoms of poor server response time	198
Hints and tips for managing IBM HTTP Server using the administrative console	198
Could not connect to IBM HTTP Server administration server error.	199
Experiencing an IBM HTTP Server Service logon failure on Windows operating systems	200
Viewing error messages for a target server that fails to start	201
Cache messages	201
Configuration messages	202
Handshake messages	203
SSL initialization messages	210
I/O error messages	217
SSL stash utility messages	218
Notices	219
Trademarks and service marks	221
Index	223

How to send your comments

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

- To send comments on articles in the WebSphere Application Server Information Center
 1. Display the article in your Web browser and scroll to the end of the article.
 2. Click on the **Feedback** link at the bottom of the article, and a separate window containing an e-mail form appears.
 3. Fill out the e-mail form as instructed, and click on **Submit feedback** .
- To send comments on PDF books, you can e-mail your comments to: **wasdoc@us.ibm.com** or fax them to 919-254-5250.

Be sure to include the document name and number, the WebSphere Application Server version you are using, and, if applicable, the specific page, table, or figure number on which you are commenting.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Changes to serve you more quickly

Print sections directly from the information center navigation

PDF books are provided as a convenience format for easy printing, reading, and offline use. The information center is the official delivery format for IBM WebSphere Application Server documentation. If you use the PDF books primarily for convenient printing, it is now easier to print various parts of the information center as needed, quickly and directly from the information center navigation tree.

To print a section of the information center navigation:

1. Hover your cursor over an entry in the information center navigation until the **Open Quick Menu** icon is displayed beside the entry.
2. Right-click the icon to display a menu for printing or searching your selected section of the navigation tree.
3. If you select **Print this topic and subtopics** from the menu, the selected section is launched in a separate browser window as one HTML file. The HTML file includes each of the topics in the section, with a table of contents at the top.
4. Print the HTML file.

For performance reasons, the number of topics you can print at one time is limited. You are notified if your selection contains too many topics. If the current limit is too restrictive, use the feedback link to suggest a preferable limit. The feedback link is available at the end of most information center pages.

Under construction!

The Information Development Team for IBM WebSphere Application Server is changing its PDF book delivery strategy to respond better to user needs. The intention is to deliver the content to you in PDF format more frequently. During a temporary transition phase, you might experience broken links. During the transition phase, expect the following link behavior:

- Links to Web addresses beginning with `http://` work
- Links that refer to specific page numbers within the same PDF book work
- The remaining links will *not* work. You receive an error message when you click them

Thanks for your patience, in the short term, to facilitate the transition to more frequent PDF book updates.

Chapter 1. Administering and configuring IBM HTTP Server

Learn how to administer and configure IBM® HTTP Server, including: Secure Socket Layer (SSL), Key management, Lightweight Directory Access Protocol (LDAP) and System Authorization Facility (SAF) for z/OS® systems

Performing required z/OS system configurations

Before starting IBM HTTP Server, there are required z/OS system configurations that you must set up.

About this task

In order to run IBM HTTP Server, you must set the following z/OS system configurations:

Procedure

- **Set the MEMLIMIT parameter.** The MEMLIMIT parameter controls the amount of virtual memory above 2 gigabytes for a particular address space. The default setting for MEMLIMIT is 0. However, all binary programs provided with IBM HTTP Server are 64-bit applications, and these applications will not be operational with the default setting for MEMLIMIT.

The MEMLIMIT parameter can be set:

- In the OMVS segment of the user ID used to run the server:
`ALTUSER WWWSERV OMVS(MEMLIMIT(512M))`
- In the parmlib member SMFPRMxx. Setting the parmlib member SMFPRMxx will establish the system-wide MEMLIMIT default.

For a complete description of how to set MEMLIMIT, refer to the section "Limiting the use of memory objects" in *z/OS MVS™ Programming Extended Addressability Guide (SA22-7614)*. You can link to this document from the *z/OS Internet Library*.

IBM HTTP Server requires approximately 5.4 megabytes of 64-bit virtual memory per thread. The minimum recommended MEMLIMIT setting for proper IBM HTTP Server operation is: $6 * (\text{ThreadsPerChild} + 3)$ megabytes.

- **Configure a mechanism for allowing access to low ports.** The Web server user ID must have access to the TCP ports on which it will handle client connections. If port values less than 1024 are used, such as Web server ports 80 and 443, special configuration is required to allow the Web server to bind to the port.

You can use one of the following mechanisms to allow access to low ports:

- Set the PORT directive in the TCP/IP configuration.
- Disable RESTRICTLOWPORTS in the TCP/IP configuration.
- Code the Web server job name on a PORT statement in the TCP/IP configuration.
- Code a wildcard for the job name on a PORT statement in the TCP/IP configuration.
- Code SAF and a safname value on the PORT statement in the TCP/IP configuration, and permit the Web server user ID read access to the SAF FACILITY class profile `EZB.PORTACCESS.sysname.stackname.safname`.

For more information on configuration methods for allowing access to low ports, refer to the sections "Port access control" and "Setting up reserved port number definitions in PROFILE.TCPIP" in *z/OS Communications Server IP Configuration Guide (SC31-8775)*. You can link to this document from the *z/OS Internet Library*.

For an explanation of how Unix System Services jobnames (such as those for IBM HTTP Server instances) are determined, refer to the section "Generating jobnames for OMVS address spaces" in *z/OS UNIX System Services Planning (GA22-7800)*. Link to this document from the *z/OS Internet Library*.

- **Required System Authorization Facility (SAF) configurations.**

– **Create a user ID and group for IBM HTTP Server.**

You can use a new or existing user ID. It must have an OMVS segment and the UID cannot be zero. The following example contains RACF® commands to create a new user and group.

Password example

```
ADDGROUP WWWGROUP OMVS(GID(999))
ADDUSER WWWSERV DFLTGRP(WWWGROUP) OMVS(UID(999)) PASSWORD(password)
```

Password phrase example

```
ADDGROUP WWWGROUP OMVS(GID(999))
ADDUSER WWWSERV DFLTGRP(WWWGROUP) OMVS(UID(999)) PHRASE('my0users@99#701_workgroup')
```

The security administrator should define the password for the Web server user ID, instead of allowing it to default, to prevent an unauthorized user from being able to log in with that user ID. The ALTUSER command can be used to modify the password of an existing user ID.

Note: If you use a JCL cataloged procedure to start an IBM HTTP Server instance, create a SAF STARTED profile to assign the server user ID and group ID to the server started task. For example, to use a cataloged procedure named WEBSRV1:

```
RDEFINE STARTED WEBSRV1.* STDATA(USER(WWWSERV) GROUP(WWWGROUP) TRACE(YES))
```

– **Set program control for required MVS data sets.**

Ensure that program control is turned on for the following MVS data sets. For *hlq*, enter the high level qualifier for your system installation, for example: SYS1.LINKLIB.

- *hlq*.LINKLIB
- *hlq*.SCEERUN
- *hlq*.SCEERUN2
- *hlq*.SCLBDLL

The following example shows how to turn on program control using RACF commands. If you are using another security product, refer to that product's documentation for instructions. If you are turning on program control for the first time, you should use RDEFINE statements instead of RALTER statements:

```
RALTER PROGRAM * ADDMEM('hlq.LINKLIB'//NOPADCHK) UACC(READ)
RALTER PROGRAM * ADDMEM('hlq.SCEERUN'//NOPADCHK) UACC(READ)
RALTER PROGRAM * ADDMEM('hlq.SCLBDLL') UACC(READ)
SETROPTS WHEN(PROGRAM) REFRESH
```

In this example, an asterisk (*) is used to specify all programs in the data set.

– **Set program control for HFS files.**

The SMP/E installation logic enables the program control bit for the provided libraries and executable files that need it. If you install custom plug-in modules, use the extattr command to enable the APF and Program Control flags. For example:

```
# extattr +ap /opt/IBM/HTTPServer/modules/mod_jauth.so
```

In this example, substitute the IBM HTTP Server installation location for /opt/IBM/HTTPServer/. (You can build custom plug-in modules using the apxs script that is provided.)

– **Set program control for z/OS System SSL.**

If you set up your IBM HTTP Server to provide secure communications over the Internet, IBM HTTP Server uses z/OS System Secure Sockets Layer (SSL) to establish the secure connections. Before IBM HTTP Server can use System SSL, you must:

- Add the System SSL load library (*hlq*.SIEALNKE) to the system link list or to the STEPLIB DD concatenation in the HTTP Server cataloged procedure
- Set program control *hlq*.SIEALNKE in RACF.

The variable *hlq* is the high level qualifier for your system installation, for example: SYS1.SIEALNKE.

To turn on program control using RACF, issue the following command:

```
RALTER PROGRAM * ADDMEM('h1q.SIEALNKE'//NOPADCHK) UACC(READ)
SETRROPTS WHEN(PROGRAM) REFRESH
```

If you are turning on program control for the first time, use the RDEFINE statements instead of the RALTER statements. If you are using another security product, refer to that product's documentation for instructions.

– **Access to SAF key rings.**

The SSL and LDAP authentication support can optionally use certificates stored in SAF key rings. This requires that the Web server user ID have certain SAF permissions. Specifically, the Web server user ID must be permitted to the IRR.DIGTCERT.LISTRING facility in order to use key rings. Here are the general steps required:

1. Define the IRR.DIGTCERT.LIST and IRR.DIGTCERT.LISTRING resources with universal access of None.
2. Permit the Web server user ID read access to the IRR.DIGTCERT.LIST and IRR.DIGTCERT.LISTRING resources in the FACILITY class.
3. Activate the FACILITY general resource class.
4. Refresh the FACILITY general resource class.

The following commands are RACF commands. Replace WWWSERV with the actual user ID under which IBM HTTP Server is started.

```
RDEFINE FACILITY IRR.DIGTCERT.LIST UACC(NONE)
PE IRR.DIGTCERT.LIST CLASS(FACILITY) ID(WWWSERV) ACCESS(READ)
RDEFINE FACILITY IRR.DIGTCERT.LISTRING UACC(NONE)
PE IRR.DIGTCERT.LISTRING CLASS(FACILITY) ID(WWWSERV) ACCESS(READ)
SETR CLASSACT(FACILITY)
SETR RACLIST(FACILITY) REFRESH
```

For a complete guide to RACF commands, refer to *z/OS Security Server RACF Security Administrator's Guide* (SA22-7683). You can link to this document from the *z/OS Internet Library*.

– **Permitting user IDs to CSFSERV for hardware encryption:**

Integrated Cryptographic Services Facility (ICSF) is the software interface to the cryptographic hardware. If you plan to run IBM HTTP Server with cryptographic hardware capability, you can restrict the use of ICSF services. To restrict the use of ICSF services, you can permit user IDs to certain profiles in the CSFSERV general resource class. CSFSERV controls the use of ICSF software. If you have defined your IBM HTTP Server to execute with a nonzero user ID, you can give the nonzero user ID READ access to CSFSERV. If you are using a security product other than RACF, refer to that product's documentation for instructions.

If you want to restrict the use of ICSF services, issue RACF commands similar to the commands in the following examples. If you have applications other than IBM HTTP Server that are using ICSF, you must customize the examples. Otherwise, the other applications will no longer have access to ICSF services.

The following example shows how to permit the WWWSERV ID and the PUBLIC ID access to profiles in CSFSERV.

```
SETRROPTS RACLIST(CSFSERV) GENERIC(CSFSERV)
RDEFINE CSFSERV CSF* UACC(NONE)
PERMIT CSF%*C CLASS(CSFSERV) ID(WWWSERV PUBLIC) ACCESS(READ)
PERMIT CSF%*PK CLASS(CSFSERV) ID(WWWSERV PUBLIC) ACCESS(READ)
PERMIT CSF%*CK CLASS(CSFSERV) ID(WWWSERV PUBLIC) ACCESS(READ)
SETRROPTS CLASSACT(CSFSERV)
SETRROPTS RACLIST(CSFSERV) GENERIC(CSFSERV) REFRESH
```

The following example shows how to give user IDs and the WWWSERV ID access to profiles in CSFSERV.

```
SETROPTS RACLIST(CSFSERV) GENERIC(CSFSERV)
RDEFINE CSFSERV CSF%%C UACC(READ)
RDEFINE CSFSERV CSFCK% UACC(READ)
RDEFINE CSFSERV CSFCK% UACC(READ)
SETROPTS CLASSACT(CSFSERV)
SETROPTS RACLIST(CSFSERV) GENERIC(CSFSERV) REFRESH
```

– **Using cryptographic hardware for key storage (optional):**

To perform key storage on cryptographic devices refer to the section "Integrated Cryptographic Service Facility (ICSF) Considerations" in *z/OS Security Server RACF Security Administrator's Guide* (SA22-7683).

For information on ICSF options refer to the section "Using Hardware Cryptographic Features with System SSL" in *z/OS Cryptographic Services System Secure Sockets Layer (SSL) Programming* (SC24-5901).

You can link to both of these documents from the *z/OS Internet Library*.

• **Setting environment variable * _BPX_JOBNAME (optional):**

IBM HTTP Server provides the file <installroot>/bin/envvars for setting environment variables for the httpd processes. You can set the environmental variable * _BPX_JOBNAME to give the server a distinct jobname. This allows you to:

- See the server in MVS operator commands and System Display and Search Facility (SDSF).
- Categorize the server in workload management (WLM) to give web traffic adequate priority.
- Use syslogd isolation for the server.
- Use PORT statements in the TCP/IP configuration that select by job name.

A typical setting is: export _BPX_JOBNAME=HTTPD. The default is to append an incrementing integer to your jobname, such as HTTPD1, HTTPD2, HTTPD3. For more information refer to the section "Generating jobnames for OMVS address spaces" in *z/OS UNIX System Services Planning* (GA22-7800). Link to this document from the *z/OS Internet Library*.

If you use the _BPX_JOBNAME variable to set the jobname, the user ID which you use to run the server must have read access to the SAF FACILITY profile BPX.JOBNAME. For example:

```
RDEFINE FACILITY BPX.JOBNAME UACC(NONE)
SETROPTS RACLIST(FACILITY) REFRESH
PERMIT BPX.JOBNAME CLASS(FACILITY) ACCESS(READ) ID(WWSERV)
SETROPTS RACLIST(FACILITY) REFRESH
RLIST FACILITY BPX.JOBNAME ALL
```

For more information refer to the section "Setting up the BPX.* FACILITY class profiles" in *z/OS UNIX System Services Planning* (GA22-7800). Link to this document from the *z/OS Internet Library*.

Starting and stopping IBM HTTP Server

You can start or stop IBM HTTP Server using the WebSphere® Application Server administrative console or using other methods depending on your platform.

Before you begin

For installation information, refer to:

- **Distributed operating systems** "Installing, updating, rolling back, and uninstalling IBM HTTP Server" on page 43
- **z/OS** "Migrating and installing IBM HTTP Server on z/OS systems" on page 71

Important: **z/OS** Before starting IBM HTTP Server, there are required z/OS system configurations that you must perform.

About this task

You can choose the following methods to start and stop IBM HTTP Server:

- 4 IBM HTTP Server for WebSphere Application Server

Procedure

- Using the WebSphere Application Server administrative console
- **AIX** **HP-UX** **Linux** **Solaris** **z/OS** Using the command line interface
- **Windows** Using the Windows service
- **z/OS** Using JCL procedures from the system console

Results

IBM HTTP Server starts successfully.

Using the administrative console to start IBM HTTP Server

You can use the WebSphere Application Server administrative console to start and stop IBM HTTP Server.

About this task

Distributed operating systems You can administer IBM HTTP Server through the WebSphere Application Server administrative console using the WebSphere Application Server node agent or using the IBM HTTP Server administration server. An IBM HTTP Server that is defined to a deployment manager (dmgr) managed node is administered using the node agent. An IBM HTTP Server that is defined to an unmanaged node is administered using the administration server.

z/OS You can administer IBM HTTP Server through the WebSphere Application Server administrative console using the WebSphere Application Server node agent. In order to enable the WebSphere Application Server administrative console for administering IBM HTTP Server, you need to specify the `-admin` option during installation of IBM HTTP Server. Or, if you already installed IBM HTTP Server without specifying the `-admin` option, you can run the `bin/enable_admin` script. For more information about enabling the administrative console for administering IBM HTTP Server on z/OS, see “Configuring an instance of IBM HTTP Server on the z/OS system” on page 78.

Important: You must start the IBM HTTP Server administration server with the same user ID that you used to start IBM HTTP Server. Also, the user ID that you used to start the IBM HTTP Server administration server must be the same as defined on the **Admin.conf** directive:

- User <admin>
- Group <admgroup>

Procedure

1. Launch the WebSphere administrative console.
2. Click **Servers > Web servers**.
3. Select your server by clicking the check box.
4. Click **Start**.
5. You can stop IBM HTTP Server by clicking **Stop**.

Results

To confirm that IBM HTTP Server started successfully, open a browser and type in your server name in the URL box.

If you are going to run Application Response Measurement (ARM) agents, make sure you have the authority to run ARM agents when you start IBM HTTP Server.

What to do next

You can configure your server for:

-
- Secure Sockets Layer (SSL)
- Lightweight Directory Access Protocol (LDAP)
- **AIX** **Windows** Fast Response Cache Accelerator (FRCA)

Using apachectl commands to start IBM HTTP Server

This topic describes how to start and stop IBM HTTP Server using the **apachectl** commands.

About this task

To start and stop IBM HTTP Server, use the **apachectl** command.

The **apachectl** command is located in the `bin` subdirectory within the IBM HTTP Server installation directory. If that directory is not in your `PATH`, the full path should be given on the command line.

z/OS Log on as the Web server user ID. This user ID must have an OMVS segment defined and a UID which is not zero. Verify that both the IBM HTTP Server product directory and the installation directory for the server instance are mounted and available.

Procedure

- **Starting and stopping IBM HTTP Server using the default configuration file.**

To start IBM HTTP Server using the default `httpd.conf` configuration file, run the `apachectl start` command.

To stop IBM HTTP Server using the default `httpd.conf` configuration file, run the `apachectl stop` command.

AIX **HP-UX** **Linux** **Solaris** Issue the commands from the default installation directories, based on your operating system.

- **AIX** `/usr/IBM/HTTPServer/bin/apachectl start|stop`
- **HP-UX** `/opt/IBM/HTTPServer/bin/apachectl start|stop`
- **Linux** `/opt/IBM/HTTPServer/bin/apachectl start|stop`
- **Solaris** `/opt/IBM/HTTPServer/bin/apachectl start|stop`

z/OS Issue the commands from the installation directory of the IBM HTTP Server instance.

- `<IHS_install_dir>/bin/apachectl start|stop`

For example, if the `apachectl` command is not in your `PATH`, the IBM HTTP Server installation directory is `/usr/IBM/HTTPServer`, and the default configuration file is used:

```
# /usr/IBM/HTTPServer/bin/apachectl start
# /usr/IBM/HTTPServer/bin/apachectl stop
```

- **Starting and stopping IBM HTTP Server using an alternate configuration file.**

To start IBM HTTP Server using an alternate configuration file, run the following command:

- `apachectl -k start -f <path_to_configuration_file>`

To stop IBM HTTP Server using an alternate configuration file, run the following command:

- `apachectl -k stop -f <path_to_configuration_file>`

For example, the `apachectl` command is not in your `PATH`, the IBM HTTP Server installation directory is `/opt/IBM/HTTPServer`, and an alternate configuration file, `/opt/IBM/HTTPServer/conf/nodeb.conf`, is used:

```
# /opt/IBM/HTTPServer/bin/apachectl -k start -f /opt/IBM/HTTPServer/conf/nodeb.conf
# /opt/IBM/HTTPServer/bin/apachectl -k stop -f /opt/IBM/HTTPServer/conf/nodeb.conf
```



Results

To confirm that IBM HTTP Server started successfully, open a browser and type in your server name in the URL box.

If you are going to run Application Response Measurement (ARM) agents, make sure you have the authority to run ARM agents when you start IBM HTTP Server.

What to do next

You can configure your server for:

- Secure Sockets Layer (SSL)
- Lightweight Directory Access (LDAP)
-  Fast Response Cache Accelerator (FRCA)

For more `apachectl` command options see Apache Hypertext Transfer Protocol Server.

Using Windows services to start IBM HTTP Server

This topic provides information on getting started with IBM HTTP Server on Windows operating systems.

About this task

Start IBM HTTP Server as a Windows service as follows:

Procedure

1. Click **Start > Programs > IBM HTTP Server > Start Server**. A message box is displayed that indicates the server has started.
2. To confirm that IBM HTTP Server started successfully by opening a browser window and type in your server name in the URL box. If you use the non-Administrator installation option, then the IBM HTTP Server does not install as a service. You have to run the `httpd.exe` file from a command line.

If IBM HTTP Server does not start:

- a. Go to **Services** in the Control Panel.
- b. Double-click **IBM HTTP Server** to start the server.
- c. To confirm that IBM HTTP Server started successfully, open a browser and type in your server name in the URL box.

If you are going to run Application Response Measurement (ARM) agents, make sure you have the authority to run ARM agents when you start IBM HTTP Server.

Results

IBM HTTP Server starts successfully.

What to do next

You can configure your server for Secure Sockets Layer (SSL), Lightweight Directory Access Protocol (LDAP), and Fast Response Cache Accelerator (FRCA).

Using JCL procedures to start IBM HTTP Server on z/OS

You can prepare JCL procedures to start and stop IBM HTTP Server from the MVS system console.

By using a JCL cataloged procedure to issue the apachectl start and stop commands, you can start and stop an IBM HTTP Server instance from the MVS system console. Other apachectl commands can be issued from the MVS system console using the same procedure.

Copy the following sample JCL procedure from SHAPJCL(HAPAPROC) to your system procedure library:

```
/*-----  
//IHSAPACH PROC ACTION='start',  
//      DIR='/path/to/IHS/runtime/directory',  
//      CONF='conf/httpd.conf'  
/*-----  
//IHS      EXEC PGM=BPXBATCH,  
// PARM='SH &DIR/bin/apachectl -k &ACTION -f &CONF -DNO_DETACH',  
// MEMLIMIT=512M  
//STDOUT DD PATH='&DIR/logs/proc.output',  
//      PATHOPTS=(OWRONLY,OCREAT,OTRUNC),  
//      PATHMODE=(SIRUSR,SIWUSR,SIRGRP,SIWGRP)  
//STDERR DD PATH='&DIR/logs/proc.errors',  
//      PATHOPTS=(OWRONLY,OCREAT,OTRUNC),  
//      PATHMODE=(SIRUSR,SIWUSR,SIRGRP,SIWGRP)  
//      PEND
```

A description of the apachectl command used in the sample JCL can be found at the Apache HTTP Server Control Interface Web site.

The default jobname for the IBM HTTP Server instance will be the same as the member name of the cataloged procedure. In the examples below, a procedure name of WEBSRV1 is used. Edit the new cataloged procedure by replacing */path/to/IHS/runtime/directory* with the actual installation directory for this instance of IBM HTTP Server. Create a SAF STARTED profile to associate the server user ID and group with the Web server started task:

```
RDEFINE STARTED WEBSRV1.* STDATA(USER(WWWSERV) GROUP(WWWGROUP) TRACE(YES))  
SETROPTS RACLIST(STARTED) GENERIC(STARTED) REFRESH
```

- To start the server from the MVS system console, enter:

```
S WEBSRV1
```

Note: The Web server name can be changed by adding jobname to the start command, for example:

```
S WEBSRV1,JOBNAME=HTTPD
```

- To stop the server, enter:

```
S WEBSRV1,ACTION='stop'
```

Note: When using SDSF.LOG, you must use the System Command Extension (command entry) screen to enter the command to stop the server.

- At the command prompt, type a forward slash (/) and then hit enter to access the System Command Extension window.
- From the System Command Extension window, enter the S WEBSRV1,ACTION='stop' command. Make sure that stop is in lower case.

- To issue other apachectl commands, enter:

```
S WEBSRV1,ACTION='<command>'
```

The output files for the start and stop commands are:

- *install_directory/logs/proc.output*
- *install_directory/logs/proc.errors*

Best Practice 1: The output files are overwritten each time the procedure is used. They might contain warning messages about the configuration or error messages for startup failures. If you want to retain a log of these messages across multiple uses of the procedure, modify the two occurrences of the PATHOPTS option in the sample procedure to

PATHOPTS=(OCREAT,OAPPEND,OWRONLY). For more information on the PATHOPTS option, refer to the *z/OS MVS JCL Reference (SA22-7597)*. Link to this document from the *z/OS Internet Library*.

Best Practice 2: The STDENV DD statement is not recommended. You might consider adding environment variable settings to the bin/envvars file within the runtime directory so that the variables are active whether IBM HTTP Server is started from JCL or from the UNIX environment.

Best Practice 3: The SH parameter of BPXBATCH is recommended instead of the PGM parameter. Processing for the PGM parameter bypasses system default settings in the /etc/profile file, including the umask setting, and files created by IBM HTTP Server do not have the correct permissions.

Configuring IBM HTTP Server

To configure the IBM HTTP Server, edit the httpd.conf configuration file.

Procedure

- **Locating the default and sample configuration files.**

The httpd.conf configuration file is in the conf directory of your server installation. There is also an httpd.conf.default file, in case you need to use another copy of the original file.

IBM HTTP Server also provides the following configuration files:

- **Distributed operating systems** admin.conf.default
- magic.default
- mime.types.default

- **Special considerations for IBM HTTP Server.** The following items regarding the configuration file should be known when using IBM HTTP Server:

- Configuration files that only support single-byte characters (SBCS) are:
 - httpd.conf (IBM HTTP Server configuration file)
 - **Distributed operating systems** admin.conf (Administration server configuration file)
- **Windows** The forward slash character (/) should be used as a path separator in the configuration file, instead of the backward slash character (\).

Apache modules (containing directives) supported by IBM HTTP Server

This section provides information on Apache modules that are supported by IBM HTTP Server. The directives defined within the supported Apache modules can be used to configure IBM HTTP Server.

Supported Apache modules

The following Apache modules were changed or are not supported.

Best Practice: **Distributed operating systems** If you are using the mod_ibm_ldap module for your LDAP configuration, consider migrating your mod_ibm_ldap directives to use the mod_ldap module. The mod_ibm_ldap module is provided with this release of IBM HTTP Server for compatibility with previous releases, however, you must migrate existing configurations to use the mod_authnz_ldap and mod_ldap modules to ensure future support for your LDAP configuration.

- The `mod_file_cache` module is provided with this release of IBM HTTP Server for compatibility with previous releases, however, you must migrate existing configurations to use the `mod_mem_cache` module to ensure future support for your LDAP configuration. These modules provide equivalent function in the memory instead of on a disk.
- The `mod_mime_magic` module is provided with this release of IBM HTTP Server for compatibility with previous releases, but might not be available in a future release. No replacement will be provided for this module.
- The `mod_proxy_ftp` module is provided with this release of IBM HTTP Server for compatibility with previous releases, but might not be available in a future release. No replacement will be provided for this module.
- The `mod_cern_meta` module is not supported. Instead use the `mod_headers` module.
- The `mod_imap` module was renamed to `mod_imagemap`. The `LoadModule` directive for the `mod_imap` module must be changed to refer to the new module name for an existing configuration file.
- You must set the `EnableExceptionHook` directive value to `On` for the `mod_backtrace` and `mod_whatkilledus` diagnostic modules.
- You may set the `McacheMinObjectSize` directive value to a minimum of 1 for the `mod_mem_cache` module. In previous releases, the minimum value was zero.
- The `Compression_Level` directive for the `mod_deflate` module was renamed to `DeflateCompressionLevel`.
- The configurations for the `mod_ldap` and the `mod_auth_ldap` modules have changed. See the procedure below about migrating from the `mod_ldap` and `mod_auth_ldap` module configurations.
- The Apache `mod_example` source is installed in the `<install>/example_module` directory.
- The `AddOutputFilterByType` directive now applies to proxy requests.
- Directory listings created by the `mod_autoindex` module now have a default character set which can be modified using the `IndexOptions` directive. If you rely on browser detection of character sets for correct display of directory listings, you might need to specify the correct character set using the `IndexOptions` directive.
- The `mod_proxy_balancer` and `mod_proxy_ajp` modules are not supported components of IBM HTTP Server. These modules are exclusively provided for, and supported by the following IBM products that include the IBM HTTP Server as part of their package:
 - WebSphere Community Edition uses both the `mod_proxy_balancer` and `mod_proxy_ajp` modules. Certain support agreements are required for these modules.
 - WebSphere Commerce uses the `mod_proxy_balancer` module in direct support of the Social Commerce functionality.

These modules are distributed with IBM HTTP Server in the following subdirectory:

Linux
 HP-UX Solaris AIX z/OS
 modules/WebSphereCE/

Windows
 C:\modules\WebSphereCE

The following table contains a list of Apache modules supported for IBM HTTP Server.

Table 1. Apache modules. The table lists the Apache module, a brief description of the module, and a web address to a detailed description of each module.

Module	Description	Web address
core	Core Apache HTTP Server features	http://publib.boulder.ibm.com/httserv/manual70/mod/core.html
Windows mpm_winnt	Multi-processing module (MPM)	http://publib.boulder.ibm.com/httserv/manual70/mod/mpm_winnt.html

Table 1. Apache modules (continued). The table lists the Apache module, a brief description of the module, and a web address to a detailed description of each module.

<p>AIX HP-UX Linux Solaris z/OS</p> <p>AIX HP-UX Linux Solaris z/OS worker</p>	MPM	http://publib.boulder.ibm.com/htpasswd/manual70/mod/worker.html
mod_actions	Provides for executing CGI scripts, based on media type or request method.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_actions.html
mod_alias	Provides for mapping different parts of the host file system in the document tree and for URL redirection.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_actions.html
mod_asis	Sends files that contain their own HTTP headers.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_asis.html
mod_auth_basic	Basic authentication	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_auth_basic.html
mod_authn_anon	Allows anonymous user access to authenticated areas.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authn_anon.html
mod_authn_dbm	User authentication using DBM files.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authn_dbm.html
mod_authn_default	Authentication fallback module	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authn_default.html
mod_authn_file	User authentication using text files	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authn_file.html
z/OS mod_authnz_ldap	Allows an LDAP directory to be used to store the database for HTTP basic authentication.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authnz_ldap.html
mod_authz_dbm	Group authorization using DBM files.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authz_dbm.html
mod_authz_default	Authorization fallback module	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authz_default.html
mod_authz_groupfile	Group authorization using text files	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authz_groupfile.html
mod_authz_host	Group authorizations based on host, such as host name or IP address	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authz_host.html
mod_authz_user	User authorization	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_authz_user.html
Windows z/OS mod_autoindex	Generates directory indexes automatically. This is similar to ls command on the UNIX platform or the Win32 dir shell command.	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_autoindex.html
AIX HP-UX Linux Solaris mod_cache	Content cache keyed to URIs	http://publib.boulder.ibm.com/htpasswd/manual70/mod/mod_cache.html

Table 1. Apache modules (continued). The table lists the Apache module, a brief description of the module, and a web address to a detailed description of each module.

<p>AIX HP-UX Linux Solaris</p> <p>Windows z/OS mod_cgi</p>	Execution of CGI scripts	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_cgi.html
<p>AIX HP-UX Linux Solaris</p> <p>AIX HP-UX Linux Solaris mod_cgid</p>	Execution of CGI scripts using an external CGI daemon.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_cgid.html
<p>z/OS</p> <p>z/OS mod_charset_lite</p>	Specifies character set translation or recoding.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_charset_lite.html
<p>Distributed operating systems</p> <p>Distributed operating systems mod_dav</p>	Distributed Authoring and Versioning (WebDAV) functionality. Tip: z/OS Although mod_dav and mod_dav_fs are not supported, IBM HTTP Server and the WebSphere plug-in can pass through WebDAV requests to WebSphere.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_dav.html
<p>Distributed operating systems</p> <p>Distributed operating systems mod_dav_fs</p>	File system provider for mod_dav.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_dav_fs.html
mod_deflate	Compress content before it is delivered to the client.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_deflate.html
mod_dir	Provides for "trailing slash" redirects and serving directory index files.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_dir.html
mod_disk_cache	Implements a disk based storage manager. It is primarily of use in conjunction mod_cache.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_disk_cache.html
mod_env	Modifies the environment which is passed to CGI scripts and SSI pages.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_env.html
mod_expires	Generation of Expires and Cache Control HTTP headers according to user-specified criteria.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_expires.html
<p>Distributed operating systems</p> <p>mod_ext_filter</p>	Pass the response body through an external program before delivery to the client.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_ext_filter.html
<p>Distributed operating systems</p> <p>Distributed operating systems mod_file_cache</p>	Caches a static list of files in memory. This module is provided with this release for compatibility with previous releases. Begin using mod_mem_cache or mod_cache to ensure compatibility with future releases of IBM HTTP Server. Tip: The recommended caching mechanism for file handling is the CacheEnable feature of the mod_cache module.	http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_file_cache.html

Table 1. Apache modules (continued). The table lists the Apache module, a brief description of the module, and a web address to a detailed description of each module.

Distributed operating systems mod_filter	Specifies the context-sensitive smart filter configuration module.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_filter.html
mod_headers	Customization of HTTP request and response headers.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_headers.html
mod_imagemap	Server-side image map processing.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_imagemap.html
mod_include	Server-parsed HTML documents (Server Side Includes).	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_include.html
mod_info	Provides a comprehensive overview of the server configuration.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_info.html
z/OS mod_ldap	Provides LDAP connection pooling and result caching services for use by other LDAP modules.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_ldap.html
mod_log_config	Logging of the requests made to the server.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_log_config.html
mod_logio	Logging of input and output bytes per request.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_logio.html
mod_mem_cache	Content cache keyed to URIs.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_mem_cache.html
mod_mime	Associates the requested file extensions with the behavior of the file (handlers and filters), and content (mime-type, language, character set and encoding).	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_mime.html
Distributed operating systems Distributed operating systems mod_mime_magic	Determines the MIME type of a file by looking at a few bytes of its contents. This module is provided with this release of IBM HTTP Server for compatibility with previous releases, but will not be supported in a future release. No replacement will be provided for this module. Important: Using mod_mime_magic can decrease performance because the file must be read and compared to a set of patterns to determine the content-type.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_mime_magic.html
mod_negotiation	Provides for content negotiation.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_negotiation.html
mod_proxy	HTTP, 1.1 proxy, and gateway server	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_proxy.html
mod_proxy_connect	Specifies the mod_proxy module extension for CONNECT request handling.	http://publib.boulder.ibm.com/httserv/manual70/mod/mod_proxy_connect.html

Table 1. Apache modules (continued). The table lists the Apache module, a brief description of the module, and a web address to a detailed description of each module.

Distributed operating systems Distributed operating systems mod_proxy_ftp	Provides FTP support for the mod_proxy module. This module is provided with this release of IBM HTTP Server for compatibility with previous releases, but will not be supported in a future release. No replacement will be provided for this module.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_proxy_ftp.html
mod_proxy_http	Provides HTTP support for the mod_proxy module.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_proxy_http.html
mod_rewrite	Provides a rule-based rewriting engine to rewrite requested URLs.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_rewrite.html
mod_setenvif	Enables the setting of environment variables based on characteristics of the request.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_setenvif.html
mod_so	Loading of executable code and modules into the server at start or restart time.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_so.html
mod_speling	Attempts to correct mistaken URLs that users might have entered by ignoring capitalization and by allowing up to one misspelling.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_speling.html
mod_status	Provides information on server activity and performance.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_status.html
mod_suexec	Allows CGI scripts to run as the specified user or group.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_suexec.html
mod_unique_id	Provides an environment variable with a unique identifier for each request.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_unique_id.html
mod_userdir	User-specific directories.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_userdir.html
mod_usertrack	Clickstream logging of user activity on a site.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_usertrack.html
mod_vhost_alias	Provides for dynamically configured mass virtual hosting.	http://publib.boulder.ibm.com/htpserv/manual70//mod/mod_vhost_alias.html

Apache programs supported by IBM HTTP Server

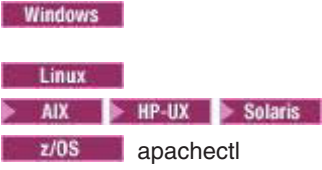
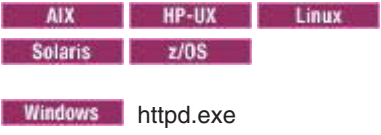

This section provides information on Apache programs that are supported by IBM HTTP Server. These supported Apache programs can be used to configure IBM HTTP Server.

Supported Apache programs

The following table contains a list of Apache commands supported for IBM HTTP Server.

Note: **Windows** The apache.exe command was replaced with the httpd.exe command. The apache.exe command is provided with this release of IBM HTTP Server for compatibility with previous releases. Migrate existing scripts and procedures to use the httpd.exe command to ensure future support for this functionality.

Program	Description	URL
---------	-------------	-----

ab	Provides benchmarking functionality for the Web server	http://publib.boulder.ibm.com/htpserv/manual70/programs/ab.html
 apachectl	Provides start, stop, and restart functionality for the Web server.	http://publib.boulder.ibm.com/htpserv/manual70/programs/apachectl.html
 httpd.exe	Provides start, stop, and restart functionality for the Web server.	http://publib.boulder.ibm.com/htpserv/manual70/programs/httpd.html
 apxs	Builds plug-in modules.	http://publib.boulder.ibm.com/htpserv/manual70/programs/apxs.html
dbmmanage	Creates and updates user authentication files in DBM format for basic authentication.	http://publib.boulder.ibm.com/htpserv/manual70/programs/dbmmanage.html
htdbm	Creates and updates user authentication files in DBM format for basic authentication.	http://publib.boulder.ibm.com/htpserv/manual70/programs/htdbm.html
htpasswd	Creates and updates user authentication files for basic authentication.	http://publib.boulder.ibm.com/htpserv/manual70/programs/htpasswd.html
httxt2dbm	Creates DMB files for use with RewriteMap.	http://publib.boulder.ibm.com/htpserv/manual70/programs/httxt2dbm.html
logresolve	Resolves host names for IP addresses in Apache log files.	http://publib.boulder.ibm.com/htpserv/manual70/programs/logresolve.html
rotatelog	Rotates log files without having to stop the server.	http://publib.boulder.ibm.com/htpserv/manual70/programs/rotatelog.html

Apache APR and APR-util libraries supported by IBM HTTP Server

This section provides information about the Apache Portable Runtime (APR) and APR-util libraries that are supported by IBM HTTP Server. IBM HTTP Server supports only the APR and APR-util libraries installed with the product. Copies of the libraries cannot be substituted.

Supported APR and APR-util libraries

The APR and APR-util libraries installed with IBM HTTP Server are provided for only IBM and third-party plug-in modules loaded into IBM HTTP Server. Use of these libraries by stand-alone applications or commands, other than those provided with IBM HTTP Server, is not supported.

The following build-time features of APR and APR-util are not provided on all platforms.

- random number support
- native atomic operation support
- i18n translation
- DBD support is not provided for any platform
- LDAP support is not provided for any platform.

The only supported APR-util library database management type is SDBM. SDBM affects the htdbm and htxt2dbm commands. It also affects the mod_authn_dbm, mod_authz_dbm, and mod_rewrite modules for






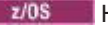
DBM map files and the mod_dav module for the lock database.

Apache MPM and addressing modes supported by IBM HTTP Server

This section provides information about Apache Multi-processing module (MPM) and addressing modes supported by IBM HTTP Server.

The following table contains a list of platforms and the MPM and addressing modes supported on those platforms by IBM HTTP Server.

Table 2. MPM and addressing modes. The table lists the platform, addressing mode, and MPM.

Platform	Addressing mode	MPM
AIX®	32-bit and 64-bit	worker MPM
 HP-UX/PA-RISC	32-bit	worker MPM
     HP-UX/ia64	64-bit	worker MPM
Linux/x86	32-bit and 64-bit	worker MPM
Linux/PPC	32-bit and 64-bit	worker MPM
Linux on System z®	32-bit and 64-bit	worker MPM
Solaris/SPARC	32-bit and 64-bit	worker MPM
Solaris/x64	64-bit	worker MPM
Windows	32-bit	WintNT MPM
z/OS	64-bit	worker MPM

IPv4 and IPv6 configuration for Windows operating systems

IBM HTTP Server supports IPv6 on Windows XP and 2003 operating systems. It does not support IPv6 on the Windows 2000 operating system.

Support for IPv6 on Windows operating systems is configured differently than other supported platforms. The Listen directive on Windows operating systems should always include either an IPv4 address or an IPv6 address. Any existing Listen directives that are not qualified with an IP address should be updated to include one, even if Windows IPv6 networking is not configured.

Use 0.0.0.0 for the default IPv4 address and [:::] for the default IPv6 address. Add the following line in httpd.conf configuration file to listen on IPv6 port 80:

```
Listen [::]:80
```

If you want to accept connections over IPv4, configure Listen 0.0.0.0:80 or AfpPort 80. Advanced fast path architecture (AFPA) is only supported for IPv4.

Configure Windows IPv6 networking before enabling the Listen directive for IPv6.

Serving static content faster with Fast Response Cache Accelerator

The fast response cache accelerator (FRCA) can improve the performance of the IBM HTTP Server when serving static content, such as text and image files. Support for FRCA is on AIX and Windows systems only.

Before you begin

The use of FRCA is discouraged, as it has been deprecated. FRCA might be removed in a future release. This release has no AFPA directives in the default configuration file.

Windows The `afpa.sys` file is no longer copied during the installation to the `C:\Windows\system32\drivers\` directory.

Windows There is no support for any Windows 64-bit operating system. FRCA is not supported on Windows Vista, Windows 2008, or any later Windows operating systems. This is a permanent restriction.

AIX FRCA on AIX is only supported with the 32-bit installation of IBM HTTP Server.

Read about deprecated features in the topic about deprecated features.

Read about Fast Response Cache Accelerator operational restrictions for information on restrictions and limitations.

About this task

When FRCA is enabled, the default configuration setting allows all static files to be cached. The cache automatically loads during server operation so that individual files do not need to be listed. Use the `AfpaCache` directive to turn caching on or off for specific directories.

FRCA will remove files from the cache when they change to avoid serving stale content.

Procedure

Configure FRCA.

- Windows** Copy the `afpa.sys` file from the IBM HTTP Server installation directory to the existing `C:\Windows\system32\drivers` directory if it is not already present. The `afpa.sys` file might already be in the directory if an earlier version of IBM HTTP Server has been previously installed on the system.
- Edit the `httpd.conf` configuration file by commenting out the `Listen` directive, and adding the following statements:

```
##  
LoadModule ibm_afpa_module modules/mod_afpa_cache.so  
  
AfpaEnable  
AfpaCache on  
AfpaPort 80  
AfpaLogFile "/logs/afpalog" V-ECLF
```

Customizing Fast Response Cache Accelerator logging

The fast response cache accelerator (FRCA) can improve the performance of the IBM HTTP Server when serving static content, such as text and image files. By default, FRCA generates an access log of all requests that are served out of the cache. In order to minimize the effect of logging on performance, this is a separate file from the normal Apache access log.

Before you begin

The use of FRCA is discouraged, as it was deprecated starting in WebSphere Application Server Version 7.0. FRCA might be removed in a future release. This release has no AFPA directives in the default configuration file.

Read about deprecated features in the Deprecated features topic.

Read the Fast Response Cache Accelerator operational restrictions topic for information on restrictions and limitations.

About this task

When FRCA is enabled, the default configuration setting allows all static files to be cached. The cache automatically loads during server operation so that individual files do not need to be listed. Use the `AfpaCache` directive to turn caching on or off for specific directories.

FRCA will remove files from the cache when they change to avoid serving stale content.

Enable the FRCA access log if you want to maintain a record of requests served by FRCA. Requests that are not served out of the cache will be logged in the FRCA access log file. The FRCA access log file provides a useful way to verify that caching is enabled and to identify cached files.

Note: Even though a particular file might be cached, it might not always be served from the cache. Therefore, not every request for a cached file will result in an FRCA access log entry.

If you do not need access logging, turn the logging off for better performance.

Procedure

- To turn FRCA logging off, edit the `httpd.conf` configuration file.

AIX Configure `AfpaLogging` off.

Windows Insert a comment character (`#`) at the beginning of the `AfpaLogFile` line. For example:

```
#AfpaLogFile "_path_to_server_/logs/afpalog" V-ECLF
```

- For each request that is served by the fast response cache accelerator, a log entry in the access log displays the following:
 - Source host address
 - **Windows** Date and time of the request
 - HTTP method of the request and what is requested
 - HTTP return code, which indicates whether the request is honored
 - Size of the returned data

A log entry can also optionally display the following:

- Target virtual host (use the formatting option `V-CLF` or `V-ECLF`)
- HTTP referer (use the formatting option `V-CLF` or `V-ECLF`)
- HTTP user agent (use the formatting option `V-CLF` or `V-ECLF`)

Note: **Windows** The log file has a date stamp that automatically appends to its name. Everyday at midnight the server closes the current access log and creates a new one. This action enables the log file to process without having to stop and restart the server. Under heavy load conditions the log file can grow rapidly. Provide sufficient space on the hard drive for storage.

Restrictions on cached content

This section discusses the caching restrictions for the fast response cache accelerator (FRCA).

Caching does not occur on the following page types:

- Default welcome pages
- Requests ending in `/"`
- Access-protected documents and pages requested over Secure Sockets Layer (SSL)

Caching limitations exist for the following situations:

- FRCA supports only limited multi-language content negotiation. Caching occurs for only a single language version, where a given URL maps to multiple translated versions.
- FRCA must not be used with locally-mounted network file systems, such as Network File System (NFS) or Windows shared drives.
- FRCA does not cache proxied content.

Fast Response Cache Accelerator operational restrictions

This section discusses the operational restrictions for the fast response cache accelerator (FRCA). FRCA is also known as Adaptive Fast Path Architecture (AFPA).

The following operational restrictions apply:

- When FRCA is enabled, the default value of 0 for the `MaxRequestsPerChild` directive should be used, because graceful server restart is not supported with the cache accelerator.
- FRCA does not support Windows 64-bit operating systems.
- FRCA cannot be used when certain antivirus software is enabled. Currently Norton Antivirus has been identified as one such program.
- FRCA access log entries are not integrated with the Apache access log.
- Only access logging facilities exist for monitoring FRCA.
- On a given machine, only one instance of the IBM HTTP Server can have FRCA enabled.
- Do not install the IBM HTTP Server on a machine running the IBM Netfinity® Web Server Accelerator.
- FRCA does not support IPv6.
- FRCA must not be used with locally-mounted network file systems, such as Network File System (NFS) or Windows shared drives.
- On Windows systems, AFPA cannot be loaded if using the `Win32DisableAcceptEx` directive
- FRCA is not supported on Windows Vista, Windows 2008, or any later Windows operating systems.
- FRCA on AIX is only supported with the 32-bit installation of IBM HTTP Server.

Due to these operational restrictions, and since FRCA/AFPA was deprecated starting in V7.0, its use is discouraged. Instead, it is recommended to use the IBM HTTP Server default configuration to serve static files. This configuration provides support for features not available with FRCA/AFPA usage, such as IPv6, SSL, Authentication and access-control, custom headers and compression.

If CPU usage with the default configuration is too high, the `mod_mem_cache` module can be configured to cache frequently accessed files in memory, or multiple web servers can be used to scale out horizontally. Additional options include the offloading of static files to a Content Delivery Network (CDN) or caching HTTP appliance, or to use the caching proxy component of WebSphere Edge Server in WebSphere Application Server Network Deployment (ND).

Servlets and JavaServer Pages files caching

You can use the fast response cache accelerator (FRCA) with WebSphere Application Server to cache certain dynamically-generated servlet and JavaServer Pages (JSP) files. This feature is only available on Windows versions of IBM HTTP Server.

Enable IBM HTTP Server for dynamic page caching by enabling the cache accelerator. In addition, the `afpaplugin_22.dll` component that is compatible with IBM HTTP Server V7.0 must be configured by WebSphere Application Server.

For details on how to enable this capability, see *Configuring high-speed external caching through the Web server* in the WebSphere Application Server product documentation.

AIX considerations for Fast Response Cache Accelerator (FRCA)

There are special considerations when using FRCA on AIX platforms. The FRCA kernel extension must load before starting IBM HTTP Server with FRCA enabled. Also, increasing the upper-bound limit of the percentage of CPU time that the FRCA kernel extension can spend in the interrupt (high priority) context is not recommended.

The following items must be considered when you use fast response cache accelerator (FRCA) on AIX platforms:

- The FRCA kernel extension must load before starting IBM HTTP server with FRCA enabled. To do this, issue the **frcactrl load** command. This is normally configured to run whenever the system boots and before IBM HTTP Server starts. See the AIX man pages for more details about the **frcactrl** command.
- In order to place an upper bound on the percentage of CPU time that the FRCA kernel extension can spend in its interrupt (high priority) context, use the **frcactrl pctionintr** command. Increasing this above the default value of 80% is not recommended in order to allow other applications a reasonable amount of time to execute. Decrease this value if more time needs to be allocated to other applications, but note that reducing the value will result in more cache misses, even if a file is in the cache.

AFPA directives

These configuration parameters control the Advanced Fast Path Architecture (AFPA) feature in IBM HTTP Server.

The fast response cache accelerator (FRCA) utilizes a special high-performance component, based on the IBM Advanced Fast Path Architecture, from which the AFPA prefix is derived. You can configure FRCA for IPV4. IPV6 is not supported.

- “Afpabindlogger directive”
- “Afpacache directive” on page 21
- “Afpadynacachemax directive” on page 21
- “Afpaenable directive” on page 21
- “Afpalogfile directive” on page 21
- “Afpalogging directive” on page 22
- “Afpamaxcache directive” on page 22
- “Afpamincache directive” on page 22
- “Afpaport directive” on page 22
- “Afparevalidationtimeout directive” on page 22
- “Afpasendserverheader” on page 23

AIX

Afpabindlogger directive

Use the Afpabindlogger directive to bind the fast response cache accelerator (FRCA) logging thread in the kernel to a specific processor.

The format of the command is Afpabindlogger [-1, 0, 1, ..., n], where -1 leaves the logging thread unbound and a number from 0 to total number of processors on the system, binds the logging thread to that processor.

Syntax
Scope
Default
Notes

Afpabindlogger [-1,0,1,...,n]
One per physical Apache server
(-1)
Valid on AIX operating systems only.

AIX

Windows

AfpaCache directive

The AfpaCache directive turns the fast response cache accelerator (FRCA) on or off for a particular scope (such as a directory). The AfpaCache directive applies to all descendants in a scope, unless otherwise modified by another directive.

Scope	Server configuration, virtual host, directory
Syntax	On or off
Usage	AfpaCache on
Override	Options
Multiple instances in the configuration file	Allowed
Notes	Valid on Windows 32-bit and AIX operating systems.

Windows

AfpaDynacacheMax directive

The AfpaDynacacheMax directive is used on Windows operating systems to control the total amount of memory utilized for caching servlets and JavaServer Pages files.

When static files are cached, there is very little overhead for each entry since the file itself does not take up space in the cache, just the file handle. However, for servlets and JavaServer Pages files, the body of the response is stored in physical memory, so care must be taken to avoid consuming all available memory. Without this directive, the fast response cache accelerator will automatically set the upper bound to be approximately one eighth of physical memory. Use the directive to override that default.

Syntax	AfpaDynacacheMax size (Megabytes)
Scope	One per physical Apache server
Notes	Valid on Windows 32-bit operating systems

AIX

Windows

AfpaEnable directive

The AfpaEnable directive enables the fast response cache accelerator (FRCA). If the AfpaEnable directive is present and mod_afpa_cache.so is loaded, FRCA listens on the port specified by the AfpaPort directive.

Syntax	AfpaEnable
Scope	One per physical Apache server
Notes	Valid on AIX and Windows operating systems.

AIX

Windows

AfpaLogFile directive

The AfpaLogFile directive defines the fast response cache accelerator (FRCA) log file name, location, and logging format.

Scope	One per physical Apache server
Syntax	AfpaLogFile <i>log_file_name</i> [CLF ECLF V-CLF V-ECLF BINARY]
Notes	Valid on AIX and Windows 32-bit operating systems. On Windows 32-bit operating systems, the current date is used as the file type for the log file, and the log file is automatically rolled over at midnight each day.

The log formats are as follows:

- CLF = Common Log Format
- ECLF = Extended Common Log Format
- V-CLF = Common Log Format with virtual host information
- V-ECLF = Extended Common Log Format with virtual host information
- **AIX** BINARY = Binary log with virtual host information

AIX

AfpaLogging directive

The AfpaLogging directive turns the fast response cache accelerator (FRCA) logging on or off.

Scope	One per physical Apache server
Syntax	AfpaLogging On Off
Notes	Valid only on AIX operating systems.

AIX

AfpaMaxCache directive

The AfpaMaxCache directive specifies the maximum file size inserted into the fast response cache accelerator (FRCA) cache.

Syntax	AfpaMaxCache [size (bytes)]
Scope	One per physical Apache server
Default	none
Notes	Valid only on AIX operating systems.

AIX

AfpaMinCache directive

The AfpaMinCache directive specifies the minimum file size inserted into the fast response cache accelerator (FRCA) cache.

Syntax	AfpaMinCache [size]
Scope	One per physical Apache server
Default	none
Notes	Valid only on AIX operating systems.

AIX

Windows

AfpaPort directive

The AfpaPort directive tells the FRCA on which TCP port to listen. The AfpaPort directive issues a listen command for all TCP network adapters that are active on the server machine. The listen command is effective for all TCP addresses.

Syntax	AfpaPort <i>port number</i>
Scope	One directive per server
Notes	Valid only on AIX and Windows 32-bit operating systems

AIX

AfpaRevalidationTimeout directive

The AfpaRevalidationTimeout directive sets the time interval for revalidation of a cached object. When the RevalidationTimeout is exceeded for a cached object, a fresh copy is cached.

Syntax	AfpaRevalidationTimeout [value]
Scope	Global
Default	60 seconds
Notes	Valid on AIX operating systems only.

AIX

AfpaSendServerHeader

The AfpaSendServerHeader directive specifies whether or not the fast response cache accelerator (FRCA) sends the HTTP Server header in the response.

Syntax	AfpaSendServerHeader true or false
Scope	One per physical Apache server
Default	true
Notes	Valid only on AIX operating systems.

Enabling IBM HTTP Server for FastCGI applications

FastCGI applications use TCP or UNIX sockets to communicate with the Web server. This scalable architecture enables applications to run on the same platform as the Web server, or on many machines scattered across an enterprise network.

About this task

You can port FastCGI applications to other Web server platforms. Most popular Web servers support FastCGI directly, or through commercial extensions.

FastCGI applications run fast because of their persistency. These applications require no per-request startup and initialization overhead. This persistency enables the development of applications, otherwise impractical within the CGI paradigm, like a huge Perl script, or an application requiring a connection to one or more databases.

Procedure

1. Load the mod_fastcgi module into the server.

```
LoadModule fastcgi_module modules/mod_fastcgi.so
```
2. Configure FastCGI using the FastCGI directives.

Example

Windows

In the following configuration example, the c:/Program Files/IBM/HTTPServer/fcgi-bin/ directory contains FastCGI echo.exe applications. Requests from Web browsers for the /fcgi-bin/echo.exe URI will be handled by the FastCGI echo.exe application :

```
LoadModule fastcgi_module modules/mod_fastcgi.so
<IfModule mod_fastcgi.c>
    AllowOverride None
    Options +ExecCGI
    SetHandler fastcgi-script
</Directory>
```

```
FastCGIServer "C:/Program Files/IBM/HTTPServer/fcgi-bin/echo.exe" -processes 1
```

```
</IfModule>
```

In the following configuration example, the /opt/IBM/HTTPServer/fcgi-bin/ directory contains FastCGI applications, including the echo application. Requests from Web browsers for the /fcgi-bin/echo URI will be handled by the FastCGI echo application :

```
LoadModule fastcgi_module modules/mod_fastcgi.so
<IfModule mod_fastcgi.c>
ScriptAlias /fcgi-bin/ "/opt/IBM/HTTPServer/fcgi-bin/"

<Directory "/opt/IBM/HTTPServer/fcgi-bin/"
    AllowOverride None
    Options +ExecCGI
    SetHandler fastcgi-script
</Directory>

FastCGIServer "/opt/IBM/HTTPServer/fcgi-bin/echo" -processes 1
</IfModule>
```

Learn about FastCGI

FastCGI is an interface between Web servers and applications which combines some of the performance characteristics of native Web server modules with the Web server independence of the Common Gateway Interface (CGI) programming interface.

FastCGI is an open extension to CGI that is language independent and is a scalable architecture. FastCGI provides high performance and persistence without the limitations of server-specific APIs. The FastCGI interface is described at <http://www.fastcgi.com/>.

IBM HTTP Server provides FastCGI support with the mod_fastcgi module. The mod_fastcgi module implements the capability for IBM HTTP Server to manage FastCGI applications and to allow them to process requests.

A FastCGI application typically uses a programming library such as the FastCGI development kit from <http://www.fastcgi.com/>. IBM HTTP Server does not provide a FastCGI programming library for use by FastCGI applications.

FastCGI applications are not limited to a particular development language. FastCGI application libraries currently exist for Perl, C/C++, Java, Python and the transmission control layer (TCL).

For more information on FastCGI, visit the FastCGI Web site. To receive FastCGI related announcements and notifications of module updates, send mail to fastcgi-announce-request@idle.com with subscribe in the Subject field. To participate in the discussion of mod_fastcgi and FastCGI application development, send mail to fastcgi-developers-request@idle.com with subscribe in the Subject field.

The IBM HTTP Server Fast CGI plug-in provides an alternative method of producing dynamic content.

FastCGI directives

These configuration parameters control the FastCGI feature in IBM HTTP Server.

- “FastCGIAccessChecker directive” on page 25
- “FastCGIAccessCheckerAuthoritative directive” on page 25
- “FastCGIAuthenticator directive” on page 26
- “FastCGIAuthenticatorAuthoritative directive” on page 26
- “FastCGIAuthorizer directive” on page 27
- “FastCGIAuthorizerAuthoritative directive” on page 27
- “FastCGIConfig directive” on page 28
- “FastCGIExternalServer directive” on page 29
- “FastCGIIPCDir directive” on page 30
- “FastCGIServer directive” on page 31
- “FastCGISuEXEC directive” on page 32

FastCGIAccessChecker directive

The FastCGIAccessChecker directive defines a FastCGI application as a per-directory access validator.

Syntax	FastCGIAccessChecker <i>file name</i> [-compat]
Scope	directory, location
Default	Directory
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	File name

The Apache Access phase precedes user authentication and the HTTP headers submitted with the request determine the decision to enable access to the requested resource. Use FastCGI-based authorizers when a dynamic component exists as part of the access validation decision, like the time, or the status of a domain account.

If the FastCGI application file name does not have a corresponding static or external server definition, the application starts as a dynamic FastCGI application. If the file name does not begin with a slash (/), then the application assumes that the file name is relative to the ServerRoot.

Use the FastCgiAccessChecker directive within Directory or Location containers. For example:

```
<Directory htdocs/protected>
FastCgiAccessChecker fcgi-bin/access-checker
</Directory>
```

Mod_fastcgi sends nearly all of the standard environment variables typically available to CGI and FastCGI request handlers. All headers returned by a FastCGI access-checker application in a successful response (Status: 200), pass to subprocesses, or CGI and FastCGI invocations, as environment variables. All headers returned in an unsuccessful response pass to the client. Obtain FastCGI specification compliant behavior by using the -compat option.

Mod_fastcgi sets the environment variable FCGI_APACHE_ROLE to ACCESS_CHECKER, to indicate the Apache-specific authorizer phase performed.

The HTTP Server does not support custom failure responses from FastCGI authorizer applications. See the ErrorDocument directive for a workaround. A FastCGI application can serve the document.

FastCGIAccessCheckerAuthoritative directive

The FastCGIAccessCheckerAuthoritative directive enables access checking passing to lower level modules.

Syntax	FastCGIAccessCheckerAuthoritative On Off
Scope	directory, location
Default	FastCGIAccessCheckerAuthoritative On
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	On or off

Setting the FastCgiAccessCheckerAuthoritative directive explicitly to Off, enables access checking passing to lower level modules, as defined in the Configuration and modules.c files, if the FastCGI application fails to enable access.

By default, control does not pass on and a failed access check results in a forbidden reply. Consider the implications carefully before disabling the default.

FastCGIAuthenticator directive

The FastCGIAuthenticator directive defines a FastCGI application as a per-directory authenticator.

Syntax	FastCGIAuthenticator file name [-compat]
Scope	directory
Default	None
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	File name

Authenticators verify the requester by matching the user name and password that is provided against a list or database of known users and passwords. Use FastCGI-based authenticators when the user database is maintained within an existing independent program, or resides on a machine other than the Web server.

If the FastCGI application file name does not have a corresponding static or external server definition, the application starts as a dynamic FastCGI application. If the file name does not begin with a slash (/), then the file name is assumed to be relative to the ServerRoot.

Use the FastCgiAuthenticator directive within directory or location containers, along with an AuthType and AuthName directive. This directive only supports the basic user authentication type. This authentication type needs a require, or FastCgiAuthorizer directive, to work correctly.

```
/Directory htdocs/protected>
AuthType Basic
AuthName ProtectedRealm
FastCgiAuthenticator fcgi-bin/authenticator
require valid-user
</Directory>
```

The Mod_fastcgi directive sends nearly all of the standard environment variables that are typically available to CGI and FastCGI request handlers. All headers returned by a FastCGI authentication application in a successful response (Status: 200) pass to subprocesses, or CGI and FastCGI invocations, as environment variables. All headers returned in an unsuccessful response are passed to the client. Obtain FastCGI specification compliant behavior by using the -compat option.

The Mod_fastcgi directive sets the FCGI_APACHE_ROLE environment variable to AUTHENTICATOR, indicating the Apache-specific authorizer phase performed.

This directive does not support custom failure responses from FastCGI authorizer applications. See the ErrorDocument directive for a workaround. A FastCGI application can serve the document.

FastCGIAuthenticatorAuthoritative directive

The FastCGIAuthenticatorAuthoritative directive enables authentication passing to lower level modules defined in the configuration and modules.c files, if explicitly set to off and the FastCGI application fails to authenticate the user.

Syntax	FastCGIAuthenticatorAuthoritative <i>On</i> <i>Off</i>
Scope	directory
Default	FastCgiAuthenticatorAuthoritative On
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	On or off

Use this directive in conjunction with a well protected AuthUserFile directive, containing a few administration-related users.

By default, control does not pass on and an unknown user results in an Authorization Required reply. Consider implications carefully before disabling the default.

FastCGIAuthorizer directive

The FastCGIAuthorizer directives defines a FastCGI application as a per-directory authorizer.

Syntax	FastCgiAuthorizer file name [<i>-compat</i>]
Scope	directory
Default	None
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	File name

Authorizers validate whether an authenticated user can access a requested resource. Use FastCGI-based authorizers when a dynamic component exists as part of the authorization decision, such as the time, or currency of the user's bills.

If the FastCGI application file name does not have a corresponding static or external server definition, the application starts as a dynamic FastCGI application. If the file name does not begin with a slash (/) then the file name is assumed relative to the ServerRoot.

Use FastCgiAuthorizer within Directory or Location containers. Include an AuthType and AuthName directive. This directive requires an authentication directive, such as FastCgiAuthenticator, AuthUserFile, AuthDBUserFile, or AuthDBMUserFile to work correctly.

```
<Directory htdocs/protected>
AuthType Basic
AuthName ProtectedRealm
AuthDBMUserFile conf/authentication-database
FastCgiAuthorizer fcgi-bin/authorizer
</Directory>
```

The Mod_fastcgi directive sends nearly all of the standard environment variables typically available to CGI and FastCGI request handlers. All headers returned by a FastCGI authentication application in a successful response (Status: 200) pass to subprocesses, or CGI and FastCGI invocations, as environment variables. All headers returned in an unsuccessful response pass on to the client. Obtain FastCGI specification compliant behavior by using the -compat option.

The Mod_fastcgi directive sets the environment variable FCGI_APACHE_ROLE to AUTHORIZER, to indicate the Apache-specific authorizer phase performed.

This directive does not support custom failure responses from FastCGI authorizer applications. See the ErrorDocument directive for a workaround. A FastCGI application can serve the document.

FastCGIAuthorizerAuthoritative directive

The FastCGIAuthorizerAuthoritative directive enables authentication passing to lower level modules, as defined in the configuration and modules.c files, when explicitly set to Off, if the FastCGI application fails to authenticate the user.

Syntax	FastCgiAuthorizerAuthoritative file name <i>On</i> <i>Off</i>
Scope	directory
Default	FastCgiAuthorizerAuthoritative file name <i>On</i>

Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	On or off

Use this directive in conjunction with a well protected AuthUserFile containing a few administration-related users.

By default, control does not pass on and an unknown user results in an Authorization Required reply. Consider the implications carefully before disabling the default.

FastCGIConfig directive

The FastCGIConfig directive defines the default parameters for all dynamic FastCGI applications.

Syntax	FastCgiConfig <i>option option...</i>
Scope	The FastCgiConfig directive does not affect static or external applications.
Default	directory
Module	None
Multiple instances in the configuration file	mod_fastcgi
Values	yes
	Dynamic applications start upon demand. Additional application instances start to accommodate heavy demand. As demand fades, the number of application instances decline. Many of the options govern this process.

Option can include one of the following (case insensitive):

- **appConnTimeout n (0 seconds).** The number of seconds to wait for a connection to the FastCGI application to complete or 0, to indicate use of a blocking connect(). If the timeout expires, a SERVER_ERROR results. For non-zero values, this amount of time used in a select() to write to the file descriptor returned by a non-blocking connect(). Non-blocking connect()s are troublesome on many platforms. See also -idle-timeout; this option produces similar results, but in a more portable manner.
- **idle-timeout n (30 seconds).** The number of seconds of FastCGI application inactivity allowed before the request aborts and the event logs at the error LogLevel. The inactivity timer applies only when a pending connection with the FastCGI application exists. If an application does not respond to a queued request within this period, the request aborts. If communication completes with the application, but not with the client (a buffered response), the timeout does not apply.
- **autoUpdate none.** This option causes the mod_fastcgi module to check the age of the application on disk before processing each request. For recent applications, this function notifies the process manager and stops all running instances of the application. Build this type of functionality into the application. A problem can occur when using this option with -restart.
- **gainValue n (0.5).** A floating point value between 0 and 1 that is used as an exponent in the computation of the exponentially decayed connection times load factor of the currently running dynamic FastCGI applications. Old values are scaled by (1 - gainValue), so making values smaller, weights them more heavily compared to the current value, which is scaled by gainValue.
- **initial-env name[=value] none.** A name-value pair passed in the initial environment when instances of the application spawn. To pass a variable from the Apache environment, do not provide the "=" (if the variable is not actually in the environment, it is defined without a value). To define a variable without a value, provide the "=" without any value. This option is repeatable.
- **init-start-delay n (1 second).** The minimum number of seconds between the spawning of instances of this application. This delay decreases the demand placed on the system at server initialization.

- **killInterval n (300 seconds).** The killInterval determines how often the dynamic application instance killing policy is implemented within the process manager. Lower numbers result in a more aggressive policy, while higher numbers result in a less aggressive policy.
- **listen-queue-depth n (100).** The depth of the listen() queue, also known as the backlog, shared by all instances of this application. A deeper listen queue allows the server to cope with transient load fluctuations without rejecting requests; it does not increase throughput. Adding additional application instances can increase throughput and performance, depending upon the application and the host.
- **maxClassProcesses n (10).** The maximum number of dynamic FastCGI application instances allowed to run for any one FastCGI application.
- **maxProcesses n (50).** The maximum number of dynamic FastCGI application instances allowed to run at any time.
- **minProcesses n (5).** The minimum number of dynamic FastCGI application instances the process manager allows to run at any time, without killing them due to lack of demand.
- **multiThreshhold n (50).** An integer between 0 and 100 used to determine whether to terminate any instance of a FastCGI application. If the application has more than one instance currently running, this attribute helps to decide whether to terminate one of them. If only one instance remains, singleThreshhold is used instead.
- **pass-header header none.** The name of an HTTP Request Header passed in the request environment. This option makes the contents of headers available to a CGI environment.
- **priority n (0).** The process priority assigned to the application instances using setpriority().
- **processSlack n (5 seconds).** If the sum of all currently running dynamic FastCGI applications exceeds maxProcesses - processSlack, the process manager invokes the killing policy. This action improves performance at higher loads, by killing some of the most inactive application instances before reaching the maxProcesses value.
- **restart none.** This option causes the process manager to restart dynamic applications upon failure, similar to static applications.
- **Restart-delay n (5 seconds).** The minimum number of seconds between the respawning of failed instances of this application. This delay prevents a broken application from soaking up too much of the system.
- **singleThreshhold n (0).** An integer between 0 and 100, used to determine whether the last instance of a FastCGI application can terminate. If the process manager computed load factor for the application is lower than the specified threshold, the last instance is terminated. Specify a value closer to 1, to make your executables run in the idle mode for a long time. If memory or CPU time is a concern, a value closer to 100 is more applicable. A value of 0, prevents the last instance of an application from terminating; this value is the default. Changing this default is not recommended, especially if you set the -appConnTimeout option.
- **startDelay n (3 seconds).** The number of seconds the Web server waits while trying to connect to a dynamic FastCGI application. If the interval expires, the process manager is notified with hope that another instance of the application starts. Set the startDelay value smaller than the appConnTimeout value, to be effective.
- **updateInterval n (300 seconds).** The updateInterval decides how often statistical analysis is performed to determine the fate of dynamic FastCGI applications.

FastCGIExternalServer directive

The FastCGIExternalServer defines file name as an external FastCGI application.

It operates the same as the Fastcgiserver directive, except that the CGI application is running in another process outside of the Web server.

Syntax

```
FastCgiExternalServer file name -host hostnameport
[-appConnTimeout n] FastCgiExternalServer file name
-socket file name [-appConnTimeout n]
```

Scope
Default
Module
Multiple instances in the configuration file
Values

Server configuration
None
mod_fastcgi
yes

- **appConnTimeout *n* (0 seconds).** The number of seconds to wait for a connection to the FastCGI application to complete, or 0, to indicate use of a blocking `connect()` method. If the timeout expires, a `SERVER_ERROR` results. For non-zero values, this indicator is the amount of time used in a `select()` method to write to the file descriptor returned by a non-blocking `connect()` method. Non-blocking `connect()` methods are troublesome on many platforms. See also `-idle-timeout`; this option produces similar results, but in a more portable manner.
- **idle-timeout *n* (30 seconds).** The number of seconds of FastCGI application inactivity allowed before the request aborts and the event is logged (at the error `LogLevel`). The inactivity timer applies only as long as a connection is pending with the FastCGI application. If a request is queued to an application, but the application does not respond by writing and flushing within this period, the request aborts. If communication is complete with the application but incomplete with the client (a buffered response), the timeout does not apply.
- **flush none.** Force a write to the client as data is received from the application. By default, the `mod_fastcgi` option buffers data to free the application quickly.
- **host hostname:port none.** The hostname, or IP address and TCP port number (1-65535) the application uses for communication with the Web server. The `-socket` and `-host` options are mutually exclusive.
- **Pass-header header none.** The name of an HTTP Request Header passed in the request environment. This option makes the header contents available, to a CGI environment.
- **socket file name none.**
 - **On UNIX operating systems.** The file name of the UNIX domain socket the application uses for communication with the Web server. The file name is relative to the `FastCgIpcDir` option. The `-socket` and `-port` options are mutually exclusive.
 - **On Windows operating systems.** The name of the pipe the application uses for communicating with the Web server. The name is relative to the `FastCgIpcDir` option. The `-socket` and `-port` options are mutually exclusive.

FastCGIipcDir directive

The `FastCGIipcDir` directive specifies directory as the place to store the UNIX socket files used for communication between the applications and the Web server.

Syntax

- On UNIX platforms - `FastCgIpcDir directory`
- On Windows operating systems - `FastCgIpcDir name`

Scope	Server configuration
Default	None
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	directory or name

AIX **HP-UX** **Linux** **Solaris** The FastCgiIpcDir directive specifies directory as the place to store and find, in the case of external FastCGI applications, the UNIX socket files that are used for communication between the applications and the Web server. If the directory does not begin with a slash (/) then it is assumed to be relative to the ServerRoot. If the directory does not exist, the function attempts to create the directive with appropriate permissions. Specify a directory on a local file system. If you use the default directory, or another directory within /tmp, mod_fastcgi breaks if your system periodically deletes files from the /tmp directory.

Windows The FastCgiIpcDir directive specifies *name* as the root for the named pipes used for communication between the application and the Web server. Define the name in the form >\\.\pipe\pipename. . The pipename syntax can contain any character other than a backslash.

The FastCgiIpcDir directive must precede any FastCgiServer or FastCgiExternalServer directives, which make use of UNIX sockets. Ensure a readable, writeable, and executable directory by the Web server. No one should have access to this directory.

FastCGIServer directive

The FastCGIServer directive defines file name as a static FastCGI application.

The Process Manager starts one instance of the application with the default configuration specified in parentheses below. Should a static application instance die for any reason, the mod_fastcgi module spawns another instance for replacement and logs the event at the warn LogLevel.

Syntax	FastCgiServer file name [options]
Scope	Server configuration
Default	None
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	directory or name

You can use one of the following case-insensitive options:

- **appConnTimeout *n* (0 seconds)**. The number of seconds to wait for a connection to the FastCGI application to complete, or 0, to indicate use of a blocking connect(). If the timeout expires, a SERVER_ERROR results. For non-zero values, this indicator is the amount of time used in a select() to write to the file descriptor returned by a non-blocking connect(). Non-blocking connect()s prove troublesome on many platforms. See the -idle-timeout option; it produces similar results but in a more portable manner.
- **Idle-timeout *n* (30 seconds)**. The number of seconds of FastCGI application inactivity allowed before the request aborts and the event logs at the error LogLevel. The inactivity timer applies only when a pending connection with the FastCGI application exists. If an application does not respond to a queued request within this period, the request aborts. If communication completes with the application, but does not complete with the client (a buffered response), the timeout does not apply.
- **initial-env name [=value] none**. A name-value pair passed in the FastCGI application initial environment. To pass a variable from the Apache environment, do not provide the "=" (variables not actually in the environment, are defined without a value). To define a variable without a value, provide the "=" without a value. You can repeat this option.

- **init-start-delay *n* (1 second)**. The minimum number of seconds between the spawning of instances of this application. This delay decreases the demand placed on the system at server initialization.
- **Flush none**. Force a write to the client as data arrives from the application. By default, mod_fastcgi buffers data to free the application quickly.
- **Listen-queue-depth *n* (100)**. The depth of the listen() queue, also known as the backlog, shared by all of the instances of this application. A deeper listen queue enables the server to cope with transient load fluctuations, without rejecting requests; this option does not increase throughput. Adding additional application instances can increase throughput and performance, depending upon the application and the host.
- **Pass-header header none**. The name of an HTTP Request Header passed in the request environment. This option makes the contents of headers available to a CGI environment.
- **processes *n* (1)**. The number of application instances to spawn at server initialization.
- **Priority *n* (0)**. The process priority assigned to the application instances, using setpriority().
- **port *n* none**. The TCP port number (1-65535) the application uses for communication with the Web server. This option makes the application accessible from other machines on the network. The -socket and -port options are mutually exclusive.
- **Restart-delay *n* (5 seconds)**. The minimum number of seconds between the respawning of failed instances of this application. This delay prevents a broken application from using too many system resources.
- **Socket file name:**
 - On UNIX platforms: The file name of the UNIX domain socket that the application uses for communication with the Web server. The module creates the socket within the directory specified by FastCgiLpcDir. This option makes the application accessible to other applications, for example, cgi-fcgi on the same machine, or through an external FastCGI application definition, FastCgiExternalServer. If neither the -socket nor the -port options are given, the module generates a UNIX domain socket file name. The -socket and -port options are mutually exclusive.
 - On Windows operating systems: The name of the pipe for the application to use for communication with the Web server. The module creates the named pipe off the named pipe root specified by the FastCgiLpcDir directive. This option makes the application accessible to other applications, like cgi-fcgi on the same machine or through an external FastCGI application definition, FastCgiExternalServer. If neither the -socket nor the -port options are given, the module generates a name for the named pipe. The -socket and -port options are mutually exclusive. If the file name does not begin with a slash (/), then this file name is assumed relative to the ServerRoot.

Distributed operating systems

FastCGIsuEXEC directive

The FastCGIsuEXEC directive supports the suEXEC-wrapper.

Syntax	FastCgiSuexec <i>On</i> <i>Off</i> <i>file name</i>
Scope	Server configuration
Default	FastCgiSuexec <i>Off</i>
Module	mod_fastcgi
Multiple instances in the configuration file	yes
Values	The FastCgiSuexec directive requires suEXEC enabling in Apache for CGI. To use the same suEXEC-wrapper used by Apache, set FastCgiSuexec to On. To use a different suEXEC-wrapper, specify the file name of the suEXEC-wrapper. If the file name does not begin with a slash (/), then the file name is assumed relative to the ServerRoot.

When you enable the FastCgiSuexec directive, the location of static or external FastCGI application definitions becomes important. These differences inherit their user and group from the User and Group directives in the virtual server in which they were defined. User and Group directives should precede FastCGI application definitions. This function does not limit the FastCGI application to the virtual server in which it was defined. The application can service requests from any virtual server with the same user and group. If a request is received for a FastCGI application, without an existing matching definition running with the correct user and group, a dynamic instance of the application starts with the correct user and group. This action can lead to multiple copies of the same application running with a different user and group. If this causes a problem, preclude navigation to the application from other virtual servers, or configure the virtual servers with the same user and group.

See the Apache documentation for more information about suEXEC and the security implications.

Managing IBM HTTP Server remotely with the WebSphere Application Server administrative console

You can remotely administer and configure IBM HTTP Server using the WebSphere Application Server administrative console.

About this task

After you define a Web server definition in the WebSphere repository to represent the installed IBM HTTP Server, an administrator can administer and configure IBM HTTP Server through the WebSphere Application Server administrative console.

Administration includes the ability to start and stop the IBM HTTP Server. You can display and edit the IBM HTTP Server configuration file, and you can view the IBM HTTP Server error and access logs. The plug-in configuration file can be generated for IBM HTTP Server and propagated to the remote, or locally-installed, IBM HTTP Server.

Note: z/OS Administration and configuration using the WebSphere Application Server administrative console is available if IBM HTTP Server is on a managed node only. The node agent must be present to perform administration because there is no support for the IBM HTTP Server administration server.

Procedure

- **IBM HTTP Server remote administration with managed nodes:** When you install IBM HTTP Server on a remote machine with a managed node, the administration interface that handles requests between the administrative console and the IBM HTTP Server is the network deployment node agent.

Windows If you are planning on managing an IBM HTTP Server on a managed node (through nodeagent), configure the Windows service for IBM HTTP Server to *log on as local system account*. You can specify this during the installation using the create services panel.

- Distributed operating systems **IBM HTTP Server remote administration with unmanaged nodes:** When you install IBM HTTP Server on a remote machine *without* a managed node, the **administration server** is necessary for remote administration. The IBM HTTP Server installation includes the administration server, which installs by default during a typical IBM HTTP Server installation.

The administration server is the interface that handles requests between the administrative console and the remote IBM HTTP Server defined on the unmanaged node. The administration server must be started by a root user and defined to an unmanaged WebSphere Application Server node.

- Distributed operating systems **IBM HTTP Server remote administration using WebSphere Application Server Express and Base:** Administration function for IBM HTTP Server with the WebSphere Application Server Express or Base product requires installation and configuration of the administration server.

Extending IBM HTTP Server functionality with third-party plug-in modules

This section contains topics on using third-party plug-in modules with IBM HTTP Server.

Distributed operating systems

Before you begin

Modules that are loaded into IBM HTTP Server, whether distributed by IBM or a third-party vendor, must comply with the following specifications:

- The openssl library cannot be loaded by IBM HTTP Server plug-in modules.
- Plug-in modules provided by IBM may use the Global Security Kit (GSKit) library for SSL communications. These plug-in modules must comply with the GSKit restrictions for using a local GSKit installation to interoperate with the current release of IBM HTTP Server.

About this task

You can build third-party plug-in modules (dynamic shared object modules) for execution with IBM HTTP Server. IBM HTTP Server ships as an installation image with executables that you cannot rebuild because the source does not ship with the installation image. However, IBM HTTP Server does ship the header files necessary to compile and build third-party plug-in modules that execute as an IBM HTTP Server module.

Important: The use of third-party plug-in modules does not prevent IBM HTTP Server from being supported, but IBM cannot support the third-party plug-in module itself. If a problem occurs when the third-party plug-in module is loaded, IBM support might ask for the problem to be reproduced without the third-party plug-in module loaded, in order to determine if the problem is specific to the configuration with the third-party plug-in module. If a problem is specific to the configuration with the third-party plug-in module, the provider of that module might need to help determine the cause of the problem. IBM cannot resolve such problems without the involvement of the provider of the module, as this requires understanding of the implementation of the module, particularly with regard to its use of the Apache APIs.

Procedure

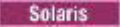


- Identify viable compilers. Apache and third-party plug-in module testing incorporated the compilers and compiler levels that are listed in this topic.
- **AIX** **HP-UX** **Linux** **Solaris** **z/OS** Determine the method to use to build the dynamic modules. Two common options for building dynamic modules are described in this topic.
- **Windows** Considerations for building dynamic modules. Restrictions apply when building a module to run with IBM HTTP Server. This topic describes the restrictions.

Viable compilers for Apache and third-party plug-in modules

There are many viable compilers and compiler levels, which have been tested, that you can use for Apache and third-party plug-in modules.

Apache modules and third-party module testing incorporated the compilers and compiler levels that are included in the following list. Other compilers may work, but testing was limited to the following environments:

- **AIX** AIX V5.0.2.3: C or VisualAge® C++ Professional
- **HP-UX** HP_UXaC++ Compiler (A.03.xx)
- **Linux** Linux platforms:
 - Linux on Intel: gcc 3.3.3

- Linux on POWER®: gcc 3.3.3
- Linux on zSeries®: gcc 3.3.3
-  SunWorkShop V5.0
-  Microsoft Visual C++ 6.0
-  z/OS V1R6.0 C/C++

The primary concern with determining if a different compiler can be used is when the third-party module, or libraries it uses, are implemented in C++. Different compiler versions may use different C++ application binary interfaces (ABI), in which case the behavior is undefined.

Build method options for dynamic modules

There are two common methods you can use to build dynamic modules: Apache extension tool (apxs) and module-provided configuration scripts.

The two common options for building dynamic modules are:

- **Apache extension tool (apxs).** IBM HTTP Server provides the apxs tool for building dynamic modules. You can build and install most modules with apxs. Here is an example:

```
# /usr/IBMIHS/bin/apxs -ci mod_example.c
```

To use the apxs tool, verify that Perl V5.004 or later is installed and that the path to the Perl executable on the first line of apxs is correct. See Apache APXS for more information.

- **Module-provided configuration scripts.** Some complex modules cannot be built directly with apxs, and instead provide their own configuration scripts for building the module. Consult the documentation provided with the module for detailed instructions. Check for special configuration options that must point to the IBM HTTP Server installation directory, or the apxs program installed with IBM HTTP Server.

The configuration scripts for some modules check specifically for the use of Apache HTTP Server and will not work properly with IBM HTTP Server. In that case, install Apache V2.2.8 and build the module for Apache V2.2.8, then use the resulting dynamic module (mod_example.so) with IBM HTTP Server.

IBM HTTP Server customers occasionally try to use third-party modules which do not build or run properly on their platform with either Apache HTTP Server or IBM HTTP Server. Whenever there are build or run-time concerns with third-party modules, first verify that it builds and operates properly with Apache HTTP Server on the same machine. If problems are encountered with Apache HTTP Server, the module cannot be expected to work with IBM HTTP Server.

Considerations for building dynamic modules on Windows platforms

There are restrictions that you must consider when building dynamic modules for Windows platforms.

The following restrictions apply when building a module to run with IBM HTTP Server:

- Link your dynamic module to the libraries that are contained in `lib` directory where the server is installed.
- The Apache HTTP Server module API is defined by the header files that are contained in the `include` directory where the server is installed. Your module should include any of these header files as needed.
- You must not modify any file or data structure that is contained in any file in the `include` directory where the server is installed.

Chapter 2. Administering and configuring the administration server

Learn how to administer and configure the administration server.

Starting and stopping the IBM HTTP Server administration server

This topic describes how to start and stop the IBM HTTP Server administration server on distributed platforms.

Before you begin

You can set up the IBM HTTP Server administration server when you install IBM HTTP Server. For more information see “Installing IBM HTTP Server using the GUI” on page 45.

About this task

Start the IBM HTTP Server administration server as follows.

Procedure

- **Windows** From the Start menu:

- Click **Start > Programs > IBM HTTP Server > Start Administration Server**. A message box displays that indicates the server has started.
- If the IBM HTTP Server administration server does not start, complete the following steps:
 1. Open the Control Panel.
 2. Click **Services**.
 3. Double-click IBM HTTP Server Administration Server to start the server.

Confirm that IBM HTTP Server administration server started successfully by checking the `admin_error.log` file for a "start successful" message. If you use the developer installation option, then the IBM HTTP Server administration server does not install as a service. You have to run the `httpd.exe` file from a command line with the `-f` option. From the default directory, type:

```
httpd -f conf\admin.conf
```

- **AIX** **HP-UX** **Linux** **Solaris** **The adminctl command starts and stops the IBM HTTP Server administration server.** You can find the `adminctl` command in the `bin` subdirectory, within the IBM HTTP Server installation directory. If that directory is not in your `PATH`, the full path should be given on the command line. Start or stop the IBM HTTP Server administration server using the default `admin.conf` configuration file as follows:

1. Run the `adminctl start` command to start the server or run the `adminctl stop` command to stop the server. Issue the commands from the default directories, based on your operating system:

- **AIX** `/usr/IBM/HTTPServer/bin/adminctl start|stop`

- **HP-UX** **Linux** **Solaris** `/opt/IBM/HTTPServer/bin/adminctl start|stop`

For example, The `adminctl` command is not in your `PATH`, the IBM HTTP Server installation directory is `/usr/IBM/HTTPServer`, and the default configuration file is used as follows:

```
# /usr/IBM/HTTPServer/bin/adminctl start
# /usr/IBM/HTTPServer/bin/adminctl stop
```

Important: The `admin.conf` configuration file supports single-byte characters (SBCS) only.

2. Confirm that IBM HTTP Server administration server started successfully by checking the `admin_error.log`.

Protecting access to the IBM HTTP Server administration server

This section describes topics on controlling access to the administration server in order to protect IBM HTTP Server configuration files.

About this task

The WebSphere Application Server administrative console can administer a remote IBM HTTP Server, on an unmanaged node, using IBM HTTPS Server administration server as the interface. Refer to the following topics for controlling access to the administration server in order to protect IBM HTTP Server configuration files.

Procedure

- Enable access to the administration server using the `htpasswd` utility
- Run the `setupadm` script for the administration server
- Set permissions manually for the administration server

Enabling access to the administration server using the `htpasswd` utility

The administration server is installed with authentication enabled. This means that the administration server will not accept a connection without a valid user ID and password. This is done to protect the IBM HTTP Server configuration file from unauthorized access.

Procedure

Launch the **htpasswd** utility that is shipped with the administration server. This utility creates and updates the files used to store user names and password for basic authentication of users who access your Web server. Locate **htpasswd** in the `bin` directory.

- **Windows** `htpasswd -cm <install_dir>\conf\admin.passwd [login name]`
- **AIX** **HP-UX** **Linux** **Solaris** `./htpasswd -cm <install_dir>/conf/admin.passwd [login name]`

where `<install_dir>` is the IBM HTTP Server installation directory and `[login name]` is the user ID that you use to log into the administration server.

Results

The password file is referenced in the `admin.conf` file with the `AuthUserFile` directive. For further information on authentication configuration, see the Apache Authentication, Authorization and Access Control documentation.

Running the `setupadm` command for the administration server

Run the `setupadm` command if you need to configure the administration server manually or you need to modify its configuration.

Before you begin

You might need to configure the administration server if you did not configure it during the IBM HTTP Server installation process or you completed a non-administrator installation.

Procedure

1. Optional: Change the user ID and password that WebSphere Application Server uses to authenticate to the administration server. If you need to change the user ID and password, use the `htpasswd` utility. For more information, see the documentation about enabling access to the administration server using the `htpasswd` utility.
2. Use the `IHS_HOME/bin/setupadm` command to set up the administration server. You can set up the administration server to run in the following scenarios:
 - A non-root user and group, which the IBM HTTP Server Administration Server will run as when started by root
 - A non-root group
 - The path to the IBM HTTP Server configuration file
 - The path to the IBM HTTP Server administration server configuration file
 - The plug-in configuration file for WebSphere Application Server

Command syntax

```
setupadm [-silent] [-create] -usr user_name  
-grp group_name -cfg IBM_HTTP_Server_configuration_file  
[-plg plug-in_configuration_file]  
-adm administration_server_configuration_file
```

-silent This parameter enables the `setupadm` command to run without message text.

-create

This parameter specifies that you want to create a user and group. If you do not specify this parameter, the values for the **-usr** and **-grp** parameters must exist.

-usr This parameter specifies the user ID that will run the administration server. This user ID value is updated in the `<User>` directive within the administration server configuration file, which is called `admin.conf`.

-grp This parameter specifies the group name that will run the administration server. When you specify a value, it is used to change the file permissions for the configuration files and the user or group authentication files. This group name value is updated in the `<Group>` directive within the administration server configuration file, which is called `admin.conf`.

Ensure that you specify a unique group name for the administration server.

-cfg This parameter defines the fully qualified path to the IBM HTTP Server web server configuration file. Within this file, the permission and group information is updated.

Note: The administration server requires both read and write access to IBM HTTP Server configuration files.

-plg This parameter specifies the fully qualified path to the `plugin-cfg.xml` configuration file. Within this file, permissions are changed.

-adm This parameter specifies the fully qualified path to the administration server configuration file. If you do not specify this parameter, a default administration configuration file is used that is based on the `install_root/conf/admin.conf` file.

Results

When you run the `setupadm` command, the following actions occur:

- Creates a new user and group is created on the system if you specify the **-create** parameter
- Changes the group owner of the configuration files to the group name that you specify and grants group write permissions to those files. This process allows the administration server to modify those configuration files.
- Updates the administration server configuration file with the user name and group name.

- Creates a backup file each time that you run the command.

What to do next

Complete the steps to start the IBM HTTP Server administration server. For more information, see the documentation about starting and stopping the IBM HTTP Server administration server.

The IBM HTTP Server administration server has to be started under the same user ID as the IBM HTTP Server to be able to restart it with **apachectl restart**.

Setting permissions manually for the administration server

For IBM HTTP Server administration server, the setupadm script creates users and groups and sets file permissions for them. This topic describes how to do this manually.

About this task

Perform the following steps to create users and groups and set file permissions.

Procedure

- Create a new user and unique group for the IBM HTTP Server administration server.

– AIX

1. Launch SMIT.
2. Click **Security and Users**.
3. Click **Groups > Add a Group**.
4. Enter the group name, for example, **admingrp**.
5. Click **OK**. Go back to **Security and Users**.
6. Click **Users > Add a User**.
7. Enter the user name, for example, **adminuser**.
8. Enter the primary group you just created.
9. Click **OK**.

– HP-UX

Linux

- Run the following command from a command line:

```
groupadd <group_name>  
useradd -g <group_name> <user_ID>
```

– Solaris

1. Launch the administration tool.
2. Click **Browse > Groups**.
3. Click **Edit > Add**.
4. Enter the group name, for example, **admingrp**.
5. Click **OK**.
6. Click **Browse > Users**.
7. Click **Edit > Add**.
8. Enter the user name, for example, **adminuser** and the primary group name, for example, **admingrp**.
9. Click **OK**.

- **AIX** **HP-UX** **Linux** **Solaris** Updating file permissions.

Once you have created a user and group, set up file permissions as follows:

1. Update the permissions for the targeted IBM HTTP Server conf directory.
 - a. At a command prompt, change to the directory where you installed IBM HTTP Server.

- b. Type the following commands:

```
chgrp <group_name> <directory_name>  
chmod g+rw <directory_name>
```
2. Update the file permission for the targeted IBM HTTP Server configuration files.
 - a. At a command prompt, change to the directory that contains the configuration files.
 - b. Type the following commands:

```
chgrp <group_name> <file_name>  
chmod g+rw <file_name>
```
3. Update the admin.conf configuration file for the IBM HTTP Server administration server.
 - a. Change to the IBM HTTP Server administration server admin.conf directory.
 - b. Search for the following lines in the admin.conf file:

```
User nobody  
Group nobody
```
 - c. Change those lines to reflect the user ID and unique group name you created. For example:

```
User userID  
Group group_name
```
4. Update the file permission for the targeted plug-in configuration files.
 - a. At a command prompt, change to the directory that contains the plug-in configuration files.
 - b. Type the following commands:

```
chgrp <group_name> <file_name>  
chmod g+rw <file_name>
```

Results

You have set up read and write access for the configuration and authentication files. Now you can perform Web server configuration data administration.

Chapter 3. Migrating and installing IBM HTTP Server

Learn how to establish the product in new and existing environments, including planning, preparing for, and completing product installations.

Installing, updating, rolling back, and uninstalling IBM HTTP Server

IBM Installation Manager is a common installer for many IBM software products that you use to install, update, roll back, and uninstall IBM HTTP Server.

Before you begin

Restrictions:

- Before you can successfully install IBM HTTP Server, ensure that your environment meets the prerequisites for the application server. For more information, see the Preparing the operating system for product installation topic.
- **HP-UX** **Solaris** The Installation Manager GUI is not supported on HP-UX PA-RISC and Solaris 10 x64 systems. Perform the following actions to install or uninstall IBM HTTP Server on these systems:
 - Use the Installation Manager GUI on a supported system to record a response file that will allow you to install or uninstall IBM HTTP Server silently.
 - Edit the recorded response file if necessary.
 - Use the response file to install or uninstall IBM HTTP Server silently on your system.
- **Linux** For any Linux system that is enabled for Security Enhanced Linux (SELinux), such as Red Hat Enterprise Linux Version 5 or SUSE Linux Enterprise Server Version 11, you must identify the Java shared libraries in the Installation Manager 1.4.2 or later installation image to the system. Also, you must identify the Java shared libraries in the Installation Manager 1.4.2 or later installation after it has been installed. For example:

```
chcon -R -t texrel_shlib_t ${IM_Image}/jre_5.0.3.sr8a_20080811b/jre/bin
chcon -R -t texrel_shlib_t ${IM_Install_root}/eclipse/jre_5.0.3.sr8a_20080811b/jre/bin
```

About this task

Complete one of these procedures to install, update, roll back, or uninstall IBM HTTP Server using Installation Manager.

Procedure

- “Installing IBM HTTP Server using the GUI” on page 45
- “Installing IBM HTTP Server silently” on page 55
- “Updating IBM HTTP Server” on page 55
- “Uninstalling IBM HTTP Server using the GUI” on page 65

Results

Notes on logging and tracing:

- An easy way to view the logs is to open Installation Manager and go to **File > View Log**. An individual log file can be opened by selecting it in the table and then clicking the **Open log file** icon.
- Logs are located in the logs directory of Installation Manager's application data location. For example:

- **Windows** **Administrative installation:**

```
C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager
```

- **Windows** **Non-administrative installation:**

```
C:\Documents and Settings\user_name\Application Data\IBM\Installation Manager
```

- **AIX** **HP-UX** **Linux** **Solaris** **Administrative installation:**

/var/ibm/InstallationManager

- **AIX** **HP-UX** **Linux** **Solaris** **Non-administrative installation:**

user_home/var/ibm/InstallationManager

- The main log files are time-stamped XML files in the logs directory, and they can be viewed using any standard Web browser.
- The log.properties file in the logs directory specifies the level of logging or tracing that Installation Manager uses.

Notes on troubleshooting:

- **HP-UX** By default, some HP-UX systems are configured to not use DNS to resolve host names. This could result in Installation Manager not being able to connect to an external repository.

You can ping the repository, but nslookup does not return anything.

Work with your system administrator to configure your machine to use DNS, or use the IP address of the repository.

- In some cases, you might need to bypass existing checking mechanisms in Installation Manager.
 - On some network file systems, disk space might not be reported correctly at times; and you might need to bypass disk-space checking and proceed with your installation.

To disable disk-space checking, specify the following system property in the config.ini file in *IM_install_root/eclipse/configuration* and restart Installation Manager:

```
cic.override.disk.space=sizeunit
```

where *size* is a positive integer and *unit* is blank for bytes, k for kilo, m for megabytes, or g for gigabytes. For example:

```
cic.override.disk.space=120 (120 bytes)
cic.override.disk.space=130k (130 kilobytes)
cic.override.disk.space=140m (140 megabytes)
cic.override.disk.space=150g (150 gigabytes)
cic.override.disk.space=true
```

Installation Manager will report a disk-space size of Long.MAX_VALUE. Instead of displaying a very large amount of available disk space, N/A is displayed.

- To bypass operating-system prerequisite checking, add `disableOSPrereqChecking=true` to the config.ini file in *IM_install_root/eclipse/configuration* and restart Installation Manager.

If you need to use any of these bypass methods, contact IBM Support for assistance in developing a solution that does not involve bypassing the Installation Manager checking mechanisms.

- Installation Manager might display a warning message during the uninstallation process.

Uninstalling IBM HTTP Server using Installation Manager requires that the data repositories remain valid and available.

A warning message is displayed by Installation Manager to alert you when repositories are not available or connected. A similar warning message might display after you add or modify data repository connection preferences in Installation Manager.

If Installation Manager detects missing data repositories or fails to connect to repositories during the uninstallation process, complete the following actions:

1. Click **Cancel** to end the uninstallation task.
2. Select **File > Preferences > Repositories**, and add the appropriate data repositories that you can connect to successfully.
3. Exit Installation Manager.
4. Restart Installation Manager.
5. Uninstall IBM HTTP Server.

- For more information on using Installation Manager, read the IBM Installation Manager Information Center.

Read the release notes to learn more about the latest version of Installation Manager. To access the release notes, complete the following task:

- **Windows** Click **Start > Programs > IBM Installation Manager > Release Notes®**.
- **AIX** **HP-UX** **Linux** **Solaris** Go to the documentation subdirectory in the directory where Installation Manager is installed, and open the `readme.html` file.

Installing IBM HTTP Server using the GUI

You can use the Installation Manager GUI to install IBM HTTP Server.

Before you begin

Install Installation Manager:

1. Perform one of the following procedures:

- If you want to use the Installation Manager that is included with this product, perform the following actions:

a. Obtain the necessary files from the physical media or the web.

There are three basic options for obtaining and installing Installation Manager and the product.

– **Access the physical media, and use local installation**

You can access Installation Manager and the product repositories on the product media. You can install Installation Manager on your system and use it to install the product from the product repositories on the media.

– **Download the files from the Passport Advantage® site, and use local installation**

Licensed customers can download Installation Manager as well as the necessary product repositories from the Passport Advantage site. You can then install Installation Manager on your system and use it to install the product from the repositories.

– **Download a file from the Installation Manager website, and use web-based installation**

You can download and unpack a compressed file containing Installation Manager from the IBM Installation Manager website. You can then install Installation Manager on your local system and use it to install the product from the web-based repository located at

<http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80>

b. Change to the location containing the Installation Manager installation files, and run one of the following commands:

Administrative installation:

- **Windows** `install.exe`
- **AIX** **HP-UX** **Linux** **Solaris** `./install`

Non-administrative installation:

- **Windows** `userinst.exe`
- **AIX** **HP-UX** **Linux** **Solaris** `./userinst`

Group-mode installation:

- **AIX** **HP-UX** **Linux** **Solaris** `./groupinst`

Notes on group mode:

- Group mode allows multiple users to use single instance of IBM Installation Manager to manage software packages. This does not mean that two people can use the single instance of IBM Installation Manager at the same time.

- **Windows** Group mode is not available on Windows operating systems.
- If you do not install Installation Manager using group mode, you will not be able to use group mode to manage any of the products that you install later using this Installation Manager.
- Make sure that you change the installation location from the default location in the current user's home directory to a location that is accessible by all users in the group.
- Set up your groups, permissions, and environment variables as described in the Group mode road maps in the IBM Installation Manager Information Center before installing in group mode.
- For more information on using group mode, read the Group mode road maps in the IBM Installation Manager Information Center.

The installer opens an **Install Packages** window.

- c. Make sure that the Installation Manager package is selected, and click **Next**.
 - d. Accept the terms in the license agreements, and click **Next**.
The program creates the directory for your installation.
 - e. Click **Next**.
 - f. Review the summary information, and click **Install**.
 - If the installation is successful, the program displays a message indicating that installation is successful.
 - If the installation is not successful, click **View Log File** to troubleshoot the problem.
- If you already have a version of Installation Manager installed on your system and you want to use it to install and maintain the product, obtain the necessary product files from the physical media or the web.

There are three basic options for installing the product.

– **Access the physical media, and use local installation**

You can access the product repositories on the product media. Use your existing Installation Manager to install the product from the product repositories on the media.

– **Download the files from the Passport Advantage site, and use local installation**

Licensed customers can download the necessary product repositories from the Passport Advantage site. You can then use your existing Installation Manager to install the product from the repositories.

– **Access the live repositories, and use web-based installation**

You can install Installation Manager on your local system and use it to install the product from the web-based repository located at

<http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80>

Whenever possible, you should use the remote web-based repositories so that you are accessing the most up-to-date installation files.

2. Add the product repository to your Installation Manager preferences.
 - a. Start Installation Manager.
 - b. In the top menu, click **File > Preferences**.
 - c. Select **Repositories**.
 - d. Perform the following actions:
 - 1) Click **Add Repository**.
 - 2) Enter the path to the `repository.config` file in the location containing the repository files.

For example:

- **Windows** C:\repositories\product_name\local-repositories
- **AIX** **HP-UX** **Linux** **Solaris** /var/repositories/product_name/local-repositories

or

<http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80>

- 3) Click **OK**.
- e. Deselect any locations listed in the Repositories window that you will not be using.
- f. Click **Apply**.
- g. Click **OK**.
- h. Click **File > Exit** to close Installation Manager.

About this task

Complete this procedure to use the Installation Manager GUI to install IBM HTTP Server.

Procedure

1. Start Installation Manager.
2. Click **Install**.

Note: If you are prompted to authenticate, use the IBM ID and password that you registered with on the program website.

Installation Manager searches its defined repositories for available packages.

3. In the **Install Packages** window, complete the appropriate actions.
 - a. Select **IBM HTTP Server** and the appropriate version.

Note: If you are installing the trial version of this product, select **IBM HTTP Server Trial**.

If you already have IBM HTTP Server installed on your system, a message displays indicating that IBM HTTP Server is already installed. To create another installation of IBM HTTP Server in another location, click **Continue**.

- b. Click **Next**.

Note: If you try to install a newer level of IBM HTTP Server with a previous version of Installation Manager, Installation Manager might prompt you to update to the latest level of Installation Manager when it connects to the repository. Update to the newer version before you continue if you are prompted to do so. Read *Installing updates* in the Installation Manager information center for information about automatic updates.

4. Accept the terms in the license agreements, and click **Next**.
5. On the Location panel, specify the installation root directory for the product binaries, which are also referred to as the core product files or system files.

The panel also displays the shared resources directory and disk-space information.

The core product files do not change unless you install maintenance.

Restrictions:

- Deleting the default target location and leaving an installation-directory field empty prevents you from continuing.
- Do not use symbolic links as the destination directory.
Symbolic links are not supported.
- Do not use spaces in the name of the installation directory.
These spaces are not supported.

- Do not use a semicolon in the directory name.
IBM HTTP Server cannot install properly if the target directory includes a semicolon.
 - **Windows** A semicolon is the character used to construct the class path on Windows systems.
 - **Windows** The maximum path length on the Windows XP, Windows Vista, and Windows 7 operating systems is 60 characters.
6. Click **Next**.
 7. **AIX** **Linux** **Solaris** If you are installing on a 64-bit system, choose between a 32-bit or 64-bit HTTP Server environment and click **Next**.

Notes:

- This option displays only if you are installing on a 64-bit system.
 - This does not apply to Solaris x86 64-bit systems.
 - You must select one of the two options.
 - You cannot modify this installation later and change this selection.
8. Click **Next** to display the Configuration for IBM HTTP Server panel,
 9. **Windows** On the Configuration for IBM HTTP Server panel, specify your Web server configuration.
 - Specify a port number on which IBM HTTP Server will communicate. The default port is 80.
 - Choose whether to use a Windows service to run IBM HTTP Server.

Note: You have the option to create a Windows service for IBM HTTP Server on this panel. You can configure the services to run as local system account or a user ID that you specify. The user ID requires the following advanced user rights:

- Act as part of the operating system
- Log on as a service

Important: If you do not select **Run IBM HTTP Server as a Windows Service**, this instance of IBM HTTP Server cannot be stopped or started by the WebSphere Application Server administrative console. At any time after installation, you can create a new service by running the following command:

```
ihp_root/bin/httpd.exe -n new_service_name -k install
```

and then updating the web server definition in the administrative console to reflect the new service name.

- Determine if your startup type will be automatic or manual.
10. **AIX** **HP-UX** **Solaris** On the Configuration for IBM HTTP Server panel, specify your Web server configuration.
Specify a port number on which IBM HTTP Server will communicate. The default port is 80.
 11. Click **Next**.
 12. Review the summary information, and click **Install**.
 - If the installation is successful, the program displays a message indicating that installation is successful.

Note: The program might also display important post-installation instructions as well.

 - If the installation is not successful, click **View Log File** to troubleshoot the problem.
 13. Click **Finish**.
 14. Click **File > Exit** to close Installation Manager.

Results

If the installation is successful, the IBM HTTP Server product is installed and the log file is located in the `/logs/install/` directory. However, if the product installation fails, see the `log.txt` file in either the `/logs/install/` directory or the `$USER/ihslogs/` directory.

What to do next

Set up IBM HTTP Server administration authentication, using the `htpasswd` utility.

You can get started using Secure Sockets Layer (SSL) connections by making only a few configuration changes, as described in [Securing with SSL communications](#).

AIX **Windows** You can configure the Fast Response Cache Accelerator to boost performance.

You can also make other configuration changes with Apache directives.

Mounting CD-ROMS on AIX, HP-UX, Linux and Solaris systems

This section describes how to mount the CD-ROM for IBM HTTP Server on AIX, HP-UX, Linux and Solaris operating systems.

Before you begin

After inserting a CD-ROM into a drive, some operating systems require you to mount the drive.

About this task

Use these procedures to mount the product discs for IBM HTTP Server.

Procedure

- **AIX** **Mount the CD-ROM using the System Management Interface Tool (SMIT) as follows:**
 1. Log in as a user with root authority.
 2. Insert the CD-ROM in the drive.
 3. Create a CD-ROM mount point by entering the `mkdir -p /cdrom` command, where `cdrom` represents the CD-ROM mount point directory.
 4. Allocate a CD-ROM file system using SMIT by entering the **smit storage** command.
 5. After SMIT starts, click **File Systems > Add / Change / Show / Delete File Systems > CDROM File Systems > Add CDROM File System**.
 6. In the Add a File System window:
 - Enter a device name for your CD-ROM file system in the **DEVICE Name** field. Device names for CD-ROM file systems must be unique. If there is a duplicate device name, you may need to delete a previously-defined CD-ROM file system or use another name for your directory. The example uses `/dev/cd0` as the device name.
 - Enter the CD-ROM mount point directory in the **MOUNT POINT** window. In our example, the mount point directory is `/cdrom`.
 - In the **Mount AUTOMATICALLY at system restart** field, select yes to enable automatic mounting of the file system.
 - Click **OK** to close the window, then click **Cancel** three times to exit SMIT.
 7. Next, mount the CD-ROM file system by entering the **smit mountfs** command.
 8. In the Mount a File System window:
 - Enter the device name for this CD-ROM file system in the **FILE SYSTEM name** field. In our example, the device name is `/dev/cd0`.

- Enter the CD-ROM mount point in the **Directory over which to mount** field. In our example, the mount point is `/cdrom`.
- Enter `cdrfs` in the **Type of Filesystem** field. To view the other kinds of file systems you can mount, click List.
- In the **Mount as READ-ONLY system** field, select yes.
- Accept the remaining default values and click **OK** to close the window.

Your CD-ROM file system is now mounted. To view the contents of the CD-ROM, place the disk in the drive and enter the `cd /cdrom` command where `cdrom` is the CD-ROM mount point directory.

- **HP-UX** **Mount the CD-ROM.** Because WebSphere Application Server contains several files with long file names, the mount command can fail. The following steps let you successfully mount your WebSphere Application Server product CD-ROM.

1. Log in as a user with root authority.
2. In the `/etc` directory, add the following line to the `pfs_fstab` file:


```
/dev/dsk/c0t2d0 mount_point pfs-rrip ro,hard
```

where `mount_point` represents the mount point of the CD-ROM.

3. Start the `pfs` daemon by entering the following commands (if they are not already running):

```
/usr/sbin/pfs_mountd &
/usr/sbin/pfsd 4 &
```

4. Insert the CD-ROM in the drive and enter the following commands:

```
mkdir /cdrom
/usr/sbin/pfs_mount /cdrom
```

The `/cdrom` variable represents the mount point of the CD-ROM.

5. Log out.

- **Linux** **Mount the CD-ROM using the following steps.**

1. Log in as a user with root authority.
2. Insert the CD-ROM in the drive and enter the following command:

```
mount -t iso9660 -o ro /dev/cdrom /cdrom
```

The `/cdrom` variable represents the mount point of the CD-ROM.

3. Log out.

Some window managers can automatically mount your CD-ROM for you. Consult your system documentation for more information.

- **Solaris** **Mount the CD-ROM using the following steps.**

1. Log in as a user with root authority.
2. Insert the CD-ROM into the drive.
3. If the Volume Manager is not running on your system, enter the following commands to mount the CD-ROM:

```
mkdir -p /cdrom/unnamed_cdrom
mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/unnamed_cdrom
```

The `/cdrom/unnamed_cdrom` variable represents the CD-ROM mount directory and the `/dev/dsk/c0t6d0s2` represents the CD-ROM drive device.

If you are mounting the CD-ROM drive from a remote system using NFS, the CD-ROM file system on the remote machine must be exported with root access. You must also mount that file system with root access on the local machine.

If the Volume Manager (`vold`) is running on your system, the CD-ROM is automatically mounted as:

```
/cdrom/unnamed_cdrom
```

4. Log out.

What to do next

Return to the installation procedure to continue.

Installing IBM HTTP Server with a non-administrator user ID

The common way to install IBM HTTP Server is to run the installation program using an administrator user ID. However, it is sometimes necessary to install IBM HTTP Server using a non-administrator (non-root) user ID. On UNIX platforms, the user ID that performs the installation of IBM HTTP Server must be the same user ID that will start the IBM HTTP Server. When IHS is started by the root, which is the recommended configuration, the user ID is changed to an unprivileged user ID for the bulk of the runtime processing.

Before you begin

You must remember to perform all of the following actions using the same user ID, whether it is an administrator user ID or non-administrator user ID:

- Install IBM HTTP Server
- Install the WebSphere plug-in
- Start and stop the IBM HTTP Server
- Start and stop the IHS administration server (when applicable)
- Start and stop the WebSphere node agent (when applicable)

Administering a non-root IBM HTTP Server from WebSphere

If you are installing a non-root IBM HTTP Server as a local Web server, it may be managed by WebSphere only if the following constraints are met:

- If the Web server is to be managed by the IHS Administration Server, you must manually run the `setupadm` command. The `setupadm` command is run in the `<IHS_HOME>/bin` directory so that you can properly administer the administrative server with the WebSphere Application Server. The format for the command is as follows:

```
setupadm -usr <userName>  
-grp <groupName>  
-cfg <IHS Web server configuration file>  
-adm <IHS administrative server configuration file>  
-plg <plug-in configuration file>
```
- If the Web server is to be managed by the IHS Administration Server, the value you specify for the administration server port value must be greater than 1024. An IBM HTTP Server that is running under a non-root user ID does not start if the port number for its listener port is 1024 or less.
- If the Web server is a local Web server managed by a WebSphere node agent, the user ID used to install and start IBM HTTP Server is the same user ID under which the non-root node agent is running.

About this task

Launching the IBM HTTP Server installation program is done the same way for a non-root installation as it is for a root installation, but there are several installation steps that require root privileges that cannot be completed or must be completed separately. Complete the installation steps as follows:

Procedure

- **Register the installed program with the operating system.** This cannot be done for a non-root installation. The non-root IBM HTTP Server installation is not listed when using operating system facilities to display installed programs.
- **Windows** **Create the Windows service entries for IBM HTTP Server and IBM HTTP Administration Server.** This cannot be done for a non-root installation. Neither of these service entries are created and IBM HTTP Server cannot start as a service.

Start IBM HTTP Server as follows:

```
<ihs_install_directory>/bin/httpd.exe
```

Start the IBM HTTP administration server as follows:

```
<ihs_install_directory>/bin/httpd.exe -f <ihs_install_directory>/conf/admin.conf
```

Stop IBM HTTP server as follows:

- Press Control+C in the IBM HTTP Server window, or
- End the httpd.exe processes using the Windows Task Manager
- **Windows** **Create an entry in Start > Programs.** This cannot be done for a non-root installation. No entries are created.
- **Windows** **Create an entry in Add/Remove programs.** This cannot be done for a non-root installation. No entry is created.
- **Windows** **Install AFP.** This cannot be done for a non-root installation. Do not enable AFP for the non-administrator IBM HTTP Server installation, even if AFP is already installed from a previous administrator installation. Only enable AFP for one instance of IBM HTTP Server.
- **Silent installations.** To enable a non-root installation, add the following option to the silent installation response file:

```
<data key='user.ihs.allowNonRootSilentInstall' value='true'/>
```

What to do next

Uninstall a non-root installation of IBM HTTP as follows:

AIX

HP-UX

Linux

Solaris

```
<ihs_install_directory>/uninstall/uninstall
```

Windows

```
<ihs_install_directory>\uninstall\uninstall
```

Creating multiple instances of IBM HTTP Server on Windows operating systems

On Windows operating systems, you can create multiple instances of IBM HTTP Server by manually creating additional service names.

Before you begin

About this task

When you install IBM HTTP Server, you create one IBM HTTP Server as a Windows service with a default name. If you need to run with more than one IBM HTTP Server instance, you can manually create additional service names.

Procedure

1. Install a new service name. Use the httpd.exe program, which is located in the bin directory of the IBM HTTP Server installation. The command syntax for installing a new service name is:

```
httpd -k install -n <new_service_name> -f  
  <path_to_new_configuration_file>
```

This command allows you to associate a unique configuration file with each service name.

2. Specify different IP addresses or ports in the Listen directives of each configuration file and specify different log file names.
3. Optional: Change settings of the new service using the Windows Services control panel. The new service name will have "Log On" set to "Local System Account" and will have "Startup Type" set to "Automatic." You can change these default settings using the Windows Services control panel. It might be necessary to change the "Log On" setting of the new service name to match the "Log On" of the main installation in order to ensure that file permissions will allow the new service name to run.

4. Disable the Fast Response Cache Accelerator (FRCA). When running multiple instances of IBM HTTP Server, you must disable the FRCA (AFPA directives) in all configuration files.

What to do next

After creating a new service name, you can add it to the WebSphere Administration Server administrative console by creating a new Web server definition and specifying the new service name and the path to the new configuration file.

The syntax for uninstalling an existing service name is:

```
httpd -k uninstall -n <service_name>
```

Running multiple instances of IBM HTTP Server from a single install

Run multiple, independent instances of IBM HTTP Server from a single installation. It is seldom necessary to run multiple instances, as features like virtual hosts allow a single instance to efficiently serve many sites, but in some cases it is necessary. If you need to securely administer your sites by different administrators, for example, you must run separate instances that each use their own configuration files.

Before you begin

This topic is primarily for AIX, HP-UX, Linux, Solaris, and Windows operating systems. On the z/OS platform, the `install_ihs` command creates a separate directory for each instance without creating another copy of the product. See the z/OS topic for configuring IBM HTTP Server for more information.

Before configuring multiple instances, consider if your problem can be solved by using virtual hosts and/or having IBM HTTP Server listen on multiple addresses and ports. The advantage of a single instance is that it uses less resources to serve the same requests as multiple instances.

Note: When you follow the examples, change "this_instance" to a unique name for each instance.

Procedure

1. Create a separate main configuration file, normally the `httpd.conf` file, for each instance.

Note: To reduce duplication, store common directives in common files and import these into the separate, main configuration files with the *Include* directive.

We'll call the configuration file `conf/this_instance.conf` for the rest of these steps.

Here is a simple example of a configuration file for an instance:

```
Listen 10.0.0.1:80
PidFile instance1/httpd.pid
ErrorLog instance1/error.log
CustomLog instance1/access.log common
# Other directives that make this instance behave uniquely
Include conf/common.conf
```

A real configuration file would have more directives in it to make this instance behave differently than the other instances.

2. Configure the port settings in the configuration files. You cannot use a combination of listen port and listen IP address for more than one instance. Check the Listen directives in each configuration file, and verify that they are unique. See information on the Listen directive for Apache HTTP Server for more information.
3. Configure settings for logging and other special files. Any files that are normally stored in the `install_root/logs` directory cannot be shared between instances. Each instance must have unique values for the following directives:

PidFile

Applicable to all configurations. See the information on the PidFile directive for Apache HTTP Server.

ScriptSock

Applicable to non-Windows configurations with mod_cgid enabled.

ErrorLog

Applicable to all configurations. See the information on the ErrorLog directive for Apache HTTP Server.

CustomLog or TransferLog

Applicable to all configurations. See the information on the CustomLog directive or the TransferLog directive for Apache HTTP Server.

SSLCachePortFilename

Applicable to all non-Windows configurations with SSL enabled. See the information on the SSLCachePortFilename directive.

SSLCachePath

Applicable when all of the conditions below are true:

- Platform is not Windows.
- SSL is enabled.
- SSLCacheDisable directive is not configured.
- bin/apachectl has been modified to specify a different -d flag, or bin/apachectl is launched with an explicit -d flag.
- The directory specified by the -d flag does not contain the file bin/sidd.

See the information on the SSLCachePath directive for Apache HTTP Server. See information on the SSLCachePath directive.

Other optional directives that specify a file path, like logging or tracing.

4. **AIX** **Windows** Ensure that no more than one IHS instance has the fast response cache accelerator (FRCA), or AFPA, enabled.

Note: FRCA/AFPA has been deprecated starting with V7.0 and its use is discouraged. There is no support for Windows Vista, Windows 2008, or any later Windows operating systems.

5. Start or stop the IHS server instance.

- **AIX** **HP-UX** **Linux** **Solaris** Use these commands to start and stop IHS:

```
# cd /install_dir
# bin/apachectl -k start -f conf/this_instance.conf
# bin/apachectl -k stop -f conf/this_instance.conf
```

Alternatively, you can create a copy of apachectl for each instance, and update the commands in each copy to include "-f conf/this_instance.conf".

- **Windows** Use these commands to setup a new instance:

```
cd \install_dir
bin\Apache.exe -f conf/this_instance.conf -k install -n IHS-this_instance
```

Choose one of these commands to start and stop IHS:

- Use this command:

```
net start IHS-this_instance
```

- Use this command:

```
cd \install_dir
bin\Apache.exe -k install -n IHS-this_instance.conf
```

- Find IHS-this_instance in the Services interface for Microsoft Windows.

See the topic on starting and stopping IBM HTTP Server for more information.

Updating IBM HTTP Server

You can use Installation Manager to update IBM HTTP Server to a later version.

Before you begin

Make sure that your Installation Manager preferences are pointing to Web-based or local repositories that contain the appropriate updates for IBM HTTP Server.

About this task

Perform this procedure to use Installation Manager to update IBM HTTP Server.

Procedure

1. Start Installation Manager.
2. Click **Update**.
3. Select the package group to update.
4. Click **Next**.

Note: If you are prompted to authenticate, use the IBM ID and password that you registered with on the program Web site.

5. Select the version to which you want to update under **IBM HTTP Server**.
6. Click **Next**.
7. Accept the terms in the license agreements, and click **Next**.
8. Review the summary information, and click **Update**.
 - If the installation is successful, the program displays a message indicating that installation is successful.
 - If the installation is not successful, click **View Log File** to troubleshoot the problem.
9. Click **Finish**.
10. Click **File > Exit** to close Installation Manager.

Installing IBM HTTP Server silently

You can use Installation Manager to install IBM HTTP Server silently.

Before you begin

Install Installation Manager on each of the systems onto which you want to install the product.

1. Perform one of the following procedures:
 - If you want to use the Installation Manager that is included with this product, perform the following actions:
 - a. Obtain the necessary files from the physical media or the web.

There are three basic options for obtaining and installing Installation Manager and the product.

 - **Access the physical media, and use local installation**

You can access Installation Manager and the product repositories on the product media. You can install Installation Manager on your system and use it to install the product from the product repositories on the media.
 - **Download the files from the Passport Advantage site, and use local installation**

Licensed customers can download Installation Manager as well as the necessary product repositories from the Passport Advantage site. You can then install Installation Manager on your system and use it to install the product from the repositories.

– **Download a file from the Installation Manager website, and use web-based installation**

You can download and unpack a compressed file containing Installation Manager from the IBM Installation Manager website. You can then install Installation Manager on your local system and use it to install the product from the web-based repository located at

<http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80>

- b. Change to the location containing the Installation Manager installation files, and run one of the following commands to install Installation Manager silently:

Administrative installation:

- **Windows** `installc.exe -acceptLicense -log log_file_path_and_name`
- **AIX** **HP-UX** **Linux** **Solaris** `./installc -acceptLicense -log log_file_path_and_name`

Non-administrative installation:

- **Windows** `userinstc.exe -acceptLicense -log log_file_path_and_name`
- **AIX** **HP-UX** **Linux** **Solaris** `./userinstc -acceptLicense -log log_file_path_and_name`

Group-mode installation:

- **AIX** **HP-UX** **Linux** **Solaris** `./groupinstc -acceptLicense -dataLocation application_data_location -log log_file_path_and_name`

Notes on group mode:

- Group mode allows multiple users to use single instance of IBM Installation Manager to manage software packages. This does not mean that two people can use the single instance of IBM Installation Manager at the same time.
 - **Windows** Group mode is not available on Windows operating systems.
 - If you do not install Installation Manager using group mode, you will not be able to use group mode to manage any of the products that you install later using this Installation Manager.
 - Make sure that you change the installation location from the default location in the current user's home directory to a location that is accessible by all users in the group.
 - Set up your groups, permissions, and environment variables as described in the Group mode road maps in the IBM Installation Manager Information Center before installing in group mode.
 - For more information on using group mode, read the Group mode road maps in the IBM Installation Manager Information Center.
- If you already have a version of Installation Manager installed on your system and you want to use it to install and maintain the product, obtain the necessary product files from the physical media or the web.

There are three basic options for installing the product.

– **Access the physical media, and use local installation**

You can access the product repositories on the product media. Use your existing Installation Manager to install the product from the product repositories on the media.

- **Download the files from the Passport Advantage site, and use local installation**
Licensed customers can download the necessary product repositories from the Passport Advantage site. You can then use your existing Installation Manager to install the product from the repositories.
 - **Access the live repositories, and use web-based installation**
You can install Installation Manager on your local system and use it to install the product from the web-based repository located at
`http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80`
Whenever possible, you should use the remote web-based repositories so that you are accessing the most up-to-date installation files.
2. Add the product repository to your Installation Manager preferences.
 - a. Start Installation Manager.
 - b. In the top menu, click **File > Preferences**.
 - c. Select **Repositories**.
 - d. Perform the following actions:
 - 1) Click **Add Repository**.
 - 2) Enter the path to the repository.config file in the location containing the repository files.
For example:
 - **Windows** `C:\repositories\product_name\local-repositories`
 - **AIX** **HP-UX** **Linux** **Solaris** `/var/repositories/product_name/local-repositories`
 or
`http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80`
 - 3) Click **OK**.
 - e. Deselect any locations listed in the Repositories window that you will not be using.
 - f. Click **Apply**.
 - g. Click **OK**.
 - h. Click **File > Exit** to close Installation Manager.

About this task

Complete this procedure to install IBM HTTP Server silently.

Procedure

1. **Record a response file to install IBM HTTP Server:** On one of your systems, complete the following actions to record a response file that will install IBM HTTP Server.
 - a. From a command line, change to the eclipse subdirectory in the directory where you installed Installation Manager.
 - b. Start Installation Manager from the command line using the -record option.
For example:
 - **Windows Administrator or non-administrator:**
`IBMIM.exe -skipInstall "C:\temp\imRegistry" -record C:\temp\install_response_file.xml`
 - **AIX HP-UX Linux Solaris Administrator:**
`./IBMIM -skipInstall /var/temp/imRegistry -record /var/temp/install_response_file.xml`
 - **AIX HP-UX Linux Solaris Non-administrator:**
`./IBMIM -skipInstall user_home/var/temp/imRegistry -record user_home/var/temp/install_response_file.xml`

Tip: When you record a new response file, you can specify the `-skipInstall` parameter. Using this parameter has the following benefits:

- No files are actually installed, and this speeds up the recording.
- If you use a temporary data location with the `-skipInstall` parameter, Installation Manager writes the installation registry to the specified data location while recording. When you start Installation Manager again without the `-skipInstall` parameter, you then can use your response file to install against the real installation registry.

The `-skipInstall` operation should not be used on the actual agent data location used by Installation Manager. This is unsupported. Use a clean writable location, and re-use that location for future recording sessions.

For more information, read the IBM Installation Manager Information Center.

c. Add the appropriate repositories to your Installation Manager preferences.

1) In the top menu, click **File > Preferences**.

2) Select **Repositories**.

3) Complete the following actions for each repository:

a) Click **Add Repository**.

b) Enter the path to the `repository.config` file in the remote Web-based repository or the local directory into which you unpacked the repository files.

For example:

- Remote repositories:

`https://downloads.mycorp.com:8080/WAS_80_repository`

or

`http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80`

- Local repositories:

– **Windows** C:\repositories\ihs\local-repositories

– **AIX** **HP-UX** **Linux** **Solaris** /var/repositories/ihs/local-repositories

c) Click **OK**.

4) Click **Apply**.

5) Click **OK**.

d. Click **Install**.

Note: If you are prompted to authenticate, use the IBM ID and password that you registered with on the program website.

Installation Manager searches its defined repositories for available packages.

e. In the **Install Packages** window, complete the appropriate actions.

1) Select **IBM HTTP Server** and the appropriate version.

Note: If you are installing the trial version of this product, select **IBM HTTP Server Trial**.

If you already have IBM HTTP Server installed on your system, a message displays indicating that IBM HTTP Server is already installed. To create another installation of IBM HTTP Server in another location, click **Continue**.

2) Click **Next**.

f. Accept the terms in the license agreements, and click **Next**.

g. On the Location panel, specify the installation root directory for IBM HTTP Server binaries, which are also referred to as the core product files or system files.

The panel also displays the shared resources directory and disk-space information.

The core product files do not change unless you install maintenance.

Restrictions:

- Deleting the default target location and leaving an installation-directory field empty prevents you from continuing.
- Do not use symbolic links as the destination directory.
Symbolic links are not supported.
- Do not use spaces in the name of the installation directory.
These spaces are not supported.
- Do not use a semicolon in the directory name.
IBM HTTP Server cannot install properly if the target directory includes a semicolon.

Windows A semicolon is the character used to construct the class path on Windows systems.

- **Windows** The maximum path length on the Windows XP, Windows Vista, and Windows 7 operating systems is 60 characters.

h. Click **Next**.

- i. **AIX** **Linux** **Solaris** If you are installing on a 64-bit system, choose between a 32-bit or 64-bit HTTP Server environment and click **Next**.

Notes:

- This option displays only if you are installing on a 64-bit system.
- This does not apply to Solaris x86 64-bit systems.
- You must select one of the two options.
- You cannot modify this installation later and change this selection.

j. Click **Next** to display the Configuration for IBM HTTP Server panel,

- k. **Windows** On the Configuration for IBM HTTP Server panel, specify your Web server configuration.

- Specify a port number on which IBM HTTP Server will communicate. The default port is 80.
- Choose whether to use a Windows service to run IBM HTTP Server.

Note: You have the option to create a Windows service for IBM HTTP Server on this panel. You can configure the services to run as local system account or a user ID that you specify. The user ID requires the following advanced user rights:

- Act as part of the operating system
- Log on as a service

Important: If you do not select **Run IBM HTTP Server as a Windows Service**, this instance of IBM HTTP Server cannot be stopped or started by the WebSphere Application Server administrative console. At any time after installation, you can create a new service by running the following command:

```
ifs_root/bin/httpd.exe -n new_service_name -k install
```

and then updating the web server definition in the administrative console to reflect the new service name.

- Determine if your startup type will be automatic or manual.

- l. **AIX** **HP-UX** **Solaris** On the Configuration for IBM HTTP Server panel, specify your Web server configuration.

Specify a port number on which IBM HTTP Server will communicate. The default port is 80.

m. Click **Next**.

n. Review the summary information, and click **Install**.

- If the installation is successful, the program displays a message indicating that installation is successful.

Note: The program might also display important post-installation instructions as well.

- If the installation is not successful, click **View Log File** to troubleshoot the problem.
- o. Click **Finish**.
- p. Click **File > Exit** to close Installation Manager.
- q. Optional: If you are using an authenticated remote repository, create a keyring file for silent installation.
 - 1) From a command line, change to the eclipse subdirectory in the directory where you installed Installation Manager.
 - 2) Start Installation Manager from the command line using the -record option.

For example:

- **Windows Administrator or non-administrator:**

```
IBMIM.exe -skipInstall "C:\temp\imRegistry"  
-keyring C:\IM\im.keyring  
-record C:\temp\keyring_response_file.xml
```

- **AIX HP-UX Linux Solaris Administrator:**

```
./IBMIM -skipInstall /var/temp/imRegistry  
-keyring /var/IM/im.keyring  
-record /var/temp/keyring_response_file.xml
```

- **AIX HP-UX Linux Solaris Non-administrator:**

```
./IBMIM -skipInstall user_home/var/temp/imRegistry  
-keyring user_home/var/IM/im.keyring  
-record user_home/var/temp/keyring_response_file.xml
```

- 3) When a window opens that requests your credentials for the authenticated remote repository, enter the correct credentials and **save** them.
- 4) Click **File > Exit** to close Installation Manager.

For more information, read the IBM Installation Manager Information Center.

2. Use the response files to install IBM HTTP Server silently:

- a. Optional: **Use the response file to install the keyring silently:** Go to a command line on each of the systems on which you want to install IBM HTTP Server, change to the eclipse/tools subdirectory in the directory where you installed Installation Manager, and install the keyring silently.

For example:

- **Windows Administrator or non-administrator:**

```
imcl.exe -acceptLicense  
-input C:\temp\keyring_response_file.xml  
-log C:\temp\keyring_log.xml
```

- **AIX HP-UX Linux Solaris Administrator:**

```
./imcl -acceptLicense  
-input /var/temp/keyring_response_file.xml  
-log /var/temp/keyring_log.xml
```

- **AIX HP-UX Linux Solaris Non-administrator:**

```
./imcl -acceptLicense  
-input user_home/var/temp/keyring_response_file.xml  
-log user_home/var/temp/keyring_log.xml
```

- b. **Use the response file to install IBM HTTP Server silently:** Go to a command line on each of the systems on which you want to install IBM HTTP Server, change to the eclipse/tools subdirectory in the directory where you installed Installation Manager, and install IBM HTTP Server silently.

For example:

- **Windows Administrator or non-administrator:**

```
imcl.exe -acceptLicense
-input C:\temp\install_response_file.xml
-log C:\temp\install_log.xml
-keyring C:\IM\im.keyring -password password
```

- **AIX** **HP-UX** **Linux** **Solaris** **Administrator:**

```
./imcl -acceptLicense
-input /var/temp/install_response_file.xml
-log /var/temp/install_log.xml
-keyring /var/IM/im.keyring -password password
```

- **AIX** **HP-UX** **Linux** **Solaris** **Non-administrator:**

```
./imcl -acceptLicense
-input user_home/var/temp/install_response_file.xml
-log user_home/var/temp/install_log.xml
-keyring user_home/var/IM/im.keyring -password password
```

Notes:

- The relevant terms and conditions, notices, and other information are provided in the license-agreement files in the `lfiles` or `product_name/lfiles` subdirectory of the installation image or repository for this product.
- The program might write important post-installation instructions to standard output.

Read the IBM Installation Manager Information Center for more information.

Windows Example

The following is an example of a response file for silently installing IBM HTTP Server.

```
<?xml version="1.0" encoding="UTF-8"?>

<!-- ##### Copyright #####
# Licensed Materials - Property of IBM (c) Copyright IBM Corp. 2011.
# All Rights Reserved. US Government Users Restricted Rights-Use, duplication
# or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
##### -->

<!-- ##### Frequently Asked Questions #####
# The latest information about using Installation Manager is
# located in the online Information Center. There you can find
# information about the commands and attributes used in
# silent installation response files.
#
# Installation Manager Information Center can be found at:
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp
#
# Question 1. How do I record a response file using Installation Manager?
# Answer 1. Start Installation Manager from the command line under the
# eclipse subdirectory with the record parameter and it will generate a
# response file containing actions it performed, repositories it used, and
# its preferences settings. Optionally use the -skipInstall parameter if
# you do not want the product to be installed to the machine. Specify a
# new agentDataLocation location value when doing a new installation. Do
# not use an existing agentDataLocation for an installation because it might
# damage the installation data and prevent you from modifying, updating,
# rolling back, or uninstalling the installed packages.
#
# Windows: IBMIM -record <responseFile> -skipInstall <agentDataLocation>
# Linux or UNIX: ./IBMIM -record <responseFile> -skipInstall <agentDataLocation>
#
# For example:
# Windows = IBMIM.exe -record c:\temp\responsefiles\WASv8.install.Win32.xml
# -skipInstall c:\temp\skipInstall\WebSphere_Temp_Registry
# Linux or UNIX = ./IBMIM -record /home/user/responsefiles/WASv8.install.RHEL64.xml
# -skipInstall c:\temp\skipInstall\WebSphere_Temp_Registry
#
# Question 2. How do I run Installation Manager silently using response file?
# Answer 2. Create a silent installation response file and run the following command
# from the eclipse\tools subdirectory in the directory where you installed
# Installation Manager:
#
# Windows = imcl.exe -acceptLicense -showProgress
# input <response_file_path_and_name> -log <log_file_path_and_name>
# Linux, UNIX, IBM i and z/OS = ./imcl -acceptLicense -showProgress
# input <response_file_path_and_name> -log <log_file_path_and_name>
#
# For example:
```

```

# Windows = imcl.exe -acceptLicense -showProgress
#   input c:\temp\responsefile\WASv8.install.Win32.xml
# Linux, UNIX, IBM i and z/OS = ./imcl -acceptLicense -showProgress
#   input /home/user/responsefile/WASv8.install.RHEL64.xml
#
# The -acceptLicense command must be included to indicate acceptance of all
# license agreements of all offerings being installed, updated or modified.
# The -showProgress command shows progress when running in silent mode.
# Additional commands can be displayed by requesting help: IBMIM -help
#
# Question 3. How do I store and pass credentials to repositories that
# require authentication?
# Answer 3. Installation Manager uses a key ring file to store encrypted
# credentials for authenticating with repositories. Follow this two-step
# process for creating and using a key ring file with Installation Manager.
#
# First, create a key ring file with your credentials by starting
# Installation Manager from the command line under eclipse subdirectory
# with the keyring parameter.
# Use the optional password parameter to password protect your file.
#
# Windows = IBMIM.exe -keyring <path and file name> -password <password>
# Linux, UNIX, IBM i and z/OS = ./IBMIM -keyring <path and file name>
#                               -password <password>
#
# Installation Manager will start in graphical mode. Verify that the
# repositories to which you need to authenticate are included in the
# preferences, File / Preferences / Repositories. If they are not
# listed, then click Add Repositories to add the URL or UNC path.
# Installation Manager will prompt for your credentials. If the repository
# is already in the list, then any attempt to access the repository location,
# such as clicking the Test Connections button, will also prompt for your
# credentials. Enter the correct credential and check the Save password
# checkbox. The credentials are saved to the key ring file you specified.
#
# Second, when you start a silent installation, run imcl under eclipse/tools
# subdirectory, and provide Installation Manager with the location of the key
# ring file and the password if the file is protected. For example:
#
# Windows = imcl.exe -acceptLicense -showProgress
#   input <path and file name of response file>
#   -keyring <path and name of key ring file> -password <password>
# Linux, UNIX, IBM i and z/OS = ./imcl -acceptLicense -showProgress
#   input <path and file name of response file>
#   -keyring <path and name of key ring file> -password <password>
#
##### -->

<!-- ##### Agent Input #####
#
# Note that the "acceptLicense" attribute has been deprecated.
# Use "-acceptLicense" command line option to accept license agreements.
#
# The clean and temporary attributes specify the repositories and other
# preferences Installation Manager uses and whether those settings
# should persist after the installation finishes.
#
# Valid values for clean:
#   true = only use the repositories and other preferences that are
#         specified in the response file.
#   false = use the repositories and other preferences that are
#         specified in the response file and Installation Manager.
#
# Valid values for temporary:
#   true = repositories and other preferences specified in the
#         response file do not persist in Installation Manager.
#   false = repositories and other preferences specified in the
#         response file persist in Installation Manager.
#
##### -->

<agent-input clean="true" temporary="true">

<!-- ##### Repositories #####
# Repositories are locations that Installation Manager queries for
# installable packages. Repositories can be local (on the machine
# with Installation Manager) or remote (on a corporate intranet or
# hosted elsewhere on the internet).
#
# If the machine using this response file has access to the internet,
# then include the IBM WebSphere Live Update Repositories in the list
# of repository locations.
#
# If the machine using this response file cannot access the internet,
# then comment out the IBM WebSphere Live Update Repositories and
# specify the URL or UNC path to custom intranet repositories and
# directory paths to local repositories to use.
#
##### -->

```



```

<server>
  <!-- ##### IBM WebSphere Live Update Repositories #####
  # These repositories contain IBM HTTP Server offerings,
  # and updates for those offerings
  #
  # To use the secure repository (https), you must have an IBM ID,
  # which can be obtained by registering at: http://www.ibm.com/account
  # or your Passport Advantage account.
  #
  # And, you must use a key ring file with your response file.
  ##### -->
<repository location="http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80" />
  <!-- <repository location="https://www.ibm.com/software/rational/repositorymanager/repositories/websphere" /> -->

  <!-- ##### Custom Repositories #####
  # Uncomment and update the repository location key below
  # to specify URLs or UNC paths to any intranet repositories
  # and directory paths to local repositories to use.
  ##### -->
  <!-- <repository location='https:\\w3.mycompany.com\repositories\' /> -->
  <!-- <repository location='/home/user/repositories/websphere/' /> -->

  <!-- ##### Local Repositories #####
  # Uncomment and update the following line when using a local
  # repository located on your own machine to install a
  # IBM HTTP Server offering.
  ##### -->
  <!-- <repository location='insert the full directory path inside single quotes' /> -->
</server>

<!-- ##### Install Packages #####
#
# Install Command
#
# Use the install command to inform Installation Manager of the
# installation packages to install.
#
# The modify attribute is optional and can be paired with an install
# command to add features or paired with an uninstall command to
# remove commands. If omitted, the default value is set to false.
#   false = indicates not to modify an existing install by adding
#           or removing features.
#   true  = indicates to modify an existing install by adding or
#           removing features.
#
# The offering ID attribute is required because it specifies the
# offering to be installed. The offering listed must be present in
# at least one of the repositories listed earlier. The example
# command below contains the offering ID for the IBM HTTP Server.
#
# The version attribute is optional. If a version number is provided,
# then the offering will be installed at the version level specified
# as long as it is available in the repositories. If the version
# attribute is not provided, then the default behavior is to install
# the latest version available in the repositories. The version number
# can be found in the repository.xml file in the repositories.
# For example, <offering ... version='8.0.0.20110617_2222'>.
#
# The profile attribute is required and typically is unique to the
# offering. If modifying or updating an existing installation, the
# profile attribute must match the profile ID of the targeted installation
# of IBM HTTP Server.
#
# The features attribute is optional. Offerings always have at least
# one feature; a required core feature which is installed regardless
# of whether it is explicitly specified. If other feature names
# are provided, then only those features will be installed.
# Features must be comma delimited without spaces.
#
# The feature values for IBM HTTP Server include:
#   com.ibm.jre.6_32bit,com.ibm.jre.6_64bit
#
# The installFixes attribute indicates whether fixes available in
# repositories are installed with the product. By default, all
# available fixes will be installed with the offering.
#
# Valid values for installFixes:
#   none = do not install available fixes with the offering.
#   recommended = installs all available recommended fixes with the offering.
#   all = installs all available fixes with the offering.
#
# Interim fixes for offerings also can be installed while they
# are being installed by including the offering ID for the interim
# fix and specifying the profile ID. A commented out example is
# provided in the install command below.
#
# Installation Manager supports installing multiple offerings at once.
# Additional offerings can be included in the install command,
# with each offering requiring its own offering ID, version, profile value,
# and feature values.

```

```

#
# Profile Command
#
# A separate profile command must be included for each offering listed
# in the install command. The profile command informs Installation
# Manager about offering specific properties or configuration values.
#
# The installLocation specifies where the offering will be installed.
# If the response file is used to modify or update an existing
# installation, then ensure the installLocation points to the
# location where the offering was installed previously.
#
# The eclipseLocation data key should use the same directory path to
# IBM HTTP Server as the installationLocation attribute.
#
# Include data keys for product specific profile properties.
#
##### -->
<install modify='false'>
<offering id='com.ibm.websphere.IHS.v80'
  profile='IBM HTTP Server for WebSphere Application Server V8.0'
  features='core.feature.com.ibm.jre.6_32bit' installFixes='none' />
<!-- <offering id='PM12345_WAS80' profile='IBM HTTP Server for WebSphere Application Server V8.0' /> -->
</install>

<profile id='IBM HTTP Server for WebSphere Application Server V8.0'
  installLocation='C:\Program Files\IBM\WebSphere\AppServer'>
<data key='eclipseLocation' value='C:\Program Files\IBM\WebSphere\AppServer' />
<data key='user.import.profile' value='false' />
<data key='user.ihs.http.server.service.name' value='none' />
<data key='user.ihs.httpPort' value='80' />
<data key='user.ihs.installHttpService' value='false' />
<data key='user.ihs.http.admin.service.name' value='none' />
<data key='user.ihs.runSetupAdmin' value='true' />
<data key='user.ihs.adminPort' value='8008' />
<data key='user.ihs.createAdminAuth' value='true' />
<data key='user.ihs.adminAuthUser' value='admin' />
<data key='user.ihs.adminAuthPassword' value='3NYFazstIQztwg6vC6xv3g==' />
<data key='user.ihs.createAdminUserGroup' value='true' />
<data key='user.ihs.setupAdminUser' value='system' />
<data key='user.ihs.setupAdminGroup' value='systemGroup' />
<data key='user.ihs.installAdminService' value='false' />
<data key='cic.selector.nl' value='en' />
</profile>

<!-- ##### Shared Data Location #####
# Uncomment the preference for eclipseCache to set the shared data
# location the first time you use Installation Manager to do an
# installation.
#
# Eclipse cache location can be obtained from the installed.xml file found in
# Linux/Unix: /var/ibm/InstallationManager
# Windows: C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager
# from the following property:
# <property name='cacheLocation' value='C:\Program Files\IBM\IMShared' />
#
# Open the installed.xml file in a text editor because the style sheet
# might hide this value if opened in a web browser.
# For further information on how to edit preferences, refer to the public library at:
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_silent_prefs.html
#
# After the shared data location is set, it cannot be changed
# using a response file or the graphical wizard.
#
# Ensure that the shared data location is a location that can be written
# to by all user accounts that are expected to use Installation Manager.
#
# By default, Installation Manager saves downloaded artifacts to
# the shared data location. This serves two purposes.
#
# First, if the same product is installed a more than once to the machine,
# then the files in the shared data location will be used rather than
# downloading them again.
#
# Second, during the rollback process, the saved artifacts are used.
# Otherwise, if the artifacts are not saved or are removed, then
# Installation Manager must have to access the repositories used to
# install the previous versions.
#
# Valid values for preserveDownloadedArtifacts:
# true = store downloaded artifacts in the shared data location
# false = remove downloaded artifacts from the shared data location
#
##### -->

<!--
<preference name='com.ibm.cic.common.core.preferences.eclipseCache' value='C:\Program Files\IBM\IMShared' />
<preference name='com.ibm.cic.common.core.preferences.preserveDownloadedArtifacts' value='true' />
-->

```

```

<!-- ##### Preferences Settings #####
# Additional preferences for Installation Manager can be specified.
# These preference correspond to those that are located in the graphical
# interface under File / Preferences.
#
# If a preference command is omitted from or commented out of the response
# file, then Installation Manager uses the preference value that was
# previously set or the default value for the preference.
#
# Preference settings might be added or deprecated in new versions of
# Installation Manager. Consult the online Installation Manager
# Information Center for the latest set of preferences and
# descriptions about how to use them.
#
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp
#
##### -->

<!--
<preference name='com.ibm.cic.common.core.preferences.connectTimeout' value='30'/>
<preference name='com.ibm.cic.common.core.preferences.readTimeout' value='45'/>
<preference name='com.ibm.cic.common.core.preferences.downloadAutoRetryCount' value='0'/>
<preference name='offering.service.repositories.areUsed' value='true'/>
<preference name='com.ibm.cic.common.core.preferences.ssl.nonsecureMode' value='false'/>
<preference name='com.ibm.cic.common.core.preferences.http.disablePreemptiveAuthentication' value='false'/>
<preference name='http.ntlm.auth.kind' value='NTLM'/>
<preference name='http.ntlm.auth.enableIntegrated.win32' value='true'/>
<preference name='com.ibm.cic.common.core.preferences.keepFetchedFiles' value='false'/>
<preference name='PassportAdvantageIsEnabled' value='false'/>
<preference name='com.ibm.cic.common.core.preferences.searchForUpdates' value='false'/>
<preference name='com.ibm.cic.agent.ui.displayInternalVersion' value='false'/>
-->

</agent-input>

```

Uninstalling IBM HTTP Server using the GUI

Use the Installation Manager GUI to uninstall IBM HTTP Server.

Procedure

Uninstall IBM HTTP Server.

1. Start Installation Manager.
2. Click **Uninstall**.
3. In the **Uninstall Packages** window, perform the following actions.
 - a. Select **IBM HTTP Server** and the appropriate version.

Note: If you are uninstalling the trial version of this product, select **IBM HTTP Server Trial**.

- b. Click **Next**.
4. Review the summary information.
5. Click **Uninstall**.
 - If the uninstallation is successful, the program displays a message that indicates success.
 - If the uninstallation is not successful, click **View log** to troubleshoot the problem.
6. Click **Finish**.
7. Click **File > Exit** to close Installation Manager.

Uninstalling IBM HTTP Server silently

You can use Installation Manager to uninstall IBM HTTP Server silently.

Before you begin

Optional: Complete or record the installation of Installation Manager and installation of IBM HTTP Server to a temporary installation registry on one of your systems so that you can use this temporary registry to record the uninstallation without using the standard registry where Installation Manager is installed.

Read the following for more information:

- “Installing IBM HTTP Server using the GUI” on page 45
- “Installing IBM HTTP Server silently” on page 55

Procedure

1. **Record a response file to uninstall IBM HTTP Server:** On one of your systems, complete the following actions to record a response file that will uninstall IBM HTTP Server:
 - a. From a command line, change to the `eclipse` subdirectory in the directory where you installed Installation Manager.
 - b. Start Installation Manager from the command line using the `-record` option.

For example:

- **Windows Administrator or non-administrator:**

```
IBMIM.exe -skipInstall "C:\temp\imRegistry" -record C:\temp\uninstall_response_file.xml
```

- **AIX HP-UX Linux Solaris Administrator:**

```
./IBMIM -skipInstall /var/temp/imRegistry -record /var/temp/uninstall_response_file.xml
```

- **AIX HP-UX Linux Solaris Non-administrator:**

```
./IBMIM -skipInstall user_home/var/temp/imRegistry -record user_home/var/temp/uninstall_response_file.xml
```

Tip: If you choose to use the `-skipInstall` parameter with a temporary installation registry created as described in "Before you begin," Installation Manager uses the temporary installation registry while recording the response file. It is important to note that when the `-skipInstall` parameter is specified, no packages are installed or uninstalled. All of the actions that you complete in Installation Manager simply update the installation data that is stored in the specified temporary registry. After the response file is generated, it can be used to uninstall IBM HTTP Server, removing IBM HTTP Server files and updating the standard installation registry.

The `-skipInstall` operation should not be used on the actual agent data location used by Installation Manager. This is unsupported. Use a clean writable location, and re-use that location for future recording sessions.

For more information, read the IBM Installation Manager Information Center.

- c. Click **Uninstall**.
- d. In the **Uninstall Packages** window, complete the following actions.
 - 1) Select **IBM HTTP Server** and the appropriate version.

Note: If you are uninstalling the trial version of this product, select **IBM HTTP Server Trial**.

- 2) Click **Next**.

- e. Review the summary information.
- f. Click **Uninstall**.
 - If the uninstallation is successful, the program displays a message that indicates success.
 - If the uninstallation is not successful, click **View log** to troubleshoot the problem.
- g. Click **Finish**.
- h. Click **File > Exit** to close Installation Manager.

2. **Use the response file to uninstall IBM HTTP Server silently:** From a command line on each of the systems from which you want to uninstall IBM HTTP Server, change to the `eclipse/tools` subdirectory in the directory where you installed Installation Manager and use the response file that you created to silently uninstall IBM HTTP Server.

For example:

- **Windows Administrator or non-administrator:**

```

imcl.exe
-input C:\temp\uninstall_response_file.xml
-log C:\temp\uninstall_log.xml
  • AIX HP-UX Linux Solaris Administrator:
./imcl
-input /var/temp/uninstall_response_file.xml
-log /var/temp/uninstall_log.xml
  • AIX HP-UX Linux Solaris Non-administrator:
./imcl
-input user_home/var/temp/uninstall_response_file.xml
-log user_home/var/temp/uninstall_log.xml
  Go to the IBM Installation Manager Information Center.

```

Windows Example

The following is an example of a response file for silently uninstalling IBM HTTP Server.

```

<?xml version="1.0" encoding="UTF-8"?>

<!-- ##### Copyright #####
# Licensed Materials - Property of IBM (c) Copyright IBM Corp. 2011.
# All Rights Reserved. US Government Users Restricted Rights-Use, duplication
# or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
##### -->

<!-- ##### Frequently Asked Questions #####
# The latest information about using Installation Manager is
# located in the online Information Center. There you can find
# information about the commands and attributes used in
# silent installation response files.
#
# Installation Manager Information Center can be found at:
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp
#
# Question 1. How do I record a response file using Installation Manager?
# Answer 1. Start Installation Manager from the command line under the
# eclipse subdirectory with the record parameter and it will generate a
# response file containing actions it performed, repositories it used, and
# its preferences settings. Optionally use the -skipInstall parameter if
# you do not want the product to be installed to the machine. Specify a
# new agentDataLocation location value when doing a new installation. Do
# not use an existing agentDataLocation for an installation because it might
# damage the installation data and prevent you from modifying, updating,
# rolling back, or uninstalling the installed packages.
#
# Windows: IBMIM -record <responseFile> -skipInstall <agentDataLocation>
# Linux or UNIX: ./IBMIM -record <responseFile> -skipInstall <agentDataLocation>
#
# For example:
# Windows = IBMIM.exe -record c:\temp\responsefiles\WASv8.install.Win32.xml
# -skipInstall c:\temp\skipInstall\WebSphere_Temp_Registry
# Linux or UNIX = ./IBMIM -record /home/user/responsefiles/WASv8.install.RHEL64.xml
# -skipInstall c:\temp\skipInstall\WebSphere_Temp_Registry
#
# Question 2. How do I run Installation Manager silently using response file?
# Answer 2. Create a silent installation response file and run the following command
# from the eclipse\tools subdirectory in the directory where you installed
# Installation Manager:
#
# Windows = imcl.exe -acceptLicense -showProgress
# input <response_file_path_and_name> -log <log_file_path_and_name>
# Linux, UNIX, IBM i and z/OS = ./imcl -acceptLicense -showProgress
# input <response_file_path_and_name> -log <log_file_path_and_name>
#
# For example:
# Windows = imcl.exe -acceptLicense -showProgress
# input c:\temp\responsefile\WASv8.install.Win32.xml
# Linux, UNIX, IBM i and z/OS = ./imcl -acceptLicense -showProgress
# input /home/user/responsefile/WASv8.install.RHEL64.xml
#
# The -acceptLicense command must be included to indicate acceptance of all
# license agreements of all offerings being installed, updated or modified.
# The -showProgress command shows progress when running in silent mode.
# Additional commands can be displayed by requesting help: IBMIM -help
#
##### -->

<!-- ##### Agent Input #####
# The clean and temporary attributes specify the repositories and other
# preferences Installation Manager uses and whether those settings

```

```

# should persist after the uninstall finishes.
#
# Valid values for clean:
#   true = only use the repositories and other preferences that are
#         specified in the response file.
#   false = use the repositories and other preferences that are
#         specified in the response file and Installation Manager.
#
# Valid values for temporary:
#   true = repositories and other preferences specified in the
#         response file do not persist in Installation Manager.
#   false = repositories and other preferences specified in the
#         response file persist in Installation Manager.
#
##### -->

<agent-input clean='true' temporary='true'>

<!-- ##### Repositories #####
# Repositories are locations that Installation Manager queries for
# installable packages. Repositories can be local (on the machine
# with Installation Manager) or remote (on a corporate intranet or
# hosted elsewhere on the internet).
#
# If the machine using this response file has access to the internet,
# then include the IBM WebSphere Live Update Repositories in the list
# of repository locations.
#
# If the machine using this response file cannot access the internet,
# then comment out the IBM WebSphere Live Update Repositories and
# specify the URL or UNC path to custom intranet repositories and
# directory paths to local repositories to use.
#
##### -->

<server>
  <!-- ##### IBM WebSphere Live Update Repositories #####
  # These repositories contain IBM HTTP Server offerings,
  # and updates for those offerings
  #
  # To use the secure repository (https), you must have an IBM ID,
  # which can be obtained by registering at: http://www.ibm.com/account
  # or your Passport Advantage account.
  #
  # And, you must use a key ring file with your response file.
  ##### -->
  <repository location="http://www.ibm.com/software/repositorymanager/com.ibm.websphere.IHS.v80" />
  <!-- <repository location="https://www.ibm.com/software/rational/repositorymanager/repositories/websphere" /> -->

  <!-- ##### Custom Repositories #####
  # Uncomment and update the repository location key below
  # to specify URLs or UNC paths to any intranet repositories
  # and directory paths to local repositories to use.
  ##### -->
  <!-- <repository location="https://w3.mycompany.com/repositories"/> -->
  <!-- <repository location='/home/user/repositories/websphere/'> -->

  <!-- ##### Local Repositories #####
  # Uncomment and update the following line when using a local
  # repository located on your own machine to install a
  # IBM HTTP Server offering.
  ##### -->
  <!-- <repository location='insert the full directory path inside single quotes'> -->
</server>

<!-- ##### Uninstall Packages #####
#
# Uninstall Command
#
# Use the uninstall command to inform Installation Manager of the
# installation packages to uninstall.
#
# The modify attribute is optional and can be paired with an install
# command to add features or paired with an uninstall command to
# remove commands. If omitted, the default value is set to false.
#   false = indicates not to modify an existing install by adding
#         or removing features.
#   true = indicates to modify an existing install by adding or
#         removing features.
#
# The offering ID attribute is required because it specifies the
# offering to be uninstalled. The example command below contains the
# offering ID for IBM HTTP Server.
#
# The version attribute is optional. If a version number is provided,
# then the offering will be uninstalled at the version level specified
# If the version attribute is not provided, then the default behavior is
# to uninstall the latest version. The version number can be found in
# the repository.xml file in the repositories.
# For example, <offering ... version='8.0.0.20110617_2222'>.

```

```

#
# The profile attribute is required and must match the package group
# name for the offering to be uninstalled.
#
# The features attribute is optional. If there is no feature attribute,
# then all features are uninstalled. If features are specified, then
# only those features will be uninstalled.
# Features must be comma delimited without spaces.
#
# The feature values for IBM HTTP Server include:
# com.ibm.jre.6_32bit,com.ibm.jre.6_64bit
#
# Installation Manager supports uninstalling multiple offerings at once.
# Additional offerings can be included in the uninstall command,
# with each offering requiring its own offering ID, version, profile value,
# and feature values.
#
# Profile Command
#
# A separate profile command must be included for each offering listed
# in the install command. The profile command informs Installation
# Manager about offering specific properties or configuration values.
#
# The installLocation specifies where the offering will be installed.
# If the response file is used to modify or update an existing
# installation, then ensure the installLocation points to the
# location where the offering was installed previously.
#
# The eclipseLocation data key should use the same directory path to
# IBM HTTP Server as the installationLocation attribute.
#
# Include data keys for product specific profile properties.
#
##### -->
<uninstall modify='false'>
<offering id='com.ibm.websphere.IHS.v80'
  profile='IBM HTTP Server for WebSphere Application Server V8.0'
  features='core.feature,com.ibm.jre.6_32bit'>
</uninstall>

<profile id='IBM HTTP Server for WebSphere Application Server V8.0'
  installLocation='C:\Program Files\IBM\HTTPServer'>
<data key='eclipseLocation' value='C:\Program Files\IBM\HTTPServer'>
<data key='user.import.profile' value='false'>
<data key='user.ihs.http.server.service.name' value='none'>
<data key='user.ihs.httpPort' value='80'>
<data key='user.ihs.installHttpService' value='false'>
<data key='user.ihs.http.admin.service.name' value='none'>
<data key='user.ihs.runSetupAdmin' value='false'>
<data key='user.ihs.createAdminAuth' value='false'>
<data key='user.ihs.installAdminService' value='false'>
<data key='user.ihs.win.adminServiceLogOnAsLocalSystem' value='false'>
<data key='user.ihs.createAdminUserGroup' value='false'>
<data key='user.ihs.adminPort' value=''>
<data key='cic.selector.nl' value='en'>
</profile>

<!-- ##### Shared Data Location #####
# Uncomment the preference for eclipseCache to set the shared data
# location the first time you use Installation Manager to do an
# installation.
#
# Eclipse cache location can be obtained from the installed.xml file found in
# Linux/Unix: /var/ibm/InstallationManager
# Windows: C:\Documents and Settings\All Users\Application Data\IBM\Installation Manager
# from the following property:
# <property name='cacheLocation' value='C:\Program Files\IBM\IMShared'>
#
# Open the installed.xml file in a text editor because the style sheet
# might hide this value if opened in a web browser.
# For further information on how to edit preferences, refer to the public library at:
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp?topic=/com.ibm.silentinstall12.doc/topics/r_silent_prefs.html
#
# After the shared data location is set, it cannot be changed
# using a response file or the graphical wizard.
#
# Ensure that the shared data location is a location that can be written
# to by all user accounts that are expected to use Installation Manager.
#
# By default, Installation Manager saves downloaded artifacts to
# the shared data location. This serves two purposes.
#
# First, if the same product is installed a more than once to the machine,
# then the files in the shared data location will be used rather than
# downloading them again.
#
# Second, during the rollback process, the saved artifacts are used.
# Otherwise, if the artifacts are not saved or are removed, then
# Installation Manager must have to access the repositories used to
# install the previous versions.

```

```

#
# Valid values for preserveDownloadedArtifacts:
#   true = store downloaded artifacts in the shared data location
#   false = remove downloaded artifacts from the shared data location
#
##### -->

<!--
<preference name='com.ibm.cic.common.core.preferences.eclipseCache' value='C:\Program Files\IBM\IMShared' />
<preference name='com.ibm.cic.common.core.preferences.preserveDownloadedArtifacts' value='true' />
-->

<!-- ##### Preferences Settings #####
# Additional preferences for Installation Manager can be specified.
# These preference correspond to those that are located in the graphical
# interface under File / Preferences.
#
# If a preference command is omitted from or commented out of the response
# file, then Installation Manager uses the preference value that was
# previously set or the default value for the preference.
#
# Preference settings might be added or deprecated in new versions of
# Installation Manager. Consult the online Installation Manager
# Information Center for the latest set of preferences and
# descriptions about how to use them.
#
# http://publib.boulder.ibm.com/infocenter/install/v1r4/index.jsp
#
##### -->

<!--
<preference name='com.ibm.cic.common.core.preferences.connectTimeout' value='30' />
<preference name='com.ibm.cic.common.core.preferences.readTimeout' value='45' />
<preference name='com.ibm.cic.common.core.preferences.downloadAutoRetryCount' value='0' />
<preference name='offering.service.repositories.areUsed' value='true' />
<preference name='com.ibm.cic.common.core.preferences.ssl.nonsecureMode' value='false' />
<preference name='com.ibm.cic.common.core.preferences.http.disablePreemptiveAuthentication' value='false' />
<preference name='http.ntlm.auth.kind' value='NTLM' />
<preference name='http.ntlm.auth.enableIntegrated.win32' value='true' />
<preference name='com.ibm.cic.common.core.preferences.keepFetchedFiles' value='false' />
<preference name='PassportAdvantageIsEnabled' value='false' />
<preference name='com.ibm.cic.common.core.preferences.searchForUpdates' value='false' />
<preference name='com.ibm.cic.agent.ui.displayInternalVersion' value='false' />
-->

</agent-input>

```

Rolling back IBM HTTP Server

You can use Installation Manager to roll back IBM HTTP Server to an earlier version.

Before you begin

Make sure that your Installation Manager preferences are pointing to Web-based or local repositories that contain the appropriate earlier version of IBM HTTP Server.

About this task

Complete this procedure to use Installation Manager to roll back IBM HTTP Server to an earlier version.

Procedure

1. Start Installation Manager.
2. Click **Roll Back**.
3. Select the package group to roll back.
4. Click **Next**.

Note: If you are prompted to authenticate, use the IBM ID and password that you registered with on the program Web site.

5. Select the version to which you want to roll back under **IBM HTTP Server**.
6. Click **Next**.
7. Review the summary information, and click **Roll Back**.

- If the roll back is successful, the program displays a message indicating that the roll back is successful.
 - If the roll back is not successful, click **View Log File** to troubleshoot the problem.
8. Click **Finish**.
 9. Click **File > Exit** to close Installation Manager.

Migrating and installing IBM HTTP Server on z/OS systems

You can install the IBM HTTP Server product on z/OS systems. Ensure that you order the current version of WebSphere Application Server that contains the IBM HTTP Server for z/OS.

Before you can successfully install IBM HTTP Server, ensure that your environment meets the prerequisites for the application server. For more information, see the Preparing the base operating system topic.

Install the IBM HTTP Server product code using IBM Installation Manager and then run an installation script to create a configuration in the installation directory:

Attention: IBM HTTP Server is different from the HTTP Server for z/OS. The information contained within the IBM HTTP Server product documentation pertains to IBM HTTP Server, not the HTTP Server for z/OS.

- The product code for the IBM HTTP Server is installed into a read-only file system by IBM Installation Manager.
- After you install the product code, choose an installation directory for each IBM HTTP Server instance that you want to run under z/OS. Each server instance requires its own installation directory.
- You can use an install script provided with IBM HTTP Server to: Copy files into this directory, perform initial customization, and create symbolic links to the product code directory.

You can use the Web server jobname or other identifier in the installation directory name. For example:

```
/opt/www/webserver1  
/var/webservers/AAST1
```

In the instruction examples in the following topic for installing IBM HTTP Server, an installation directory of `/etc/websrv1` is used.

Learn more about the installation features for IBM HTTP Server by reading the appropriate topics:

- Creating Installation Manager on a z/OS operating system using downloaded zip files or SMP/E files
- Configuring an instance of IBM HTTP Server
- Uninstalling IBM HTTP Server

Migrating the `mod_auth_saf` module directives to use the `mod_authnz_saf` module directives.

Use the `mod_authnz_saf` directive for your SAF configuration instead of the `mod_auth_saf` directive. In addition, SAF password authentication is now enabled by specifying the SAF basic authentication provider directive.

```
AuthBasicProvider saf
```

Also, the `SAFRequire` and `AuthSAF` directives are not supported in this release of IBM HTTP Server. For information about SAF directives, see the information center topic about SAF directives.

z/OS

Migrating from mod_auth_ldap to mod_authnz_ldap

The mod_auth_ldap directive in IBM HTTP Server Version 6.1 has been renamed to mod_authnz_ldap, and has undergone several changes, such as:

- AuthBasicProvider ldap is now mandatory to have LDAP handle authentication.
- AuthZLDAPAuthoritative replaces AuthLDAPAuthoritative for configuration, if other types of authorization are permissible.
- Require directives have been updated to be uniquely identified as being LDAP related, for example, Require ldap-user and Require ldap-group.

For further information, see the mod_authnz_ldap documentation.

Installing IBM HTTP Server for WebSphere Application Server for z/OS

You install IBM HTTP Server for WebSphere Application Server for z/OS using IBM Installation Manager. Installation Manager installs products from one or more product repositories.

About this task

Note: The IBM Installation Manager is now used on all platforms to install IBM HTTP Server for WebSphere Application Server.

For more information on using Installation Manager, read the IBM Installation Manager Information Center.

To install IBM HTTP Server for z/OS Version 8, you need both of the following:

- Installation Manager running on a z/OS system
To create an Installation Manager on z/OS, perform one of the following procedures:
 - Install the Installation Manager install kit (FMID HGIN140) using SMP/E, and run batch jobs to create the Installation Manager.
 - Download the Installation Manager installation kit as a zip file, unzip it to your z/OS system, and invoke shell commands to create the Installation Manager.
- Access to a copy of the product repository
Install the for z/OS V8 repository (FMID HBBO800) using SMP/E.

Procedure

1. Create an Installation Manager and obtain the product repositories.
Perform the following procedures:
 - “Obtaining an Installation Manager installation kit for installing IBM HTTP Server for z/OS” on page 73
 - “Creating an Installation Manager on z/OS for installing IBM HTTP Server” on page 73
 - “Obtaining product repositories for installing IBM HTTP Server for z/OS” on page 76
2. Install IBM HTTP Server for z/OS Version 8.
Perform the following procedure: “Installing IBM HTTP Server for WebSphere Application Server for z/OS” on page 76.

What to do next

You can use Installation Manager to install the DMZ Secure Proxy Server for IBM for z/OS, the Web Server Plug-ins for on z/OS, and the IBM HTTP Server for z/OS.

Obtaining an Installation Manager installation kit for installing IBM HTTP Server for z/OS

The installation kit for the IBM Installation Manager is provided with the WebSphere Application Server for z/OS Version 8 product as FMID HGIN140. You can also download the IBM Installation Manager installation kit to your z/OS system.

Procedure

- To install the IBM Installation Manager installation kit with SMP/E:
 - If you ordered the WebSphere Application Server for z/OS Version 8 product as part of a ServerPac or SystemPac, the IBM Installation Manager installation kit will already be installed.
Mount the installation kit file system at `/usr/lpp/InstallationManager/V1R4` or a location of your choice.
 - If you ordered the WebSphere Application Server for z/OS Version 8 product as part of a Custom-Built Product Delivery Offering (CBPDO), the IBM Installation Manager installation kit will be included in the CBPDO as FMID HGIN140. Install this product following the instructions in the Installation Manager program directory.
Mount the installation kit file system at `/usr/lpp/InstallationManager/V1R4` or a location of your choice.
- To install the IBM Installation Manager installation kit from a downloaded compressed file:
See the IBM Installation Manager Information Center for download and extraction instructions.
Mount the resulting installation kit file system at `/usr/lpp/InstallationManager/V1R4` or a location of your choice.
The installation kit file system can be mounted read-only after it is installed with SMP/E or downloaded and extracted. It is not modified during Installation Manager processing.

What to do next

When the Installation Manager installation kit is available on your z/OS system, you can use it to create an Installation Manager. See “Creating an Installation Manager on z/OS for installing IBM HTTP Server.”

Creating an Installation Manager on z/OS for installing IBM HTTP Server

You can create one or more Installation Managers on your z/OS system to install and maintain software products.

Before you begin

Install the fix for z/OS APAR OA34228 on each z/OS system that will run IBM Installation Manager to allow the copying of files with extended attributes.

Decide in which of the following modes you want to run the Installation Manager:

admin mode

In admin mode, the Installation Manager is installed from a superuser ID (`uid=0`) and can be invoked from any superuser ID. There can only be one admin-mode Installation Manager on a system.

user mode

In user mode (also called "nonAdmin mode"), the Installation Manager can be invoked only by the user that installed it. There can only be one user-mode Installation Manager for a user.

group mode

In group mode, the Installation Manager can be invoked by any user ID that is connected to the "owning group" for the Installation Manager (the default group of the user ID that creates it). There is no limit to the number of group-mode Installation Managers that you can have on a system.

The Installation Manager will consist of two sets of files—a set of executable files that are copied or updated from the installation kit, and a set of runtime data files that describe the products installed by this Installation Manager. Both sets of files must be writeable by the Installation Manager. You must select locations for both the executable and runtime data for each Installation Manager.

Table 3. Default locations for Installation Manager files. The following table shows the default locations for the Installation Manager executable files ("binaries") and runtime data on z/OS.

Files	Admin or group mode	User mode
Binaries	/InstallationManager/bin	\$HOME/InstallationManager/bin
Runtime data (also called "agent data")	/InstallationManager/appdata	\$HOME/InstallationManager/appdata

These locations are assumed in the Installation Manager documentation and sample jobs. If these names are not appropriate for your system or if you choose to have several Installation Managers, you can choose different names and specify them when you create the Installation Manager.

Procedure

1. Make sure that the fix for z/OS APAR OA34228 is installed on your z/OS system.
2. Create a user ID and group to own the Installation Manager.

This user ID must have the following attributes:

- Read/write home directory
- Read access to FACILITY profile BPX.FILEATTR.APF
- Read access to FACILITY profile BPX.FILEATTR.PROGCTL
- Read access to FACILITY profile BPX.FILEATTR.SHARELIB
- Read access to UNIXPRIV profile SUPERUSER.FILESYS.CHOWN
- Read access to UNIXPRIV profile SUPERUSER.FILESYS.CHANGEPERMS

The user ID that creates the Installation Manager will become the initial (possibly only) user ID that can invoke that particular Installation Manager. If you create an Installation Manager in group mode, the default group for this user will become the "owning group" for the Installation Manager.

You can use an existing user ID if it meets these requirements.

If you installed the Installation Manager installation kit with SMP/E, you can use the Installation Manager sample job GIN2ADMN in SGINJCL to create this user ID and group as well as to assign appropriate permissions.

Tip: If you are creating a group-mode Installation Manager, consider putting the following line in a `.profile` script in the home directory for each user ID that will invoke the Installation Manager:

`umask 002`

This will ensure that all files created in the Installation Manager runtime data and installed products are group writable. Otherwise, you might have to issue `chown 775` commands against these directories whenever you use a different user ID to invoke the group-mode Installation Manager.

3. If the Installation Manager binaries and runtime data will not reside in existing read/write file systems, create file systems for the data and mount the file systems read/write.

The file systems should be owned by the user ID and group that will create the Installation Manager and have permissions 755 for an admin or user-mode Installation Manager or 775 for a group-mode Installation Manager.

If you installed the Installation Manager installation kit with SMP/E, you can use the Installation Manager sample job GIN2CFS in SGINJCL to allocate and mount a file system to hold the binaries and runtime data.

The Installation Manager creation process described below will create the binaries and runtime data directories if they do not already exist.

4. Log in to the Unix system services shell under the owning user ID for the Installation Manager, and change the directory to the location of the Installation Manager installation kit.

```
cd /usr/lpp/InstallationManager/V1R4
```

5. Run the `installc`, `userinstc`, or `groupinstc` command from the installation kit to create the Installation Manager.

- To create an Installation Manager in admin mode, issue the following command from the shell:

```
installc -acceptLicense -installationDirectory binaries_location -dataLocation appdata_location
```

- To create an Installation Manager in user mode, issue the following command from the shell:

```
userinstc -acceptLicense -installationDirectory binaries_location -dataLocation appdata_location
```

- To create an Installation Manager in group mode, issue the following command from the shell:

```
groupinstc -acceptLicense -installationDirectory binaries_location -dataLocation appdata_location
```

You can omit the `-installationDirectory` and `-dataLocation` parameters if you use the default locations.

If you used SMP/E to install the Installation Manager installation kit, you can use sample job GIN2INST in SGINJCL to create an Installation Manager.

What to do next

You can verify that the Installation Manager is correctly installed by logging in to the Unix System Services shell under the user ID that created the Installation Manager and running the Installation Manager `imcl` command from the `eclipse/tools` subdirectory of the Installation Manager's binaries location. For example:

```
cd /InstallationManager/bin/eclipse/tools
```

```
imcl -version
```

You are now ready to install products using IBM Installation Manager.

Authorizing additional users to a group-mode Installation Manager: To allow additional users to access a group-mode Installation Manager, make sure that they meet the requirements listed in the first step of the procedure described above and then connect them to the owning group for the Installation Manager using the TSO `CONNECT` command:

```
CONNECT user2 GROUP(IMGROUP)
```

To create an additional Installation Manager, follow the steps in the procedure described above, selecting a new user ID and group (if appropriate) and new binaries and runtime data locations. Do not share binaries or runtime data locations between separate Installation Managers.

Correcting file ownership or permission problems: If you accidentally invoke an Installation Manager from the wrong user ID, some files might end up with ownerships that prevent normal use of the Installation Manager. To correct this problem, log on to a super user or other privileged user ID and reset the file ownership and permissions for the Installation Manager binaries and runtime data. For example:

```
chown IMADMIN:IMGROUP /InstallationManager/bin
chmod 775 /InstallationManager/bin
```

```
chown IMADMIN:IMGROUP /InstallationManager/appdata
chmod 775 /InstallationManager/appdata
```

If the users of a group-mode Installation Manager do not have `umask` set to allow group-write permission on created files, you might also have to perform this step when switching from one user ID to another. You might also need to set permissions and owners for the product files that you install with the Installation Manager to ensure that maintenance can be performed from other user IDs in the group.

Upgrading the Installation Manager: To upgrade an Installation Manager to a new level of the Installation Manager product, download or install the new level of the IBM Installation Manager installation kit and mount it on your system. Then, change directory to the new level of the installation kit and reissue the same `installc`, `userinstc`, or `groupinstc` command that you used to create the Installation Manager. This will update the Installation Manager's binaries from the new installation kit.

Obtaining product repositories for installing IBM HTTP Server for z/OS

WebSphere Application Server Version 8 products are distributed as IBM Installation Manager repositories. These repositories contain the metadata and files that are required to create one or more levels of a particular product. The product repository for IBM HTTP Server Version 8 is part of the product repository for WebSphere Application Server for z/OS Version 8.

About this task

The initial repository for the WebSphere Application Server for z/OS Version 8 product is installed using SMP/E. These SMP/E-installed product repositories are updated through normal service to include additional WebSphere Application Server for z/OS Version 8 fix packs and feature packs as they become available. Fix packs, feature packs, and ifixes can also be downloaded directly by IBM Installation Manager from a central IBM service website as needed.

Perform this procedure to obtain the repository for WebSphere Application Server for z/OS Version 8.

Procedure

- If you ordered the WebSphere Application Server for z/OS Version 8 product as part of a ServerPac or SystemPac[®], the WebSphere Application Server Version 8 product repository will already be installed. Mount the repository at `/usr/lpp/InstallationManagerRepository/HBB0800` or a location of your choice.
- If you ordered the WebSphere Application Server for z/OS Version 8 product as part of a Custom-Built Product Delivery Offering (CBPDO), the IBM Installation Manager installation kit will be included in the CBPDO as FMID HBB0800. Install this product following the instructions in the WebSphere Application Server for z/OS Version 8 program directory. Mount the repository at `/usr/lpp/InstallationManagerRepository/HBB0800` or a location of your choice.

Results

This repository contains the necessary files to install the product code for WebSphere Application Server for z/OS, including the DMZ Secure Proxy Server for IBM WebSphere Application Server for z/OS, IBM HTTP Server for z/OS, and Web Server Plug-ins for WebSphere Application Server for z/OS.

What to do next

When you use IBM Installation Manager to install WebSphere Application Server Version 8 products, specify the path to the appropriate repository in the `-repositories` parameter of the `imcl` command.

Installing IBM HTTP Server for WebSphere Application Server for z/OS

The code for IBM HTTP Server for WebSphere Application Server for z/OS Version 8 is installed using IBM Installation Manager.

Before you begin

Create an Installation Manager on your z/OS system. You will need to know the location of the binaries directory for the Installation Manager and have access to a user ID that can invoke the Installation Manager.

Obtain the repository for WebSphere Application Server for z/OS Version 8. The following instructions assume that the repository is mounted at `/usr/lpp/InstallationManagerRepository/HBB0800`. The repository can be mounted read-only.

Procedure

1. Choose an installation location for this copy of IBM HTTP Server for WebSphere Application Server for z/OS Version 8.

This copy of IBM HTTP Server must be mounted at this location every time Installation Manager accesses it to install, uninstall, or modify it. This does not have to be the same location at which the product will be mounted when used in production.

Installation Manager requires that every installed product or group of products have its own installation location. Do not install IBM HTTP Server for WebSphere Application Server for z/OS into a location used by any other product.

2. Mount an empty file system read/write at this location.

It will require a minimum of 1,000 tracks (3390) or 50 megabytes. Set the ownership for the file system to that of the Installation Manager user ID, and set the permissions to allow group-write if it will be access by a group-mode Installation Manager. For example:

```
chown IMADMIN:IMGROUP /usr/lpp/IHSA/V8R0
```

```
chmod 775 /usr/lpp/IHSA/V8R0
```

You can use the `zCreateFileSystem.sh` script in the `eclipse/tools` subdirectory of the Installation Manager binaries location to create this file system. For example:

```
cd /InstallationManager/bin/eclipse/tools
```

```
zCreateFileSystem.sh -name WAS.V80.SHAPHFS -type ZFS  
-megabytes 50 10 -volume PRV005  
-mountpoint /usr/lpp/IHSA/V8R0  
-owner IMADMIN -group IMGROUP
```

You can use sample job BBO4CFS in the SBBOJCL dataset to allocate and mount this file system.

3. Log in to the Unix System Services shell under the Installation Manager user ID, and change the directory to the `eclipse/tools` subdirectory of the Installation Manager binaries location.

For example:

```
cd /InstallationManager/bin/eclipse/tools
```

4. Verify that the repository is available.

You do this by issuing the following Installation Manager command-line command.

```
imcl listAvailablePackages -repositories path_to_repository
```

For example:

```
imcl listAvailablePackages -repositories /usr/lpp/InstallationManagerRepositories/HBB0800
```

You should see one or more levels of the IBM HTTP Server for WebSphere Application Server for z/OS Version 8 offering, `com.ibm.websphere.IHS.zOS.v80`.

5. Read the product license, which can be found in the `1afiles` subdirectory of the product repository.
6. Run the Installation Manager command-line tool to install IBM HTTP Server for WebSphere Application Server for z/OS.

```
imcl install com.ibm.websphere.IHS.zOS.v80  
-installationDirectory installation_location  
-repositories repository_location  
-sharedResourcesDirectory shared_data_location  
-acceptLicense
```

The `-sharedResourcesDirectory` parameter points to a directory in which Installation Manager will store artifacts from the repository during installation processing. This value is set the first time a product is installed with a particular Installation Manager. This directory should have at least 30,000 tracks of free space. You can omit this parameter after the shared resources directory has been set.

By specifying `-acceptLicense`, you accept the terms of the product license contained in the product repository.

You can use sample job BBO4INST in the SBBOJCL dataset to perform the product installation.

7. Installation is complete when the Installation Manager completes without error messages.
Logs for the installation can be found in the `logs` subdirectory of the Installation Manager runtime data location.
8. When installation is complete, unmount the file system and remount it read-only for use by WebSphere Application Server nodes and servers.

Configuring an instance of IBM HTTP Server on the z/OS system

You can configure an instance of IBM HTTP Server on the z/OS operating system after installing IBM HTTP Server code using IBM Installation Manager.

Before you begin

Prior to using the installer program:

- Ensure that your environment meets the prerequisites for the application server. For more information, see the Preparing the base operating system topic.
- Install the IBM HTTP Server product code using IBM Installation Manager.
- Mount the file system containing this directory on the z/OS system where the IBM HTTP Server instance will run.
- Perform the z/OS system configurations that are required for IBM HTTP Server.
- If you are installing the product for the first time, then create a System Authorization Facility (SAF) user ID and group for IBM HTTP Server. For information, see the topic about required z/OS system configurations.

The examples that follow in this topic assume a server user ID of `WWWSERV` and a server group of `WWWGROUP`.

- Create an installation directory for the configuration files for the server instance. For more information, see the topic about migrating and installing IBM HTTP Server on z/OS systems.

The examples that follow in this topic assume an installation directory of `/etc/websrv1`. Set the directory permissions to `770` and the directory ownership to the server user ID and group:

```
mkdir /etc/websrv1
chown WWWSERV:WWWGROUP /etc/websrv1
chmod 770 /etc/websrv1
```

- If you are installing the product for the first time, then enable the administrative console to modify the `httpd.conf` file by adding the WebSphere Application Server control region user ID to the IBM HTTP Server group using SAF. For example, to add a user `ASCR1` to the group `WWWGROUP`, type the following command:

```
CONNECT ASCR1 GROUP (WWWGROUP) OWNER (WWWGROUP)
```

About this task

Using the installer program, perform the following tasks to install a running instance of IBM HTTP Server for z/OS on your machine.

Procedure

1. Log in to the z/OS UNIX System Services shell with the user ID that runs the installer. (See the *Before you begin* section for this topic.) Change the directory to the IBM HTTP Server product code directory:

```
cd /usr/lpp/IHSA/V8R0/server
```

2. Set the `umask` value to `022` by specifying `umask 022`. To verify that the `umask` value is set to `022`, run the `umask` command.

3. Run the installer program to install the product files into the installation directory, perform initial customization, and create symbolic links from the installation directory to the product directory.

```
bin/install_ihs <-admin> server_installation_directory <server_port>
```

Three parameters can be used to invoke the installer program.

- Optional: The `-admin` keyword, which allows you to use the administrative console to modify the `httpd.conf` file.
- The installation directory for the server instance. This must not be the same as the product directory.
- Optional: The non-SSL port for the web server. The default port is `80`. You can also change the port on the `Listen` directive.

The following examples invoke the installer program from the administrative console. You can invoke the command with or without support for modifying the `httpd.conf` file. For both examples, `/etc/websrv1` is the installation directory, and 80 is the non-SSL port for the Web server.

- This example invokes the command with support for modifying the `httpd.conf` file.

```
bin/install_ihs -admin /etc/websrv1 80
```

- This example invokes the command without support for modifying the `httpd.conf` file.

```
bin/install_ihs /etc/websrv1 80
```

Note: If your product directory path contains symbolic links, point the symbolic links to the following default product directory: `/usr/lpp/IHSA/V8R0/server`. If you do not use the default product directory, you must invoke the installation script using its absolute path, such as `/WebSphere/8.0/SMPE/bin/install_ihs`. If you do not use of the two options, IBM HTTP Server creates physical links, not logical links, when it creates the symbolic links for the installation directory.

4. Optional: This step is optional unless the administrative console is configured to start and stop IBM HTTP Server. You can start the IBM HTTP Server instance from the MVS console by creating a JCL cataloged procedure for the instance. For more information, see the topic about using JCL procedures to start IBM HTTP Server on z/OS. Ensure that the JCL procedure is assigned to the user and group you defined for IBM HTTP Server, as described in the topic about performing required z/OS system configurations.
5. Optional: You can create multiple instances of IBM HTTP Server by running the IBM HTTP Server installer program more than once. However, you must specify a different installation directory each time you run the installer program.

Results

Perform the following steps to confirm that you have successfully installed a running version of the product on your machine:

1. Log in to the OMVS shell using the server user ID. Verify that the server user ID has a non-zero UID value. Change the directory to the server instance's installation directory:

```
cd /etc/websrv1
```

2. Run the following commands to verify the installation of the program: `apachectl -v` and `apachectl configtest`

The following sample output is an example of a successful program installation:

```
# bin/apachectl -v
Server version: IBM_HTTP_Server/8.0.0.0 (Unix)
Server built:   Jan 9 2008 11:20:34
# bin/apachectl configtest
Syntax OK
```

The actual version string and build date varies.

3. Start IBM HTTP Server.

```
bin/apachectl start
```

4. Point a web browser to the IP name or address of your z/OS system, using either the non-SSL port number you specified when running the installer program, or the default port of 80. You should see the IBM HTTP Server default home page.

5. Stop IBM HTTP Server by running the following command:

```
bin/apachectl stop
```

What to do next

- Install and configure the WebSphere Application Server plug-in for IBM HTTP Server.
- For information about editing the IBM HTTP Server configuration file, `httpd.conf`, and information about supported Apache modules, see the topic about configuring IBM HTTP Server.

Typical changes that you can make to the configuration file are:

- Edit the DocumentRoot directive to point to the Web pages for your site.
- Enable the WebSphere Application Server plug-in for IBM HTTP Server by adding the following directives to the end of httpd.conf:

```
LoadModule was_ap22_module <plugin_config_hfs>/bin/mod_was_ap22_http.so
WebSpherePluginConfig /path/to/existing/plugin-cfg.xml
```

If the plug-in configuration file has been used with a WebSphere Application Server Version 5.0 or 5.1 plug-in, then the file is in EBCDIC. Before using the file with this WebSphere Application Server Version 6.0 or higher plug-in, you need to convert it to ASCII. The following example is for converting the plug-in configuration file from EBCDIC to ASCII:

```
$ iconv -f IBM1047 -t ISO8859-1 < /path/to/existing/plugin-cfg.xml \
> /path/to/ascii/plugin-cfg.xml
```

- Enable SSL support by adding the following directives to the end of httpd.conf:

```
LoadModule ibm_ssl_module modules/mod_ibm_ssl.so
Listen 443
<VirtualHost *:443>
SSLEnable
</VirtualHost>
SSLDisable
Keyfile /saf saf-keyring-name
```

The Keyfile directive can instead specify an HFS file name using the syntax: Keyfile /path/to/keyfile.kdb. The .sth file must be in the same directory as the .kdb file. For more information, see “Securing with SSL communications” on page 90 and “SSL directives” on page 103.

- Enable mod_status by removing the comment delimiters in the default configuration file which are highlighted below:

```
<IfModule mod_status.c>
ExtendedStatus On
</IfModule>
...
#<Location /server-status>
#   SetHandler server-status
#   Order deny,allow
#   Deny from all
#   Allow from .example.com
#</Location>
```

If you want to restrict access to specific networks, uncomment the sample mod_access configuration, but modify the Allow from directive to specify the proper domain or network.

- You can install the Web server to an HFS shared R/W by multiple hosts in a sysplex. There are special configuration requirements for components of the Web server which utilize AF_UNIX sockets. AF_UNIX sockets are not supported by an HFS which are shared R/W, so configuration directives are used to place the AF_UNIX sockets on a filesystem owned by the host on which the Web server runs.
 - If mod_ibm_ssl is loaded, use the SSLCachePortFilename directive to specify a file on a filesystem owned by the local host.
 - If mod_fastcgi is loaded, use the FastCGIipcDir directive to specify a directory on a filesystem owned by the local host.
- Add support for the administrative console after the initial installation.
 - Run the bin/enable_admin script to set the permissions needed to modify the httpd.conf file from the administrative console.
 - To modify the httpd.conf file from the administrative console, you must add the control region user ID to the IBM HTTP Server group using SAF. For example, to add a user *ASCR1* to the group *WWWGROUP*, type the following command:

```
CONNECT ASCR1 GROUP (WWWGROUP) OWNER (WWWGROUP)
```

- To use the administrative console to start and stop IBM HTTP Server, you must create a cataloged JCL procedure. For information, see the topic about using JCL procedures to start IBM HTTP Server on z/OS. Ensure that the JCL procedure is assigned to the user and group you defined for IBM HTTP Server, as described in the topic about performing required z/OS system configurations.

Running multiple instances of IBM HTTP Server from a single install

Run multiple, independent instances of IBM HTTP Server from a single installation. It is seldom necessary to run multiple instances, as features like virtual hosts allow a single instance to efficiently serve many sites, but in some cases it is necessary. If you need to securely administer your sites by different administrators, for example, you must run separate instances that each use their own configuration files.

Before you begin

This topic is primarily for AIX, HP-UX, Linux, Solaris, and Windows operating systems. On the z/OS platform, the `install_ihs` command creates a separate directory for each instance without creating another copy of the product. See the z/OS topic for configuring IBM HTTP Server for more information.

Before configuring multiple instances, consider if your problem can be solved by using virtual hosts and/or having IBM HTTP Server listen on multiple addresses and ports. The advantage of a single instance is that it uses less resources to serve the same requests as multiple instances.

Note: When you follow the examples, change "this_instance" to a unique name for each instance.

Procedure

1. Create a separate main configuration file, normally the `httpd.conf` file, for each instance.

Note: To reduce duplication, store common directives in common files and import these into the separate, main configuration files with the *Include* directive.

We'll call the configuration file `conf/this_instance.conf` for the rest of these steps.

Here is a simple example of a configuration file for an instance:

```
Listen 10.0.0.1:80
PidFile instance1/httpd.pid
ErrorLog instance1/error.log
CustomLog instance1/access.log common
# Other directives that make this instance behave uniquely
Include conf/common.conf
```

A real configuration file would have more directives in it to make this instance behave differently than the other instances.

2. Configure the port settings in the configuration files. You cannot use a combination of listen port and listen IP address for more than one instance. Check the Listen directives in each configuration file, and verify that they are unique. See information on the Listen directive for Apache HTTP Server for more information.
3. Configure settings for logging and other special files. Any files that are normally stored in the `install_root/logs` directory cannot be shared between instances. Each instance must have unique values for the following directives:

PidFile

Applicable to all configurations. See the information on the PidFile directive for Apache HTTP Server.

ScriptSock

Applicable to non-Windows configurations with `mod_cgid` enabled.

ErrorLog

Applicable to all configurations. See the information on the ErrorLog directive for Apache HTTP Server.

CustomLog or TransferLog

Applicable to all configurations. See the information on the CustomLog directive or the TransferLog directive for Apache HTTP Server.

SSLCachePortFilename

Applicable to all non-Windows configurations with SSL enabled. See the information on the SSLCachePortFilename directive.

SSLCachePath

Applicable when all of the conditions below are true:

- Platform is not Windows.
- SSL is enabled.
- SSLCacheDisable directive is not configured.
- bin/apachectl has been modified to specify a different -d flag, or bin/apachectl is launched with an explicit -d flag.
- The directory specified by the -d flag does not contain the file bin/sidd.

See the information on the SSLCachePath directive for Apache HTTP Server. See information on the SSLCachePath directive.

Other optional directives that specify a file path, like logging or tracing.

4. **AIX** **Windows** Ensure that no more than one IHS instance has the fast response cache accelerator (FRCA), or AFPA, enabled.

Note: FRCA/AFPA has been deprecated starting with V7.0 and its use is discouraged. There is no support for Windows Vista, Windows 2008, or any later Windows operating systems.

5. Start or stop the IHS server instance.

- **AIX** **HP-UX** **Linux** **Solaris** Use these commands to start and stop IHS:

```
# cd /install_dir
# bin/apachectl -k start -f conf/this_instance.conf
# bin/apachectl -k stop -f conf/this_instance.conf
```

Alternatively, you can create a copy of apachectl for each instance, and update the commands in each copy to include "-f conf/this_instance.conf".

- **Windows** Use these commands to setup a new instance:

```
cd \install_dir
bin\Apache.exe -f conf/this_instance.conf -k install -n IHS-this_instance
```

Choose one of these commands to start and stop IHS:

- Use this command:

```
net start IHS-this_instance
```

- Use this command:

```
cd \install_dir
bin\Apache.exe -k install -n IHS-this_instance.conf
```

- Find IHS-this_instance in the Services interface for Microsoft Windows.

See the topic on starting and stopping IBM HTTP Server for more information.

Uninstalling IBM HTTP Server for z/OS

Use IBM Installation Manager to uninstall the product code for IBM HTTP Server for z/OS.

Before you begin

Make sure that you no longer need IBM HTTP Server for z/OS.

Procedure

- Mount the file system containing the product code to be uninstalled at the installation location that Installation Manager used to install it.
- Log in to the Unix System Services shell under the Installation Manager user ID, and change the directory to the `eclipse/tools` subdirectory of the Installation Manager binaries location.

For example:

```
cd /InstallationManager/bin/eclipse/tools
```

- Invoke the Installation Manager `uninstall` command to perform the uninstallation.

```
imcl uninstall package_ID  
-installationDirectory installation_location
```

For example:

```
imcl uninstall com.ibm.websphere.IHS.zOS.v80  
-installationDirectory /usr/lpp/zWebSphere/V8R0
```

Uninstallation is complete when the Installation Manager completes without error messages. Logs for the uninstallation can be found in the `logs` subdirectory of the Installation Manager runtime data location.

- When uninstallation is complete, delete any remaining files from the product location.

Chapter 4. Product overview and quick start

This section provides shortcuts to information for obtaining a high level understanding of the product, then getting started quickly.

What is new in this release

IBM HTTP Server contains some new functions. Review this topic to find out about what is new in this release.

Distributed operating systems

IBM Installation Manager

IBM® Installation Manager is a common installer for many IBM software products that you use to install, update, roll back, and uninstall IBM HTTP Server.

Distributed operating systems

Bundled Global Security Toolkit

The bundled Global Security Toolkit (GSKit) has been upgraded for Version 8.

- Transport Layer Security (TLS) Version 1.1 and Version 1.2 connections are supported.
- New cipher suites that use SHA-2 have been added.
- The gsk7cmd and gsk7capiCmd command-wrappers have been consolidated into a single command-line key management utility, gskcapiCmd.
- On distributed platforms, the SSLCipherSpec directive has an optional protocol name to control individual applicability to Secure Sockets Layer (SSL) Version 3, TLS Version 1.0, TLS Version 1.1, and TLS Version 1.2. In previous releases, SSL Version 3 and TLS Version 1.0 always shared a common set of configured cipher suites.

Distributed operating systems

IKEYMAN utility

The IKEYMAN utility has been upgraded for Version 8.

- The utility is now a pure Java application with no dependency on the native GSKit run time. Unlike prior releases, do not move or remove the bundled <IHS_root>/java/jre/lib/gskikm.jar library.
- Significant error reporting enhancements have been made within the utility.
- The PKCS11 software for the cryptographic hardware has been updated with significant administrative changes.

▶ AIX ▶ HP-UX ▶ Solaris ▶ Linux ▶ z/OS

64-bit addressing

The ability to use the IBM HTTP Server in 64-bit addressing mode has been added for AIX, LINUX/x86, Linux/PPC, Linux on System z, and Solaris/SPARC platforms. As in prior releases, the IBM HTTP Server is available exclusively in 64-bit addressing mode on the HP-UX/ia64, Solaris/x64, and z/OS platforms.

Distributed operating systems

Web server hardening

Some common web server hardening is now the default behavior:

- The default for the TraceEnable directive is OFF.
- The SSL Version 2 protocol and weak and export SSL ciphers are disabled by default.

Key differences from the Apache HTTP Server

This section takes a high-level look at the main differences between IBM HTTP Server and the Apache HTTP Server.

IBM HTTP Server is based on the open source Apache Web server (httpd.apache.org). The Apache Web server can be built with many different capabilities and configuration options. IBM HTTP Server includes a set of features from the available options. For information about Apache Web server features supported in IBM HTTP Server, see the information center topics about Apache modules (containing directives), programs, Apache Portable Runtime (APR) and APR-util libraries, and Multi-processing module (MPM) and addressing modes.

Key features added with IBM HTTP Server

- Support for the WebSphere administrative console.
- **Distributed operating systems** InstallShield for multiple platforms enables consistent installation of the IBM HTTP Server on different platforms.
- **AIX** **Windows** Fast Response Cache Accelerator (FRCA) is available for AIX 5.x and later and certain Windows operating systems. It significantly improves HTTP Server performance when serving static content such as HTML files or image files.
- Dynamic content generation with FastCGI.
- Installation of IBM HTTP Server in multiple languages on all platforms.

Operational differences between Apache and IBM HTTP Server:

- **AIX** **HP-UX** **Linux** **Solaris** **z/OS** The **apachectl** command is the only supported command to start IBM HTTP Server. You cannot directly invoke the **httpd** command because it will not find the required libraries. The **apachectl** command is the preferred command to start Apache V2.0 and higher, but the **httpd** command might work on the Apache server as expected, depending on the platform and how Apache was built. You can specify **httpd** options on the **apachectl** command line.
- **AIX** **HP-UX** **Linux** **Solaris** IBM HTTP Server supports the **suEXEC** program, which provides for execution of CGI scripts under a particular user ID.
 - If you use the **suEXEC** program, you must install the IBM HTTP Server to the default installation directory only. The **suEXEC** program uses the security model which requires that all configuration paths are hard-coded in the executable file, and the paths chosen for IBM HTTP Server are those of the default installation directory.
 - When an Apache user chooses an installation location for Apache at compile time, the **suEXEC** program is pre-built with the chosen paths, so this issue is seen by the Apache users.
 - Customers who need to use the **suEXEC** program with arbitrary configuration paths can build it with Apache on their platform and use the generated **suEXEC** binary with IBM HTTP Server. Customers must save and restore their custom **suEXEC** file when applying IBM HTTP Server maintenance.
- **z/OS** As a more flexible replacement for the **suEXEC** program (which is for other platforms), IBM HTTP Server supports **SAFRunAs**. The **SAFRunAs** directive provides for execution of CGI scripts and access to static files under a particular user ID. You can enable the **SAFRunAs** directive from the **mod_authnz_saf** load module.
- IBM HTTP Server provides the new **AddServerHeader** directive to allow the server response header to be suppressed. For more information about this directive, see the **AddServerHeader** directive Web page.

Chapter 5. Securing IBM HTTP Server

Learn about IBM HTTP Server security, including: Secure Socket Layer (SSL), Key management, Lightweight Directory Access Protocol (LDAP) and System Authorization Facility (SAF) for z/OS systems.

Securing IBM HTTP Server

This section lists topic overviews for securing IBM HTTP Server.

About this task

The following topics describe specific tasks for you to secure IBM HTTP Server.

Procedure

- “Configure SSL between the IBM HTTP Server Administration Server and the deployment manager”
- “Securing with SSL communications” on page 90. For secure communication, you can set up the Secure Sockets Layer (SSL) directives in the default `httpd.conf` configuration file.
- “Setting advanced SSL options” on page 123. More advanced SSL options to secure your IBM HTTP Server are also available. Advanced SSL options include: setting the level and type of client authentication, setting cipher specifications, defining SSL for multiple-IP virtual hosts, and configuring reverse proxy setup with SSL.
- **Distributed operating systems** “Managing keys with the IKEYMAN graphical interface (Distributed systems)” on page 132. You can set up the Key Management utility (IKEYMAN) with IBM HTTP Server to create key databases, public and private key pairs and certificate requests. Use the IKEYMAN graphical user interface rather than using the command line interface.
- **Distributed operating systems** “Managing keys with the gskcmd command line interface (Distributed systems)” on page 141. You can use IKEYCMD, which is the Java command line interface to IKEYMAN. Use the command line only if you are unable to use the graphical user interface.
- **z/OS** “Managing keys with the native key database gskkyman (z/OS systems)” on page 153 You can use the native z/OS key management (gskkyman key database) with IBM HTTP Server to create key databases, public and private key pairs and certificate requests.
- “Getting started with the cryptographic hardware for SSL (Distributed systems)” on page 154. You can use cryptographic hardware for SSL. The IBM 4758 requires the PKCS11 software for the host machine and internal firmware.
- **Distributed operating systems** “Authenticating with LDAP on IBM HTTP Server using **mod_ibm_idap** (Distributed systems)” on page 159. You can configure LDAP to protect files on IBM HTTP Server.
- **z/OS** “Authenticating with LDAP on IBM HTTP Server using **mod_idap**” on page 185 You can configure LDAP to protect files on IBM HTTP Server.
- **z/OS** “Authenticating with SAF on IBM HTTP Server (z/OS systems)” on page 187. You can provide IBM HTTP Server with user authentication using the System Authorization Facility security product.

Results

Your IBM HTTP Server is secured.

Configure SSL between the IBM HTTP Server Administration Server and the deployment manager

Configure Secure Sockets Layer (SSL) between the deployment manager for WebSphere Application Server and the IBM HTTP Server administration server, which is called `adminctl`.

About this task

The Application Server has new SSL management functions that need to be managed properly in order for IBM HTTP Server to connect with an SSL request. In earlier releases, SSL connections used default dummy certificates that were exchanged between IBM HTTP Server and the Application Server. In WebSphere Application Server, you must configure the Application Server to accept a self-signed certificate from IBM HTTP Server so SSL connections are accepted and transactions are completed.

If the Application Server and the IBM HTTP Server administration server are not configured correctly, the Application Server shows any errors that are received in the log file for the deployment manager. In situations where the IBM HTTP Server administration server is attempting to connect through SSL and the Application Server is not configured, you might receive an error that is similar to the following message:

```
-CWPKI0022E: SSL HANDSHAKE FAILURE: A signer with
SubjectDN "CN=localhost" was sent from target host:port "null:null".
The signer may need to be added to local trust store "c:/619/app2/profiles/Dmgr01/config/cells/rjrCell02/trust.p12"
located in SSL configuration alias "CellDefaultSSLSettings"
loaded from SSL configuration file "security.xml".
The extended error message from the SSL handshake
exception is: "No trusted certificate found".
```

```
-IOException javax.net.ssl.SSLHandshakeException:
com.ibm.jsse2.util.h: No trusted certificate found
```

Procedure

1. Obtain a server certificate. You can generate a new self-signed certificate or use the existing certificate from the IBM HTTP Server Web server plugin.
 - Use the existing self-signed certificate from the IBM HTTP Server Web server plugin.
 - Create a CMS key database file and a self-signed server certificate. Use the iKeyman utility for distributed operating systems and the gskkyman tool for z/OS operating systems. This step and later steps will assume that you are using the iKeyman utility.

- **Distributed operating systems** Use the IBM HTTP Server iKeyman utility graphical user interface or command line to create a CMS key database file and a self-signed server certificate.

Use the iKeyman utility to create a self-signed certificate for the IBM HTTP Server Administration Server and save the certificate as `/conf/admin.kdb`.

Note: Make note of the password and select **Stash password to a file**.

The following fields are required for the certificate:

Label adminselfSigned

Common Name

fully_qualified_host_name

- **z/OS** IBM HTTP Server uses the z/OS gskkyman tool for key management to create a CMS key database file, public and private key pairs, and self-signed certificates. Alternatively, you can create a SAF keyring in place of a CMS key database file.
 - For information on gskkyman, see Key management using the native z/OS key database.
 - For information on creating SAF keyrings, see Authenticating with SAF on IBM HTTP Server and SSL keyfile directive.

2. Extract the certificate to a file using iKeyman utility.
 - a. Select the certificate that you created in Step 1. For example, adminselfSigned.
 - b. Click **Extract Certificate**. The recommended file name for extraction is `C:\Program Files\IBM\HTTPServer\conf\cert.arm`.

Note: Do not change the data type.

3. Modify the Administration Server configuration File, which is named `admin.conf`.

- a. Configure the file to load the IBM SSL module. Uncomment the following line:

```
LoadModule ibm_ssl_module      modules/mod_ibm_ssl.so
```
- b. Enable SSL and define a key file to use. Uncomment the following lines to enable SSL and define a key file to use:

```
SSLEnable
SSLServerCert default
Keyfile "C:/Program Files/IBM/HTTPServer5/conf/admin.kdb"
```

Note: Be aware of the following:

- The key file directive must match the name and location of a valid key file that is installed on your system.
 - You must have IBM SSL support installed for this to work.
 - The "default" in SSLServerCert is the label, or name, of the self-signed certificate that is created when the plugin-key.kdb file was created.
 - The previous example uses SSLServerCert because the default self-signed certificate in the plugin-key.kdb is not flagged as the default certificate.
4. Start the administration server for IBM HTTP Server. Verify that the log file does not contain GSKIT errors.
 5. Configure WebSphere Application Server.
 - a. Log into the Administrative Console for the Application Server and start the deployment manager.
 - b. Select **Security > SSL certificate and key management**.
 - c. Select **Manage endpoint security configurations**. You are directed to a list of inbound and outbound endpoints.
 - d. Select the outbound cell (cellDefaultSSLSettings,null). Select outbound cells because, in this setup, the Administration Console for the Application Server is the client, and the IBM HTTP Server Administration Server is the server.

Note: This setup is the opposite configuration from an SSL setup with the IBM HTTP Server plugin and the Application Server.

- e. In the Related Items section, click **Key stores and certificates**.
- f. Click **CellDefaultTrustStore**.
- g. In the Additional Properties section, click **Signer Certificates**.
- h. FTP the certificate file to the Application Server. Do not change the data type.
- i. In the collection panel for Signer Certificates, click **Add**. Enter the following information in the fields.

Table 4. Signer Certificate information

Name	Value
Alias	adminselfSigned
File name	<i>file_name</i> For example, enter the following: c:\program files\ibm\httpserver\conf\cert.arm

- j. Save the configuration changes to the administrative console.
- k. Stop the deployment manager.
- l. Start the deployment manager.

Results

The IBM HTTP Server administration server and Application Server are now configured to use SSL transactions.

Securing with SSL communications

This section provides information to help you set up Secure Sockets Layer (SSL), using the default `httpd.conf` configuration file.

About this task

For each virtual host, set the cipher specification to use during secure transactions. The specified cipher specifications validate against the level of the Global Security Kit (GSK) toolkit that is installed on your system. Invalid cipher specifications cause an error to log in the error log. If the client issuing the request does not support the ciphers specified, the request fails and the connection closes to the client.

IBM HTTP Server has a built-in list of cipher specifications to use for communicating with clients over Secure Sockets Layer (SSL). The actual cipher specification that is used for a particular client connection is selected from those cipher specifications that both IBM HTTP Server and the client support.

Some cipher specifications provide a weaker level of security than others, and might need to be avoided for security reasons. Some of the stronger cipher specifications are more computationally intensive than weaker cipher specifications and might be avoided if required for performance reasons. You can use the `SSLCipherSpec` directive to provide a customized list of cipher specifications that are supported by the Web server in order to avoid the selection of cipher specifications that are considered too weak or too computationally intensive.

If you do not specify cipher specifications using the `SSLCipherSpec` directive, IBM HTTP Server Version 8.0 and later uses a conservative set of default ciphers. The default set of ciphers excludes SSL Version 2, null ciphers, and weak ciphers. The weak ciphers include export-grade ciphers. These defaults can be viewed at runtime in the error log by enabling `LogLevel` debug and `SSLTrace`.

Procedure

1. **Distributed operating systems** Use the IBM HTTP Server `IKEYMAN` utility (graphical user interface) or `IKEYMAN` utility (command line) to create a CMS key database file and server certificate.
2. **z/OS** IBM HTTP Server uses the `z/OS gskkyman` tool for key management to create a CMS key database file, public and private key pairs, and server certificates. Or, you can create a SAF keyring in place of a CMS key database file.
 - For information on `gskkyman`, see [Key management using the native z/OS key database](#).
 - For information on creating SAF keyrings, see [“Authenticating with SAF on IBM HTTP Server \(z/OS systems\)”](#) on page 187 and `SSL keyfile` directive.
3. Enable SSL directives in the IBM HTTP Server `httpd.conf` configuration file.
 - a. Uncomment the `LoadModule ibm_ssl_module modules/mod_ibm_ssl.so` configuration directive.
 - b. Create an SSL virtual host stanza in the `httpd.conf` file using the following examples and directives.

```
LoadModule ibm_ssl_module modules/mod_ibm_ssl.so
Listen 443
<VirtualHost *:443>
    SSLEnable
</VirtualHost>
SSLDisable
KeyFile "c:/Program Files/IBM HTTP Server/key.kdb"
```

This second example assumes that you are enabling a single Web site to use SSL, and the server name is different from the server name that is defined in the global scope for non-SSL (port 80). Both host names must be registered in a domain name server (DNS) to a separate IP address, and you must configure both IP addresses on local network interface cards.

```
Listen 80
ServerName www.mycompany.com
```

```

<Directory "c:/Program Files/IBM HTTP Server/htdocs">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>

DocumentRoot "c:/program files/ibm http server/htdocs"
DirectoryIndex index.html

<VirtualHost 192.168.1.103:80>
ServerName www.mycompany2.com
<Directory "c:/Program Files/IBM HTTP Server/htdocs2">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs2"
DirectoryIndex index2.html
</VirtualHost>

Listen 443
<VirtualHost 192.168.1.103:443>
ServerName www.mycompany2.com
SSLEnable
SSLClientAuth None
<Directory "c:/Program Files/IBM HTTP Server/htdocs2">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs2"
DirectoryIndex index2.html
</VirtualHost>

SSLDisable
KeyFile "c:/program files/ibm http server/key.kdb"
SSLV2Timeout 100
SSLV3Timeout 1000

```

This third example assumes that you are enabling multiple Web sites to use SSL. All host names must be registered in the domain name server (DNS) to a separate IP address. Also, you must configure all of the IP addresses on a local network interface card. Use the `SSLServerCert` directive to identify which personal server certificate in the key database file passes to the client browser during the SSL handshake for each Web site. If you have not defined the `SSLServerCert` directive, IBM HTTP Server passes the certificate in the key database file that is marked (*) as the "default key".

```

Listen 80
ServerName www.mycompany.com

<Directory "c:/Program Files/IBM HTTP Server/htdocs">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>

DocumentRoot "c:/program files/ibm http server/htdocs"
DirectoryIndex index.html

<VirtualHost 192.168.1.103:80>
ServerName www.mycompany2.com
<Directory "c:/Program Files/IBM HTTP Server/htdocs2">

```

```

Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs2"
DirectoryIndex index2.html
</VirtualHost>

<VirtualHost 192.168.1.104:80>
ServerName www.mycompany3.com
<Directory "c:/Program Files/IBM HTTP Server/htdocs3">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs3"
DirectoryIndex index3.html
</VirtualHost>

Listen 443
<VirtualHost 192.168.1.102:443>
ServerName www.mycompany.com
SSLEnable
SSLClientAuth None
SSLServerCert mycompany
<Directory "c:/Program Files/IBM HTTP Server/htdocs">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs"
DirectoryIndex index.html
</VirtualHost>

<VirtualHost 192.168.1.103:443>
ServerName www.mycompany2.com
SSLEnable
SSLClientAuth None
SSLServerCert mycompany2
<Directory "c:/Program Files/IBM HTTP Server/htdocs2">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs2"
DirectoryIndex index2.html
</VirtualHost>

<VirtualHost 192.168.1.104:443>
ServerName www.mycompany3.com
SSLEnable
SSLClientAuth None
SSLServerCert mycompany3
<Directory "c:/Program Files/IBM HTTP Server/htdocs3">
Options Indexes
AllowOverride None
order allow,deny
allow from all
</Directory>
DocumentRoot "c:/program files/ibm http server/htdocs3"
DirectoryIndex index3.html
</VirtualHost>

```

```
SSLDisable
KeyFile "c:/program files/ibm http server/key.kdb"
SSLV2Timeout 100
SSLV3Timeout 1000
```

Secure Sockets Layer (SSL) protocol

The Secure Sockets Layer (SSL) protocol was developed by Netscape Communications Corporation.

SSL ensures the data that is transferred between a client and a server remains private. This protocol enables the client to authenticate the identity of the server.

When your server has a digital certificate, SSL-enabled browsers can communicate securely with your server, using SSL. With SSL, you can easily establish a security-enabled Web site on the Internet, or on your private intranet. A browser that does not support HTTP over SSL cannot request URLs using HTTPS. The non-SSL browsers do not allow submission of forms that require secure communications.

SSL uses a *security handshake* to initiate a secure connection between the client and the server. During the handshake, the client and server agree on the security keys to use for the session and the algorithms to use for encryption. The client authenticates the server; optionally, the server can request the client certificate. After the handshake, SSL encrypts and decrypts all the information in both the HTTPS request and the server response, including:

- The URL requested by the client
- The contents of any submitted form
- Access authorization information, like user names and passwords
- All data sent between the client and the server

HTTPS represents a unique protocol that combines SSL and HTTP. Specify `https://` as an anchor in HTML documents that link to SSL-protected documents. A client user can also open a URL by specifying `https://` to request an SSL-protected document.

Because HTTPS (HTTP + SSL) and HTTP are different protocols and use different ports (443 and 80, respectively), you can run both SSL and non-SSL requests simultaneously. This capability enables you to provide information to users without security, while providing specific information only to browsers making secure requests. With this functionality, a retail company on the Internet can support users looking through their company merchandise without security, but then fill out order forms and send their credit card numbers using security.

Certificates

This topic provides information on Secure Sockets Layer certificates.

Distributed operating systems Use the IBM HTTP Server IKEYMAN utility to create a CMS key database file and server certificate.

z/OS For IBM HTTP Server, use the native z/OS key management (gskkyman key database) to create a CMS key database file and server certificate.

Production Web servers must use signed certificates purchased from a Certificate Authority that supports IBM HTTP Server such as VeriSign or Thawte. The default certificate request file name is `certreq.arm`. The certificate request file is a PKCS 10 file, in Base64-encoded format.

Distributed operating systems You can use the IKEYMAN Key Management utility or IKEYMAN Key Management utility command line interface that is provided with IBM HTTP Server to create server certificates.

z/OS You can use the native z/OS key management (gskkyman key database) to create server certificates.

Self-signed certificates are useful for test purposes but should not be used in a production Web server.

For your convenience, IBM HTTP Server includes several default signer certificates. Be aware that these default signer certificates have expiration dates. It is important to verify the expiration dates of all your certificates and manage them appropriately. When you purchase a signed certificate from a CA, they will provide you access to their most recent signer certificates.

List of trusted certificate authorities on the IBM HTTP Server:

Associate your public key with a digitally signed certificate from a certificate authority (CA) that is designated as a trusted root CA on your server. You can buy a signed certificate by submitting a certificate request to a certificate authority provider. The default certificate request file name is `certreq.arm`. The certificate request file is a PKCS 10 file, in Base64-encoded format.

You can create a new `.kdb` keystore file and view the list of designated trusted certificate authorities (CAs). If you are using a personal certificate and the signer is not in the list, you must obtain a signer certificate from the associated trusted certificate authority. IBM HTTP Server supports the following certificate authority (CA) software:

- Any X.509-compliant certificate authority
- Entrust
- Netscape Certificate Server
- Tivoli® PKI
- XCert

Certificate expiration dates:

You can display expiration dates of certificates in your key database by viewing the certificate information with the IKEYMAN Key Management utility GUI or using the `gskcmd` command.

The following is an example of how to use the `gskcmd` command to display the validity dates on all certificates in the `key.kdb` certificate key file that will expire within 1825 days (5 years):

```
<ihsinst>/bin/gskcmd -cert -list all -expiry 1825 -db key.kdb -pw <password>
```

```
Certificates in database: key.kdb  
VeriSign Class 1 CA Individual Subscriber-Persona Not Validated  
Validity  
Not Before: Mon May 11 20:00:00 EDT 1998  
Not After: Mon May 12 19:59:59 EDT 2008
```

where `<password>` is the password you specified when creating the `key.kdb` key database file.

SSL certificate revocation list:

This section provides information on identifying directives for certificate revocation list (CRL) and those supported in global servers and virtual hosts.

Certificate revocation provides the ability to revoke a client certificate given to IBM HTTP Server by the browser when the key becomes compromised or when access permission to the key gets revoked. CRL represents a database which contains a list of certificates revoked before their scheduled expiration date.

If you want to enable certificate revocation in IBM HTTP Server, publish the CRL on a Lightweight Directory Access Protocol (LDAP) server. Once the CRL is published to an LDAP server, you can access the CRL using the IBM HTTP Server configuration file. The CRL determines the access permission status of the requested client certificate. Be aware, however, that it's not always possible to determine the revocation status of a client certificate if the backend server, the source of revocation data, is not available or not communicating properly with IBM HTTP Server.

Identifying directives needed to set up a certificate revocation list. The SSLClientAuth directive can include two options at once:

- SSLClientAuth 2 crl
- SSLClientAuth 1 crl

The CRL option turns CRL on and off inside an SSL virtual host. If you specify CRL as an option, then you elect to turn CRL on. If you do not specify CRL as an option, then CRL remains off. If the first option for SSLClientAuth equals 0/none, then you cannot use the second option, CRL. If you do not have client authentication on, then CRL processing does not take place.

Identifying directives supported in global or server and virtual host. Global server and virtual host support the following directives:

- SSLCRLHostname: The IP Address and host of the LDAP server, where the CRL database resides. Currently, you must configure any static CRL repositories to allow for checking of other URI forms in the CRLDistributionPoint fields.

z/OS Only an explicitly configured LDAP server can be queried for CRL, and the SSL handshake fails if the backend server is not reachable.

- SSLCRLPort: The port of the LDAP server where the CRL database resides; the default equals 389.
- SSLCRLUserID: The user ID to send to the LDAP server where the CRL database resides; defaults to anonymous if you do not specify the bind.
- SSLStashfile: The fully qualified path to file where the password for the user name on the LDAP server resides. This directive is not required for an anonymous bind. Use when you specify a user ID.

Use the **sslstash** command, located in the bin directory of IBM HTTP Server, to create your CRL password stash file. The password you specify using the **sslstash** command should equal the one you use to log in to your LDAP server.

Usage:

```
sslstash [-c] &lt;directory_to_password_file_and_file_name>; <function_name> <password>
```

where:

- **-c**: Creates a new stash file. If not specified, an existing file updates.
- **File**: Represents the fully qualified name of the file to create, or update.
- **Function**: Indicates the function for which to use the password. Valid values include crl, or crypto.
- **Password**: Represents the password to stash.

- **Distributed operating systems** SSLUnknownRevocationStatus: This directive allows you to configure how IBM HTTP Server will respond when fresh Certificate Revocation List (CRL) information or OCSP (Online Certificate Status Protocol) information is not available, and the client certificate that is currently offered is not known to be revoked from a previous query. Certificates are presumed not to be revoked, by default, which means they are valid, and a temporary failure to obtain CRL or OCSP information does not automatically result in an SSL handshake failure. This directive is provided to respond to circumstances in which a certificate has been accepted without IBM HTTP Server being able to reliably confirm the revocation status.

This directive has an effect only when all of these conditions are true:

- IBM HTTP Server is configured to accept client certificates with the SSLClientAuth directive.
- IBM HTTP Server is configured with one of the following directives: SSLOCSPEnable, SSLOCSPUrl, or SSLCRLHostname.
- An SSL client certificate is provided.
- IBM HTTP Server does not receive a valid OCSP or CRL response from the configured backend server, and the client certificate does not appear as revoked in a cached, but expired, CRL response.

IBM HTTP Server uses a cached CRL that is beyond its published expiration time when a current version is not available. When a certificate has been revoked in such an expired CRL, this will result in a direct SSL handshake failure that is outside the scope of the SSLUnknownRevocationStatus directive.

See the “SSL directives” on page 103 topic for more information.

CRL checking follows the URIDistributionPoint X509 extension in the client certificate as well as trying the DN constructed from the issuer of the client certificate. If the certificate contains a CRL Distribution Point (CDP), then that information is given precedence. The order in which the information is used is as follows:

1. CDP LDAP X.500 name
2. CDP LDAP URI
3. Issuer name combined with the value from the SSLCRLHostname directive

gotcha: If your certificates use the LDAP or HTTP URI forms of the CertificateDistributionPoint or AIA extensions, be sure that the IBM HTTP Server system can establish outgoing connections of this type; you might need to adjust the settings for your firewall.

Obtaining certificates:

This section provides information to help you get started with secure connections on the Web server. Obtaining certificates is the first step in securing your Web server.

About this task

When you set up secure connections, associate your public key with a digitally-signed certificate from a certificate authority (CA) that is designated as a trusted CA on your server.

Procedure

- **Buy a certificate from an external certificate authority provider.** You can buy a signed certificate by submitting a certificate request to a CA provider. The IBM HTTP Server supports several external certificate authorities. By default, many CAs exist as trusted CAs on the IBM HTTP Server. See “List of trusted certificate authorities on the IBM HTTP Server” on page 94.

Use the key management utility to create a new key pair and certificate request to send to an external CA, then define SSL settings in the `httpd.conf` file.

- **Distributed operating systems** IKEYMAN graphical user interface. If you are unable to use the IKEYMAN interface, use the command line interface `gskcmd` command.
- **z/OS** Native z/OS key management (`gskkyman` key database).
- **Create a self-signed certificate.** Use the key management utility or purchase certificate authority software from a CA provider.

Public Key Infrastructure

A Public Key Infrastructure (PKI) represents a system of digital certificates, certificate authorities, registration authorities, a certificate management service, and X.500 directories.

A PKI verifies the identity and the authority of each party that is involved in an Internet transaction, either financial or operational, with requirements for identity verification. Examples of these transactions include confirming the origin of proposal bids, or the author of e-mail messages.

A PKI supports the use of *certificate revocation lists* (CRLs). A CRL is a list of revoked certificates. CRLs provide a more global method for authenticating client identity by certificate, and can verify the validity of trusted CA certificates.

An X.500 directory server stores and retrieves CRLs and trusted CA certificates. The protocols used for storing and retrieving information from an X.500 directory server include Directory Access Protocol (DAP) and Lightweight Directory Access Protocol (LDAP). The IBM HTTP Server supports LDAP.

You can distribute information on multiple directory servers over the Internet and intranets, enabling an organization to manage certificates, trust policy, and CRLs from either a central location, or in a distributed manner. This capability makes the trust policy more dynamic because you can add or delete trusted CAs from a network of secure servers, without having to reconfigure each of the servers.

Session ID cache

IBM HTTP Server caches secure sockets layer (SSL) session IDs when Web clients establish secure connections with the Web server. Cached session IDs enable subsequent SSL session requests to use a shortened SSL handshake during session establishment. Session ID caching is enabled by default on all supported platforms.

AIX **HP-UX** **Linux** **Solaris** The session ID cache is implemented as a daemon process named **sidd**. You will see this process running when IBM HTTP Server is started with SSL enabled.

Distributed operating systems In most cases, you will not need to take an additional configuration steps to effectively use SSL session ID caching in IBM HTTP Server.

z/OS It is recommended that you disable IBM HTTP Server session ID caching (**sidd**). The z/OS System SSL provides an equivalent function that can perform better with some additional configuration.

- Disable the IBM HTTP Server **sidd** with the SSLCacheDisable directive and remove any existing SSLCacheEnable directives in `httpd.conf`.
- Enable "SSL Started Task" for z/OS System SSL. For more information on the following setup instructions, refer to the section "SSL Started Task" in *z/OS Cryptographic Services System Secure Sockets Layer (SSL) Programming (SC24-5901)*, which you can link to from the *z/OS Internet Library*:
 - Set the following environment variables in `bin/envvars`:
 - GSK_V3_SIDCACHE_SIZE=2048
 - GSK_V2_SIDCACHE_SIZE=2048
 - GSK_SYSPLEX_SIDCACHE=ON
 - export GSK_V3_SIDCACHE_SIZE GSK_V2_SIDCACHE_SIZE GSK_SYSPLEX_SIDCACHE
 - Configure the limits in the started task by editing `/etc/gskssl/server/envar`.
 - GSK_LOCAL_THREADS
 - GSK_SIDCACHE_SIZE

SSL directive considerations

When using SSL directives, you should consider the following: Limiting encryption to 128 bits or higher, rewriting HTTP (port 80) requests to HTTPS (port 443), logging SSL request information in the access log, and enabling certificate revocation lists (CRL).

You should consider the following when you want to enable SSL directives in the IBM HTTP Server `httpd.conf` configuration file:

- **Limiting IBM HTTP Server to encrypt at only 128 bits or higher.** There are several methods of configuring IBM HTTP Server to restrict and limit SSL to allow only 128 bit browsers and 128,168 bit ciphers access to Web content. For complete information, refer to Limiting IBM HTTP Server to encrypt at only 128 bits or higher .
- **How to rewrite HTTP (port 80) requests to HTTPS (port 443).** The `mod_rewrite.c` rewrite module provided with IBM HTTP Server can be used as an effective way to automatically rewrite all HTTP requests to HTTPS. For complete information refer to How to rewrite HTTP (port 80) requests to HTTPS (port 443).

- **Logging SSL request information in the access log for IBM HTTP Server.** The IBM HTTP Server implementation provides Secure Sockets Layer (SSL) environment variables that are configurable with the LogFormat directive in the httpd.conf configuration file. For complete information refer to Logging SSL request information in the access log for IBM HTTP Server.
- **Enabling certificate revocation lists (CRL) in IBM HTTP Server.** Certificate revocation provides the ability to revoke a client certificate given to the IBM HTTP Server by the browser when the key is compromised or when access permission to the key is revoked. CRL represents a database that contains a list of certificates revoked before their scheduled expiration date. For complete information refer to “SSL certificate revocation list” on page 94.

Authentication

Authentication verifies identity.

The server uses authentication in two ways:

- **Digital signature.** A digital signature represents a unique mathematically computed signature that ensures accountability. Think of a digital signature as similar to a credit card, on which your photo displays. To verify the identity of the person that is sending you a message, look at the digital certificate of the sender.
- **Digital certificate.** A digital certificate, or digital ID, is similar to having a credit card with a picture of the bank president with his arm around you. A merchant trusts you more because not only do you look like the picture on the credit card, the bank president trusts you, too.

You base your trust of the sender authenticity on whether you trust the third party, a person, or agency that certified the sender. The third party issuing digital certificates is called a certificate authority (CA) or *certificate signer*.

A digital certificate contains:

- The public key of the person getting certified
- The name and address of the person or organization getting certified, also known as the *distinguished name*
- The digital signature of the CA
- The issue date of the certificate
- The expiration date of the certificate

You enter your distinguished name as part of a certificate request. The digitally signed certificate includes your distinguished name and the distinguished name of the CA.

You can request one of the following certificates:

- A server certificate to do commercial business on the Internet from VeriSign or some other CA. For a list of supported CAs, see *Buying a certificate from an external CA provider*.
- A server certificate that you create for your own private Web network.

CAs broadcast their public key and distinguished name bundled together so that people add them to their Web servers and browsers, as a trusted CA certificate. When you designate the public key and certificate from a CA to become a trusted CA certificate, your server trusts anyone who has a certificate from that CA. You can have many trusted CAs as part of your server. The HTTP Server includes several default trusted CA certificates.

Distributed operating systems You can add or remove trusted CAs using the IBM Key Management utility (ikeyman) that is included with your server.

z/OS You can add or remove trusted CAs using the native z/OS key management (gskkyman).

To communicate securely, the receiver in a transmission must trust the CA who issued the sender certificate. This situation remains true whether the receiver is a Web server or a browser. When a sender signs a message, the receiver must have the corresponding CA-signed certificate and public key designated as a trusted CA certificate.

Encryption

Encryption in its simplest form involves scrambling a message so that no one can read the message until it is unscrambled by the receiver.

The sender uses an algorithmic pattern, or a key to scramble, or encrypt the message. The receiver has the decryption key. Encryption ensures privacy and confidentiality in transmissions sent over the Internet.

Use two different kinds of keys for encryption:

Asymmetric keys. You create a key pair with asymmetric keys. The key pair consists of a public key and a private key, which differ from each other. The private key holds more of the secret encryption pattern than the public key. Do not share your private key with anyone.

The server uses its private key to sign messages to clients. The server sends its public key to clients so that they can encrypt messages to the server, which the server decrypts with its private key. Only you can decrypt a message that is encrypted with your public key because only you have the private key. Key pairs are stored in a key database that is protected by a password.

Symmetric keys. Symmetric keys follow an older model of the sender and receiver sharing some kind of pattern. The sender uses this same pattern to encrypt the message and the receiver uses this pattern to decrypt the message. The risk involved with symmetric keys centers around finding a safe transportation method to use, when sharing your secret key with the people to which you want to communicate.

The Secure Sockets Layer (SSL) protocol uses both asymmetric and symmetric key exchange. Use asymmetric keys for the *SSL handshake*. During the handshake, the master key, encrypted with the receiver public key passes from the client to the server. The client and server make their own session keys using the master key. The session keys encrypt and decrypt data for the remainder of the session. Symmetric key exchange occurs during the exchange of the cipher specification, or encryption level.

The server needs a *digital certificate*, which is an encrypted message that authenticates Web content, to send its public key to clients. A certificate authority (CA), which signs all certificates that it issues with a private key, issues this certificate and verifies the identity of the server.

Secure Sockets Layer environment variables

The `mod_ibm_ssl` parameter provides access to information about an Secure Sockets Layer (SSL) session by setting variables in the Apache API `subprocess_env` table for the active request. These variables are considered environment variables because of how information is accessed when the variables are passed to CGI applications.

You can categorize SSL environment variables into three types based on the type of information that is accessed when the variable is passed to the application.

- Variables for information regarding the SSL handshake
- Variables for exposing the server certificate information
- Variables for exposing client certificate information, when client authentication is enabled.

The following table provides the types of access to information as well as the mechanisms used to access information using SSL environment variables.

Table 5. Types of access and mechanisms for SSL environment variables

Access type	Mechanism
access from a CGI or FastCGI application	The information is passed to the CGI application as an environment variable. Use the method provided by the implementation language for accessing environments, such as <code>getenv ("HTTPS")</code> in C or <code>\$ENV{'HTTPS'}</code> in Perl. For a SSL environment variable to be used in CGI or FastCGI, there must be a corresponding <code>PassEnv</code> directive.
access from a plug-in module	The information is available in the <code>subprocess_env</code> table after the quick handler has run. Access it with a call such as <code>apr_table_lookup (r->subprocess_env, "HTTPS")</code>
logging in the access log with other information about the request	Use the following <code>%{varname}e</code> example. <pre>LogFormat "%h %l %u %t \ "%r\ " %>s %b %{HTTPS}e" ssl-custom</pre> <p>If the information is not available, <code>mod_log_config</code> logs a dash (-) for the field.</p>
use with the <code>setenvif</code> variable	# Silly example, don't compress SSL connections <pre>SetEnvIf HTTPS no-gzip</pre>
use as part of a <code>mod_rewrite</code> rule variable	<pre>RewriteEngine On RewriteCond %{ENV:HTTPS} ^OFF\$ RewriteRule .* /no-ssl.html</pre>
access in an SSI document	In order for an SSL environment variable to be used in an SSI document, there must be a corresponding <code>PassEnv</code> directive. <pre>SSL is <!--#echo var="HTTPS" --></pre>
access control	<code>Allow from env=HTTPS</code>

SSL handshake environment variables

Secure Sockets Layer (SSL) handshake environment variables are used to access server certificate information. When an SSL handshake is successfully completed, the SSL handshake environment variables are automatically set.

Variables

Table 6. SSL handshake environment variables. The table provides a list of SSL handshake environment variables with their descriptions and values.

SSL handshake environment variable	Description	Value
HTTPS	Indicates SSL connection	String contains either ON, for an SSL connection, or OFF, if not.
HTTPS_CIPHER	Contains the cipher used in the SSL handshake.	See the following table.
HTTPS_KEYSIZE	Indicates the size of the key.	See the following table.
HTTPS_SECRETKEYSIZE	Indicates the strength of the key.	See the following table.

Table 6. SSL handshake environment variables (continued). The table provides a list of SSL handshake environment variables with their descriptions and values.

SSL_PROTOCOL_VERSION	Contains the protocol version.	<p>z/OS String contains either SSLV2, SSLV3, or TLSV1 for Transport Layer Security (TLS) Version 1.0).</p> <p>Distributed operating systems String contains SSLV2, SSLV3, TLSV1 for TLS Version 1.0, or TLSV1.1 for TLS Version 1.1.</p>
----------------------	--------------------------------	--

Table 7. Variables for HTTPS_KEYSIZE and HTTPS_SECRETKEYSIZE in Secure Sockets Layer V3 and Transport Layer Security V1. The table provides the cipher suite, the key size, and the secret key size.

Cipher suite	Key size	Secret key size
SSL_RSA_WITH_NULL_MD5	0	0
SSL_RSA_WITH_NULL_SHA	0	0
SSL_RSA_EXPORT_WITH_RC4_40_MD5	128	40
SSL_RSA_WITH_RC4_128_MD5	128	128
SSL_RSA_WITH_RC4_128_SHA	128	128
SSL_RSA_EXPORT_WITH_RC2_CBC_40_MD5	128	40
SSL_RSA_WITH_DES_CBC_SHA	64	56
SSL_RSA_WITH_3DES_EDE_CBC_SHA	192	168
SSL_NULL_WITH_NULL_NULL	0	0
TLS_RSA_EXPORT1024_WITH_RC4_56_SHA	56	20
TLS_RSA_EXPORT1024_WITH_DES_CBC_SHA	56	20

Table 8. Variables for HTTPS_CIPHER in Secure Sockets Layer V2.. The table provides the cipher suite, the key size, and the secret key size.

Cipher suite	Key size	Secret key size
RC4_128_WITH_MD5	128	128
RC4_128_EXPORT40_WITH_MD5	128	40
RC2_128_CBC_WITH_MD5	128	128
RC2_128_CBC_EXPORT40_WITH_MD5	128	40
DES_64_CBC_WITH_MD5	64	56
DES_192_EDE3_CBC_WITH_MD5	192	168

Server certificate environment variables

Server certificate environment variables are used to access server certificate information. The server certificate environment variables are automatically set. If client authentication is not configured, references to these values are empty.

Variables

The following table provides a list of server certificate environment variables with their descriptions and values.

Server certificate environment variable	Description	Value
SSL_SERVER_C	Contains the country attribute of the server certificate	String
SSL_SERVER_CN	Contains the common name attribute of the server certificate	String
SSL_SERVER_DN	Contains the distinguished name of the server certificate used in the IP-based virtual host which received the request	String
SSL_SERVER_EMAIL	Contains the e-mail attribute of the server certificate	String
SSL_SERVER_L	Contains the locality attribute of the server certificate	String
SSL_SERVER_O	Contains the organization attribute of the server certificate	String
SSL_SERVER_OU	Contains the organizational unit attribute of the server certificate	String
SSL_SERVER_ST	Contains the state or province attribute of the server certificate	String

Client certificate environment variables

Client certificate environment variables are used to access client certificate information when client authentication is enabled. If client authentication is not enabled, references to these values are empty.

Variables

The following table provides a list of client certificate environment variables and their descriptions and values.

SSL client certificate environment variable	Description	Value
SSL_CLIENT_C	Contains the client certificate country	String
SSL_CLIENT_CERTBODY	Contains the client certificate	This value is the unformatted body of the client certificate, if a certificate was provided by the client
SSL_CLIENT_CERTBODYLEN	Contains the length of the client certificate	Integer
SSL_CLIENT_CN	Contains the client certificate common name	String
SSL_CLIENT_DN	Contains the distinguished name from the client certificate	String
SSL_CLIENT_EMAIL	Contains the client certificate e-mail	String
SSL_CLIENT_IC	Contains the country name of the client certificate issuer	String

SSL_CLIENT_ICN	Contains the common name of the client certificate issuer	String
SSL_CLIENT_IDN	Contains the distinguished name of the client certificate issuer	String
SSL_CLIENT_IEMAIL	Contains the e-mail address of the client certificate issuer	String
SSL_CLIENT_IL	Contains the locality of the client certificate issuer	String
SSL_CLIENT_IO	Contains the organization name of the client certificate issuer	String
SSL_CLIENT_IOU	Contains the organizational unit name of the client certificate issuer	String
SSL_CLIENT_IPC	Contains the postal code of the client certificate issuer	String
SSL_CLIENT_IST	Contains the state or province of the client certificate issuer	String
SSL_CLIENT_L	Contains the client certificate locality	String
SSL_CLIENT_NEWSESSIONID	Indicates whether this session ID is new	String. This value must be TRUE or FALSE.
SSL_CLIENT_O	Contains the client certificate organization	String
SSL_CLIENT_OU	Contains the client certificate organizational unit	String
SSL_CLIENT_PC	Contains the client certificate postal code	String
SSL_CLIENT_SERIALNUM	Contains the client certificate serial number	String
SSL_CLIENT_SESSIONID	Contains the session ID	String
SSL_CLIENT_ST	Contains the client certificate state or province	String

SSL directives

Secure Sockets Layer (SSL) directives are the configuration parameters that control SSL features in IBM HTTP Server.

Most SSL directives in IBM HTTP Server have the same behavior. A directive specified for a given virtual host configuration overrides a directive specified in the base server configuration. Also, a directive specified for a child directory overrides a directive specified for its parent directory. However, there are exceptions.

For example, when no directive is specified for a virtual host, the directive specified in the base server configuration might be copied to the virtual host configuration. In this case, the directive in the base server configuration overrides the virtual host configuration.

Attention: The SSLEnable directive should not be specified in the base server configuration if you do not want the directive automatically copied to a given virtual host configuration.

Also, a directive specified for a child directory might be appended to the directive specified for its parent directory. In this case, the directive for the parent directory does not override the directive for the child directory, but instead is appended to it and both directives are applied to the child directory.

The following list contains the SSL directives for IBM HTTP Server.

- “SSLOCSResponderURL”
- “SSLOCSPEnable” on page 105
- “Keyfile directive” on page 105
- “SSLAcceleratorDisable directive” on page 106
- “SSLAllowNonCriticalBasicConstraints directive” on page 107
- “SSLCacheDisable directive” on page 107
- “SSLCacheEnable directive” on page 107
- “SSLCacheErrorLog directive” on page 108
- “SSLCachePath directive” on page 108
- “SSLCachePortFilename directive” on page 108
- “SSLCacheTraceLog directive” on page 108
- “SSLCipherBan directive” on page 109
- “SSLCipherRequire directive” on page 109
- “SSLCipherSpec directive” on page 109
- “SSLClientAuth directive” on page 110
- “SSLClientAuthGroup directive” on page 112
- “SSLClientAuthRequire directive” on page 113
- “SSLClientAuthVerify directive” on page 114
- “SSLCRLHostname directive” on page 115
- “SSLCRLPort directive” on page 116
- “SSLCRLUserID directive” on page 116
- “SSLDisable directive” on page 116
- “SSLEnable directive” on page 117
- “SSLFakeBasicAuth directive” on page 117
- “SSLFIPSDisable directive” on page 117
- “SSLFIPSEnable directive” on page 117
- “SSLPKCSDriver directive” on page 118
- “SSLProtocolDisable directive” on page 118
- “SSLProxyEngine directive” on page 119
- “SSLServerCert directive” on page 120
- “SSLStashfile directive” on page 120
- “SSLTrace directive” on page 121
- **Distributed operating systems** “SSLUnknownRevocationStatus” on page 121
- “SSLV2Timeout directive” on page 122
- “SSLV3Timeout directive” on page 122
- “SSLVersion directive” on page 122

SSLOCSResponderURL

The SSLOCSResponderURL directive enables checking of client certificates through a statically configured online certificate status protocol (OCSP) responder.

Syntax

Distributed operating systems SSLOCSResponderURL<URL>

Scope

Virtual host

Default

Disabled

Module

mod_ibm_ssl

Multiple instances in the configuration file

One per virtual host

Values

A fully qualified URL that points to an OCSP responder, for example, http://hostname:2560/. The path portion of the URL is not used when submitting OCSP requests.

Even if CRL checking is configured, OCSP checking is performed before any CRL checking. CRL checking occurs only if the result of the CRL is unknown or inconclusive.

If SSLOCSPResponderURL is set, IBM HTTP Server uses the supplied URL to check for certificate revocation status when an SSL client certificate is provided.

If both SSLOCSPEnable and SSLOCSPResponderURL are configured, the responder defined by SSLOCSPResponderURL is checked first. If the revocation status is unknown or inconclusive, IBM HTTP Server checks OCSP responders for SSLOCSPEnable.

Note: In some cases IBM HTTP Server might not be able to determine the revocation status of a client certificate, because the backend server, which is the source of the revocation data, is not available. You should be aware that:

- A static CRL repository (SSLCRLHost) must be configured to enable checking of other URI forms in the CRLDistributionPoint fields.
- If your certificates use the LDAP or HTTP URI forms of the CertificateDistributionPoint or AIA extensions, be sure that the IBM HTTP Server system can establish outgoing connections of this type; you must adjust the settings for your firewall.
- **Distributed operating systems** The SSLUnknownRevocationStatus directive is provided for cases in which recoverable errors occur in IBM HTTP Server when it is communicating with the backend server, and the IBM HTTP Server cannot determine the revocation status of a certificate. The default behavior is to continue processing the handshake unless the backend server can successfully indicate that the certificate is revoked.
- Only an explicitly configured LDAP server can be queried for CRL, and the SSL handshake fails if the backend server is not reachable.

SSLOCSPEnable

The SSLOCSPEnable directive enables checking of client certificates through OCSP responders defined in the Authority Information Access (AIA) extension of their certificate.

Syntax

Distributed operating systems SSLOCSPEnable

Scope

Virtual host

Default

Disabled

Module

mod_ibm_ssl

Multiple instances in the configuration file

One instance permitted for each virtual host

Values

None

If SSLOCSPEnable is set, and an SSL client certificate chain contains an AIA extension, IBM HTTP Server contacts the OCSP responder indicated by the AIA extension to check revocation status of the client certificate. The path portion of the URL is ignored.

If both OCSP and CRL checking is configured, OCSP checking is performed before any CRL checking. CRL checking occurs only if the result of the OCSP checking is unknown or inconclusive.

If both SSLOCSPEnable and SSLOCSPResponderURL are configured, the responder defined by SSLOCSPResponderURL is checked first. If the revocation status is unknown or inconclusive, IBM HTTP Server checks OCSP responders for SSLOCSPEnable.

Keyfile directive

The keyfile directive sets the key file to use.

Attention: This directive might be overridden by the base server configuration.

Syntax

AIX **Solaris** **Linux** **Windows** Keyfile
[/prompt] /fully qualified path to key
file/keyfile.kdb

Attention: **z/OS** The /prompt function is only supported when running from a USS shell, not from a JCL started job. If you attempt to use the /prompt function from a JCL started job, then a configuration error occurs.

z/OS You can use a keyring stored in the Hierarchical File System (HFS) or in the System Authorization Facility (SAF). To use a keyring stored in HFS:

- Keyfile /fully qualified path to key
file/keyfile.kdb

To use a keyring stored in SAF:

- Keyfile /saf WASKeyring

Important: With SAF keyrings:

- There is no stash file when using SAF, and access is controlled by SAF rules. Therefore, if you attempt to use the keyfile/prompt/saf argument, the argument is not supported. An attempt to use this argument results in a configuration error.
- The ID that is used to start IBM HTTP Server must have access to the keyring named in this directive. If the ID does not have access, SSL initialization fails.

Global base and virtual host

None

mod_ibm_ssl

One instance per virtual host and global server

File name of the key file.

Distributed operating systems Use the prompt option to enable the HTTP server to prompt you for the Key file password at start time.

z/OS File system protection can be used to limit access. Use the SAF (System Authorization Facility) keyrings for limiting access to SSL certificates.

Scope

Default

Module

Multiple instances in the configuration file

Values

Important: **z/OS** The z/OS system does not support key database files created on other platforms. Key database files used for z/OS systems must be created on the z/OS platform.

You can use only one of the following configurations for the key file type:

- **Distributed operating systems** Certificate Management Services (CMS)
- **z/OS** CMS or Resource Access Control Facility (RACF)

SSLAcceleratorDisable directive

The SSLAcceleratorDisable directive disables the accelerator device.

Syntax

SSLAcceleratorDisable

Scope

Virtual and global

Default

Accelerator device is enabled

Module

mod_ibm_ssl

Multiple instances in the configuration file
Values

One instance per virtual host.
None. Place this directive anywhere inside of the configuration file, including inside a virtual host. During initialization, if the system determines that an accelerator device is installed on the machine, the system uses that accelerator to increase number of secure transactions. This directive does not take arguments.

Distributed operating systems

SSLAllowNonCriticalBasicConstraints directive

The SSLAllowNonCriticalBasicConstraints directive enables compatibility with one aspect of the GPKI specification from the government of Japan that conflicts with RFC3280.

Syntax
Scope
Default
Module
Multiple instances in the configuration file
Values

SSLAllowNonCriticalBasicConstraints *on|off*
Global server or virtual host
Off
mod_ibm_ssl
One instance per virtual host and global server
None. This directive changes the behavior of the certificate validation algorithm such that a non-critical basic constraints extension on an issuer certificate authority (CA) certificate does not cause a validation failure. This enables compatibility with one aspect of the GPKI specification from the government of Japan that conflicts with RFC3280.

Attention: RFC3280 states that this extension must appear as a critical extension in all CA certificates that contain public keys used to validate digital signatures on certificates.

AIX HP-UX Linux Solaris z/OS

SSLCacheDisable directive

The SSLCacheDisable directive disables the external SSL session ID cache.

Syntax
Scope
Default
Module
Multiple instances in the configuration file
Values

SSLCacheDisable
One per physical Apache server instance, enabled only outside of virtual host stanzas.
None
mod_ibm_ssl
Not permitted.
None.

AIX HP-UX Linux Solaris z/OS

SSLCacheEnable directive

The SSLCacheEnable directive enables the external SSL session ID cache.

Syntax
Scope
Default
Module

SSLCacheEnable
One per physical Apache server instance, enabled only outside of virtual host stanzas.
None
mod_ibm_ssl

Multiple instances in the configuration file Not permitted.
Values None.

AIX HP-UX Linux Solaris z/OS

SSLCacheErrorLog directive

The SSLCacheErrorLog directive sets the file name for session ID cache.

Syntax SSLCacheErrorLog /usr/HTTPServer/logs/sidd_logg
Scope Server configuration outside of virtual host.
Default None
Module mod_ibm_ssl
Multiple instances in the configuration file Not permitted.
Values Valid file name.

AIX HP-UX Linux Solaris z/OS

SSLCachePath directive

The SSLCachePath directive specifies the path to the session ID caching daemon.

Syntax SSLCachePath /usr/HTTPServer/bin/sidd
Scope Server configuration outside of virtual host.
Default <server-root>/bin/sidd
Module mod_ibm_ssl
Multiple instances in the configuration file Not permitted.
Values Valid path name.

AIX HP-UX Linux Solaris z/OS

SSLCachePortFilename directive

The SSLCachePortFilename directive sets the file name for the UNIX domain socket that is used for communication between the server instances and the session ID cache daemon. You must set this directive if you run two instances of IBM HTTP Server from the same installation directory and both instances are configured for SSL. Otherwise, you do not need to set this directive.

Syntax SSLCachePath /usr/HTTPServer/logs/sidd
Scope Server configuration outside of virtual host.
Default If this directive is not specified and the cache is enabled, the server attempts to use the <server-root>/logs/siddport file.
Module mod_ibm_ssl
Multiple instances in the configuration file Not permitted.
Values Valid path name. The web server deletes this file during startup; do not name.

AIX HP-UX Linux Solaris z/OS

SSLCacheTraceLog directive

The SSLCacheTraceLog directive specifies the file to which the session ID trace messages are written. Without this directive, tracing is disabled.

Syntax SSLCacheTraceLog /usr/HTTPServer/logs/sidd-trace.log
Scope Server configuration outside of virtual host.

Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	Not permitted.
Values	Valid path name.

SSLCipherBan directive

The SSLCipherBan directive denies access to an object if the client has connected using one of the specified ciphers. The request fails with a 403 status code.

Attention: This directive, when specified for a child directory, does not override the directive specified for the parent directory. Instead, both directories are applied to the child directory.

Syntax	SSLCipherBan < <i>cipher_specification</i> >
Scope	Multiple instances per directory stanza.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	Permitted per directory stanza. Order of preference is top to bottom.
Values	See the SSL cipher specification topic for values.

SSLCipherRequire directive

The SSLCipherRequire directive restricts access to objects to clients that have connected using one of the specified ciphers. If access is denied, the request fails with a '403' status code.

Attention: This directive, when specified for a child directory, does not override the directive specified for the parent directory. Instead, both directories are applied to the child directory.

Syntax	SSLCipherRequire < <i>cipher_specification</i> >
Scope	Multiple instances per directory stanza.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	Permitted per directory stanza.
Values	See the SSL cipher specification topic for values

SSLCipherSpec directive

The SSLCipherSpec directive enables you to customize the SSL ciphers supported during the handshake. You can customize the set of SSL ciphers and the order of preference of the SSL ciphers.

Distributed operating systems On distributed platforms, each protocol has its own ordered list of ciphers. The supported protocols are SSL version 2, SSL version 3, TLS version 1.0, TLS version 1.1, and TLS version 1.2.

z/OS On z/OS, there are only two lists of enabled ciphers, one for SSL version 2 and one for the other protocols. The supported protocols are SSL version 2, SSL version 3, and TLS version 1.0.

SSL Version 2 ciphers default to no ciphers, which means that the protocol is disabled. The other protocols default to a set of SSL ciphers that excludes null ciphers, export ciphers, and weak ciphers.

When you use the single-argument form of SSLCipherSpec, the given cipher is enabled in all protocols for which it is valid. The first time such a change is made for each protocol, the default ciphers for the protocol are discarded.

When you use the multiple-argument form of SSLCipherSpec, specifying the name of an SSL protocol (or "ALL") as the first argument, you can use an enhanced syntax with the following benefits:

- Multiple ciphers can be listed with each occurrence of SSLCipherSpec
- Individual ciphers can be removed from the current set of enabled ciphers by prefixing the cipher name with "-".
- The first time a given protocol cipher list is being modified, the given cipher can be added to the end of the defaults, instead of replacing them, by prefixing the cipher name with "+".

If you provide a protocol name of "ALL", then the adding or removing specified for each cipher name is applied to each protocol where that cipher is valid.

As a special case, to empty all the cipher lists with a single command, you can use SSLCipherSpec ALL NONE. Using this command is a good way to start a configuration anytime you do not want to use the default ciphers.

<p>z/OS Syntax</p> <p>Distributed operating systems Syntax</p> <p>Scope</p> <p>Default</p> <p>Module</p> <p>Multiple instances in the configuration file</p> <p>Distributed operating systems Values for protocol name on distributed platforms</p> <p>z/OS Values for protocol name on z/OS</p> <p>Values for cipher names</p> <p>Example 1</p>	<p>SSLCipherSpec <i>short name</i> or SSLCipherSpec <i>long name</i></p> <p>SSLCipherSpec [<i>protocol_name</i>] [+ -]<i>short name</i> <i>long name</i> [[+ -]<i>short name</i> <i>long name</i> ...]</p> <p>Virtual host.</p> <p>If nothing is specified, the server uses all non-NULL, non-export, non-weak cipher specifications available.</p> <p>mod_ibm_ssl</p> <p>Permitted. Order of preference is top to bottom, first to last.</p> <p>SSLv2, SSLv3, TLSv10, TLSv11, TLSv12, ALL</p> <p>SSLv2, SSLv3, TLSv10, ALL</p> <p>See the SSL cipher specification topic for values</p> <p>If you want to select just a few ciphers, it is best to start by resetting all the cipher lists, then adding the ones you want to use:</p> <pre># Delete all ciphers from the cipher lists SSLCipherSpec ALL NONE # Add a few specific ciphers SSLCipherSpec ALL +SSL_RSA_WITH_3DES_EDE_CBC_SHA SSLCipherSpec ALL +TLS_RSA_WITH_AES_256_CBC_SHA</pre> <p>If you want to use most of the defaults, but there are one or two ciphers that you do not want, you can remove those from any lists that they are currently in:</p> <pre># Delete some specific ciphers from the protocols where they are valid SSLCipherSpec ALL -SSL_RSA_WITH_RC4_128_MD5 SSLCipherSpec ALL -SSL_RSA_WITH_RC4_128_SHA</pre>
<p>Example 2</p>	

SSLClientAuth directive

The SSLClientAuth directive sets the mode of client authentication to use (none (0), optional (1), or required (2)).

<p>Syntax</p> <p>Scope</p> <p>Default</p> <p>Module</p> <p>Multiple instances in the configuration file</p>	<p>SSLClientAuth <<i>level required</i>> [<i>cr</i>]</p> <p>Virtual host.</p> <p>SSLClientAuth none</p> <p>mod_ibm_ssl</p> <p>One instance per virtual host.</p>
--	--

Values

- 0/None: No client certificate requested.
- 1/Optional: Client certificate requested, but not required.
- 2/Required: Valid client certificate required.
- Required_reset: The server requires a valid certificate from all clients, and if no certificate is available, the server sends an SSL alert to the client. This alert enables the client to understand that the SSL failure is client-certificate related, and causes browsers to re-prompt for client certificate information about subsequent access. This option requires GSKit version 7.0.4.19 or later, or z/OS V1R8 or later.
Important: No numeric option is provided so it will not look like the existing options.
- CRL: Turns certificate revocation list (CRL) on and off inside an SSL virtual host. If you use CRL, you must specify `cr1` as a second argument for `SSLClientAuth`. For example: `SSLClientAuth 2 cr1`. If you do not specify `cr1`, you cannot perform CRL in an SSL virtual host.
- `noverify`: Enables SSL handshake to succeed and establish a connection, even if the certificate provided by the client fails validation (for example, the certificate is expired or revoked). Use this directive with `SSLClientAuthVerify` to provide a user-friendly web page, instead of the default browser error message. This option is only valid with `Optional`. Use `SSLClientAuthVerify` to fail requests received on the connection with the invalid client certificate.

If you specify the value 0/None, you cannot use the CRL option.

The server requires a valid certificate from all clients, and if no certificate is available, the server sends an SSL alert to the client. This enables the client to understand that the SSL failure is client-certificate related, and causes browsers to re-prompt for client certificate information about subsequent access. This option requires GSKit version 7.0.4.19 or later, or z/OS V1R8 or later.

Required_reset

Attention: In some cases IBM HTTP Server might not be able to determine the revocation status of a client certificate, because the backend server, which is the source of the revocation data, is not available. You should be aware that:

- A static CRL repository (`SSLCRLHost`) must be configured to enable checking of other URI forms in the `CRLDistributionPoint` fields.
- If your certificates use the LDAP or HTTP URI forms of the `CertificateDistributionPoint` or `AIA` extensions, be sure that the IBM HTTP Server system can establish outgoing connections of this type; you must adjust the settings for your firewall.
- **Distributed operating systems** The `SSLUnknownRevocationStatus` directive is provided for cases in which recoverable errors occur in IBM HTTP Server when it is communicating with the backend server, and the IBM HTTP Server cannot determine the revocation status of a certificate. The default behavior is to continue processing the handshake unless the backend server can successfully indicate that the certificate is revoked.
- Only an explicitly configured LDAP server can be queried for CRL, and the SSL handshake fails if the backend server is not reachable.

SSLClientAuthGroup directive

The SSLClientAuthGroup directive defines a named expression group that contains a set of specific client certificate attribute and value pairs. This named group can be used by the SSLClientAuthRequire directives. A certificate must be provided by the client, which passes this expression, before the server allows access to the protected resource.

Syntax	SSLClientAuthGroup <i>group name attribute expression</i>
Scope	Server config, virtual host.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	Permitted.
Override	None.
Values	Logical expression consisting of attribute checks linked with AND, OR, NOT, and parentheses. For example: SSLClientAuthGroup IBMUSpeople (Org = IBM) AND (Country = US)

The following section provides a description of examples with valid logical expressions. For example: SSLClientAuthGroup ((CommonName = "Fred Smith") OR (CommonName = "John Deere")) AND (Org = IBM) means that the object is not served, unless the client certificate contains a common name of either Fred Smith or John Deere and the organization is IBM. The only valid comparisons for the attribute checks, are equal and not equal (= and !=). You can link each attribute check with AND, OR, or NOT (also &&, ||, and !). Any comparisons that you link with AND, OR, or NOT must be contained within parentheses. If the value of the attribute contains a non-alphanumeric character, you must delimit the value with quotation marks.

This list contains attribute values that you can specify for this directive:

Table 9. Attribute values for the SSLClientAuthGroup directive. The table lists each attribute value as a long name and short name.

Long name	Short name
CommonName	CN
Country	C
Email	E
IssuerCommonName	ICN
IssuerEmail	IE
IssuerLocality	IL
IssuerOrg	IO
IssuerOrgUnit	IOU
IssuerPostalCode	IPC
IssuerStateOrProvince	IST
Locality	L
Org	O
OrgUnit	OU
PostalCode	PC
StateOrProvince	ST

The long name or the short name can be used in this directive.

The user specifies a logical expression of specific client certificate attributes. You can logically use AND , OR, or NOT for multiple expressions if you must specify groupings of client certificate attribute values. Any comparisons that are linked with AND, OR, or NOT must be contained within parentheses. Valid operators include '=' and '!='. For example:

```
SSLClientAuthGroup IBMpeople Org = IBM)
```

or

```
SSLClientAuthGroup
NotMNIIBM (ST != MN) && (Org = IBM)
```

A group name cannot include spaces. See “SSLClientAuthRequire directive” for more information.

SSLClientAuthRequire directive

The SSLClientAuthRequire directive specifies attribute values, or groups of attribute values, that must be validated against a client certificate before the server allows access to the protected resource.

Syntax	SSLClientAuthRequire <i>attribute expression</i>
Scope	server config, virtual host
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	Permitted. The function joins these directives by "AND".
Override	AuthConfig
Values	Logical expression consisting of attribute checks linked with AND, OR, NOT, and parentheses. For example: SSLClientAuthRequire (group != IBMpeople) && (ST = M)

If the certificate you received does not have a particular attribute, then there is no verification for an attribute match. Even if the specified matching value is " ", this might still not be the same as not having the attribute there at all. Any attribute specified on the SSLClientAuthRequire directive that is not available on the certificate causes the request to be rejected.

The list contains attribute values that you can specify for this directive:

Table 10. Attribute values for the SSLClientAuthRequire directive. The table lists each attribute value as a long name and short name.

Long name	Short name
CommonName	CN
Country	C
Email	E
IssuerCommonName	ICN
IssuerEmail	IE
IssuerLocality	IL
IssuerOrg	IO
IssuerOrgUnit	IOU
IssuerPostalCode	IPC
IssuerStateOrProvince	IST
Locality	L
Org	O

Table 10. Attribute values for the SSLClientAuthRequire directive (continued). The table lists each attribute value as a long name and short name.

Long name	Short name
OrgUnit	OU
PostalCode	PC
StateOrProvince	ST

The long name or the short name can be used in this directive.

The user specifies a logical expression of specific client certificate attributes. You can logically use AND , OR, or NOT for multiple expressions if you must specify groupings of client certificate attribute values. Any comparisons that are linked with AND, OR, or NOT must be contained within parentheses. Valid operators include '=' and '!='. The user can also specify a group name, that is configured using the “SSLClientAuthGroup directive” on page 112, to configure a group of attributes.

You can specify multiple SSLClientAuthRequire directives within the same scope. The logical expressions for each directive are used to evaluate access rights for each certificate, and the results of the individual evaluations are logically ANDed together. For example:

```
SSLClientAuthRequire
((CommonName="John Doe") || (StateOrProvince=MN)) && (Org
!=IBM)
```

or

```
SSLClientAuthRequire
(group!=IBMpeople) && (ST=MN)
```

You can put quotes around the short and long names. For example:

```
SSLClientAuthRequire (group
!= IBMpeople) && ("ST= MN")
```

See “SSLClientAuthGroup directive” on page 112 for more information.

SSLClientAuthVerify directive

The SSLClientAuthVerify directive controls whether IBM HTTP Server fails requests when a client certificate is received, but it fails validation (for example, it is expired or revoked).

Syntax	SSLClientAuthVerify statuscode OFF
Scope	Global server or virtual host.
Default	500
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per directory stanza.
Values	HTTP response status code, or OFF

Use this directive with SSLClientAuth Optional Noverify to provide a user friendly web page, instead of the default browser error message.

If you configure a virtual host with SSLClientAuth Optional Noverify, an SSL connection can be established when a client certificate is received, but it fails validation (for example, it is expired or revoked).

Use this directive in a context such as Location or Directory to fail requests that are received on that connection with a specific error code, or handled normally by setting OFF.

By providing a custom error document for that status, the administrator can control the page that is presented to the user, for example, to tell the user their certificate is invalid and provide further instructions.

If the error document is an internal redirect to another URL in the same virtual host, you must ensure that URL has `SSLClientAuthVerify OFF` in its context so it does not immediately fail, as well. An example of this scenario follows.

The specified status code must be a response status that is valid in HTTP and known to IBM HTTP Server. The values are between 100 and 599, and are typically defined in an RFC or standards proposal. If you are unsure, try a status code in a test configuration and use `apachectl -t` to see if it is valid. Other unused codes that are valid and would be good choices include: 418, 419, 420, and 421.

Because the client certificate was invalid, the error document will not have any of the environment variables available that would contain information about the client certificate. The cause of the client certificate validation failure is available in the `SSL_LAST_VALIDATION_ERROR` environment variable. The variable could be `GSKVAL_ERROR_REVOKED_CERT` or `GSKVAL_ERROR_CERT_EXPIRED`. If the certificate has multiple validation problems, the reported cause is often `GSKVAL_ERROR_CA_MISSING_CRITICAL_BASIC_CONSTRAINT`.

Each time a client certificate validation fails, two messages are logged in the error log at `LogLevel Error`. The second message includes the cause, for example:

```
[Tue Jun 08 08:54:25 2010] [error] [client 9.37.243.128] [9e44c28] [731] SSL0208E: SSL Handshake Failed,
Certificate validation error. [9.37.243.128:60347 -> 9.37.243.67:443] [08:54:25.000223331]
[Tue Jun 08 08:54:25 2010] [error] [client 9.37.243.128] [9e44c28] [731] Certificate validation error
during handshake, last PKIX/RFC3280 certificate validation error was
GSKVAL_ERROR_CA_MISSING_CRITICAL_BASIC_CONSTRAINT
[9.37.243.128:60347 -> 9.37.243.67:443] [08:54:25.000223331]
```

Example configuration:

```
<VirtualHost *:443
SSLClientAuth Optional Noverify
<Location />
SSLClientAuthVerify 419
</Location>
ErrorDocument 419 /error419.html
<Location /error419.html>
SSLClientAuthVerify OFF
</Location>
</VirtualHost>
```

SSLCRLHostname directive

The `SSLCRLHostname` directive specifies the TCP/IP name or address of LDAP server where the Certificate Revocation List (CRL) database resides.

Syntax	<code><SSLCRLHostName <TCP/IP name or address></code>
Scope	Global server or virtual host.
Default	Disabled by default.
Module	<code>mod_ibm_ssl</code>
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	TCP/IP name or address of the LDAP Server

Use the `SSLCRLHostname` directive, along with `SSLCRLPort`, `SSLCRLUserID`, and `SSLStashfile` directives, for static configuration of an LDAP-based CRL repository. It is only necessary to use these directives to query the LDAP-based CRL repository if an explicit `CRLDistributionPoint X.509v3` certificate extension is absent or the server specified in the extension is unresponsive (unavailable).

If a CRLDistributionPoint extension is present in the certificate and the server specified in the extension is responsive (available), then the LDAP server specified in the CRLDistributionPoint is queried anonymously, without using these directives.

SSLCRLPort directive

The SSLCRLPort directive specifies the port of the LDAP server where the Certificate Revocation List (CRL) database resides.

Syntax	SSLCRL<port>
Scope	Global server or virtual host.
Default	Disabled by default.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	Port of LDAP server; default = 389.

Use the SSLCRLPort directive, along with SSLCRLUserID, SSLCRLHostname, and SSLStashfile directives, for static configuration of an LDAP-based CRL repository. It is only necessary to use these directives to query the LDAP-based CRL repository if an explicit CRLDistributionPoint X.509v3 certificate extension is absent or the server specified in the extension is unresponsive (unavailable).

If a CRLDistributionPoint extension is present in the certificate and the server specified in the extension is responsive (available), then the LDAP server specified in the CRLDistributionPoint is queried anonymously, without using these directives.

SSLCRLUserID directive

The SSLCRLUserID directive specifies the user ID to send to the LDAP server, where the Certificate Revocation List (CRL) database resides.

Syntax	SSLCRLUserID <[prompt] <userid>
Scope	Global server or virtual host.
Default	Defaults to anonymous if you do not specify a user ID.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	User ID of LDAP server. Use the prompt option to enable the HTTP server to prompt you for the password to access the LDAP server during start up.

Use the SSLCRLUserID directive, along with SSLCRLPort, SSLCRLHostname, and SSLStashfile directives, for static configuration of an LDAP-based CRL repository. It is only necessary to use these directives to query the LDAP-based CRL repository if an explicit CRLDistributionPoint X.509v3 certificate extension is absent or the server specified in the extension is unresponsive (unavailable).

If a CRLDistributionPoint extension is present in the certificate and the server specified in the extension is responsive (available), then the LDAP server specified in the CRLDistributionPoint is queried anonymously, without using these directives.

SSLDisable directive

The SSLDisable directive disables SSL for the virtual host.

Syntax	SSLDisable
Scope	Global server or virtual host.
Default	Disabled by default.

Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	None.

SSLEnable directive

The SSLEnable directive enables SSL for the virtual host.

Attention: This directive should not be specified in the base server configuration if you do not want the directive automatically copied to a given virtual host configuration.

Syntax	SSLEnable
Scope	Global server or virtual host.
Default	Disabled by default.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	None.

SSLFakeBasicAuth directive

The SSLFakeBasicAuth directive enables the fake basic authentication support.

This support enables the client certificate distinguished name to become the user portion of the user and password basic authentication pair. Use **password** for the password.

Attention: This directive might be overridden by the base server configuration.

Syntax	SSLFakeBasicAuth
Scope	Within a directory stanza, used along with AuthName, AuthType, and require directives.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per directory stanza.
Values	None.

Distributed operating systems

SSLFIPSDisable directive

The SSLFIPSDisable directive disables Federal Information Processing Standards (FIPS).

Syntax	SSLFIPSDisable
Scope	Virtual and global.
Default	Disabled by default.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	None.

Distributed operating systems

SSLFIPSEnable directive

The SSLFIPSEnable directive enables Federal Information Processing Standards (FIPS).

Syntax	SSLFIPSEnable
---------------	---------------

Scope	Virtual and global.
Default	Disabled by default.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	None.

Attention: See the SSL cipher specification topic for values.

Distributed operating systems

SSLPKCSDriver directive

The SSLPKCSDriver directive identifies the fully qualified name to the module, or driver used to access the PKCS11 device.

Syntax	<i>Fully qualified name to module used to access PKCS11 device>.</i> If the module exists in the user path, then specify just the name of the module.
Scope	Global server or virtual host.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	Path and name of PKCS11 module or driver.

The default locations of the modules for each PKCS11 device follow:

- nCipher
 - AIX: /opt/nfast/toolkits/pkcs11/libcknfast.so
 - HP: /opt/nfast/toolkits/pkcs11/libcknfast.sl
 - Solaris: /opt/nfast/toolkits/pkcs11/libcknfast.so
 - Windows: c:\nfast\toolkits\pkcs11\cknfast.dll
- IBM 4758
 - AIX: /usr/lib/pkcs11/PKCS11_API.so
 - Windows: \$PKCS11_HOME\bin\nt\cryptoki.dll
- IBM e-business Cryptographic Accelerator
 - AIX: /usr/lib/pkcs11/PKCS11_API.so

SSLProtocolDisable directive

The SSLProtocolDisable directive enables you to specify one or more SSL protocols which cannot be used by the client for a specific virtual host. This directive must be located in a <VirtualHost> container.

Supported protocols for a virtual host are supported separately. If all supported protocols are disabled, clients cannot complete an SSL handshake.

Syntax	SSLProtocolDisable <protocolname>
Scope	Virtual host
Default	Disabled

Attention: The SSL Version 2 protocol is disabled by default with other methods.

Module	mod_ibm_ssl
Multiple instances in the configuration file	Multiple instances permitted per virtual host.

Values

The following possible values are available for this directive.

SSLv2

SSLv3

TLS

TLSv1

Distributed operating systems TLSv1.1

Distributed operating systems TLSv1.2

A value of TLS disables all TLS versions.

A value of TLSv1 disables TLS Version 1.0.

Distributed operating systems A value of TLSv1.1 disables TLS version 1.1.

Distributed operating systems A value of TLSv1.2 disables TLS version 1.2.

The following example disables support for multiple protocols on a virtual host.

```
<VirtualHost *:443> SSLEnable SSLProtocolDisable SSLv2
SSLv3 (any other directives) </VirtualHost>
```

Attention: SSL0230I is logged for each SSL connection attempt if the client and server do not share at least one protocol and cipher combination.

SSLProxyEngine directive

The SSLProxyEngine toggles whether the server uses SSL for proxied connections. SSLProxyEngine *on* is required if your server is acting as a reverse proxy for an SSL resource.

Syntax	SSLProxyEngine <i>on off</i>
Scope	IP-based virtual hosts
Default	Off
Module	mod_ibm_ssl
Multiple instances in the configuration file	One per virtual host and global server
Values	<i>on off</i>

SSLInsecureRenegotiation directive

The SSLInsecureRenegotiation directive determines whether insecure (pre RFC5746) SSL renegotiation is permitted. SSL Renegotiation of any kind is not common, and this directive should not be changed from its default value of *off*.

Distributed operating systems

Attention: For distributed platforms, RFC5746 support is automatically enabled.

When *on* is specified, insecure SSL renegotiation is permitted. When *off* is specified (the default), insecure SSL renegotiation is not permitted.

Syntax	SSLInsecureRenegotiation directive <i>on off</i>
Scope	Virtual hosts
Default	<i>off</i>

Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server
Values	on off

SSLServerCert directive

The SSLServerCert directive sets the server certificate to use for this virtual host.

Syntax	SSLServerCert [prompt] <i>my_certificate_label</i> ; on PKCS11 device - SSLServerCert <i>mytokenlabel:mykeylabel</i>
Scope	IP-based virtual hosts.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	Certificate label. Use the /prompt option to enable the HTTP server to prompt you for the Crypto token password during start up. Use no delimiters around the certificate label. Ensure that the label is contained on one line; leading and trailing white space is ignored.

SSLStashfile directive

The SSLStashfile directive indicates path to file with file name containing the encrypted password for opening the PKCS11 device.

Syntax	SSLStashFile /usr/HTTPServer/mystashfile.sth
Scope	Virtual host and global server.
Default	None.
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	File name of an LDAP and/or PKCS11 stash file that is created with the sslstash command.

The SSLStashFile does not point to a stash file for the KeyFile in use, as that is calculated automatically based on the name of the KeyFile, and is a different type of stashfile.

Use the **sslstash** command, located in the bin directory of IBM HTTP Server, to create your CRL password stash file. The password you specify using the **sslstash** command should equal the one you use to log in to your LDAP server.

Usage: sslstash [-c] <directory_to_password_file_and_file_name> <function_name> <password>

where:

- **-c**: Creates a new stash file. If not specified, an existing file updates.
- **File**: Represents the fully qualified name of the file to create, or update.
- **Function**: Indicates the function for which to use the password. Valid values include `crl`, or `crypto`.
- **Password**: Represents the password to stash.

Attention: See also “SSL certificate revocation list” on page 94.

Use the SSLStashFile directive, along with SSLCRLPort, SSLCRLHostname, and SSLCRLUserID directives, for static configuration of an LDAP-based CRL repository. It is only necessary to use these

directives to query the LDAP-based CRL repository if an explicit CRLDistributionPoint X.509v3 certificate extension is absent or the server specified in the extension is unresponsive (unavailable).

If a CRLDistributionPoint extension is present in the certificate and the server specified in the extension is responsive (available), then the LDAP server specified in the CRLDistributionPoint is queried anonymously, without using these directives.

SSLTrace directive

The SSLTrace directive enables debug logging in mod_ibm_ssl. It is used in conjunction with the LogLevel directive. To enable debug logging in mod_ibm_ssl, set LogLevel to debug and add the SSLTrace directive to global scope in the IBM HTTP Server configuration file, after the LoadModule directive for mod_ibm_ssl. This directive is typically used at the request of IBM support while investigating a suspected problem with mod_ibm_ssl. It is not recommended to enable this directive under normal working conditions.

Syntax	SSLTrace
Scope	Global
Default	mod_ibm_ssl debug logging in not enabled
Module	mod_ibm_ssl
Multiple instances in the configuration file	Ignored
Values	None

Attention: See also LogLevel Directive.

Distributed operating systems

SSLUnknownRevocationStatus

The SSLUnknownRevocationStatus directive specifies how IBM HTTP Server reacts when IBM HTTP Server cannot readily determine the revocation status, which is coming through CRL or OCSP.

Syntax	SSLUnknownRevocationStatus ignore log log_always deny
Scope	Virtual host
Default	ignore
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance permitted for each virtual host
Values	<p>ignore Specifies that a debug level message is issued when a handshake completes and the revocation status is not known. This message is not re-issued when the SSL session is resumed.</p> <p>log Specifies that a notice-level message is issued when a handshake completes and the revocation status is not known. This message is not re-issued when the SSL session is resumed.</p> <p>log_always Specifies that a notice-level message is issued when a handshake completes and the revocation status is not known. IBM HTTP Server issues the same message for subsequent handshakes.</p> <p>deny Specifies that a notice-level message is issued when a handshake completes, the revocation status is not known, the session is not resumable, and the HTTPS connection is immediately closed. IBM HTTP Server reports the same message for subsequent handshakes.</p>

config: Whenever a message is logged for UnknownRevocationStatus, the SSL_UNKNOWNREVOCACTION_SUBJECT variable, an internal SSL environment variable, is set. You can log this variable with the following syntax:

```
%{SSL_UNKNOWNREVOCACTION_SUBJECT}e
```

You could also use the variable in mod_rewrite expressions when the SSLUnknownRevocationStatus directive has any value other than deny. Use the following variable name:

```
%{ENV:SSL_UNKNOWNREVOCACTION_SUBJECT}
```

SSLV2Timeout directive

The SSLV2Timeout directive sets the timeout for SSL Version 2 session IDs.

Syntax	SSLV2Timeout 60
Scope	Global base and virtual host.
Default	40
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	0 to 100 seconds.

SSLV3Timeout directive

The SSLV3Timeout directive sets the timeout for SSL Version 3 and TLS session IDs.

Syntax	SSLV3Timeout 1000
Scope	Global base and virtual host.

Windows The virtual host scope or global scope are applicable.

AIX **HP-UX** **Linux** **Solaris** The virtual host scope is applicable if the SSLCacheDisable directive is also being used. Otherwise, only the global scope is allowed.

Default	120
Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per virtual host and global server.
Values	0 to 86400 seconds.

SSLVersion directive

The SSLVersion directive causes object access rejection with a 403 response if the client has connected with an SSL protocol version other than the one specified.

In most cases, the SSLProtocolDisable directive is a better choice than the SSLVersion directive for ensuring use of particular SSL protocol versions. The SSLProtocolDisable directive enables the client browser to negotiate another protocol version if possible whereas the SSLVersion directive causes IBM HTTP Server to send a 403 response, which might confuse the user.

Syntax	SSLVersion ALL
Scope	One per directory stanza.
Default	None.

Module	mod_ibm_ssl
Multiple instances in the configuration file	One instance per <Directory> or <Location> stanza.
Values	<div style="display: flex; justify-content: space-between;"> Distributed operating systems z/OS </div> SSLV2 SSLV3 TLS TL SV1 TL SV11 TL SV12 SSLV2 SSLV3 TLS TL SV1 ALL

Setting advanced SSL options

You can enable advanced security options such as: client authentication, setting and viewing cipher specifications, defining SSL for multiple-IP virtual hosts, and setting up a reverse proxy configuration with SSL.

About this task

After setting up secure connections, follow these instructions to enable advanced security options:

Procedure

1. Enable client authentication. If you enable client authentication, the server validates clients by checking for trusted certificate authority (CA) root certificates in the local key database.
2. Set and view cipher specifications.

Important: If you specify V3 or TLS ciphers and no SSL V2 ciphers, SSL V2 support is disabled. Also, if you specify SSL V2 ciphers and no SSL V3 or TLS ciphers, SSL V3 and TLS support is disabled.

3. Define Secure Sockets Layer (SSL) for multiple-IP virtual hosts.

Choosing the level of client authentication

If you enable client authentication, the server validates clients by requesting a certificate from the client and verifying that is signed by a trusted certificate authority (CA) root certificate in the server key database.

About this task

For each virtual host, choose the level of client authentication:

Procedure

1. Specify one of the following values in the configuration file on the SSLClientAuth directive, for each virtual host stanza . A virtual host stanza represents a section of the configuration file that applies to one virtual host.

Table 11. Client authentication level. The table lists the value for the client authentication level and a description of the value

Value	Description
None	The server requests no client certificate from the client.
Optional	The server requests, but does not require, a client certificate. If presented, the client certificate must prove valid.
Required	The server requires a valid certificate from all clients, returning a 403 status code if no certificate is present.

Table 11. Client authentication level (continued). The table lists the value for the client authentication level and a description of the value

Value	Description
Required_reset	The server requires a valid certificate from all clients, and if no certificate is available, the server sends an SSL alert to the client. This enables the client to understand that the SSL failure is client-certificate related, and will cause browsers to re-prompt for client certificate information on subsequent access.

For example, SSLClientAuth required.

If you want to use a certificate revocation list (CRL), add `cr1`, as a second argument for SSLClientAuth. For example: SSLClientAuth required cr1.

2. Save the configuration file and restart the server.

Choosing the type of client authentication protection

If you enable client authentication, the server validates clients by checking for trusted certificate authority (CA) root certificates in the local key database.

About this task

For each virtual host, choose the type of client authentication:

Procedure

1. Specify one of the following directives in the configuration file, for each virtual host stanza:
 - a. SSLClientAuthRequire. For example, SSLClientAuthRequire CommonName=Richard
 - b. SSLFakeBasicAuth. If you specify SSLFakeBasicAuth, verify that the `mod_ibm_ssl` module is displayed last in the module list.
2. Save the configuration file and restart the server.

Viewing cipher specifications

This section describes viewing cipher specifications for secure transactions and for a specific HTTP request.

About this task

To see which cipher specifications the server uses for secure transactions or for a specific HTTP request, complete one of the following steps.

Procedure

1. **To see which cipher specifications the server uses for secure transactions.** Specify `LogLevel info` in the configuration file to include informational messages in the error log using the `LogLevel` directive. The error log is specified by the `ErrorLog` directive in the `http` configuration file. The location is set by the `ErrorLog` directive, which can be configured. Look in the error log for messages in this format: *TimeStamp info_message mod_ibm_ssl: Using Version 2/3 Cipher:longname/shortname*. The order that the cipher specifications are displayed in the error log from top to bottom represents the attempted order of the cipher specifications.
2. **To see which cipher specification was negotiated with a specific client for a specific request.** Change the `LogFormat` directive to include the cipher specification as part of the information logged for each request. The format string `%{HTTPS_CIPHER}e` will log the name of the cipher (for example, "TLS_RSA_WITH_AES_256_CBC_SHA"). Be sure that the `LogFormat` directive you change is for the format used on the `CustomLog` directive. Here is an example:

```
LogFormat "%h %l %u %t \"%r\" %>s %b %{HTTPS_CIPHER}e" common
CustomLog logs/access_log common
```

Look in the access log to find the cipher used. The position of the cipher will depend on where the `%{HTTPS_CIPHER}e` format string was placed in the `LogFormat` directive. Following are some example `access_log` entries, using the example `LogFormat` directive above:

```
9.48.108.152 - - [17/Feb/2005:15:37:39 -0500]
"GET / HTTP/1.1" 200 1507 SSL_RSA_WITH_RC4_128_SHA

9.48.108.152 - - [17/Feb/2005:15:37:40 -0500]
"GET /httpTech.view1.gif HTTP/1.1" 200 1814 SSL_RSA_WITH_RC4_128_SHA

9.48.108.152 - - [17/Feb/2005:15:37:40 -0500]
"GET /httpTech.masthead.gif HTTP/1.1" 200 11844 SSL_RSA_WITH_RC4_128_SHA

9.48.108.152 - - [17/Feb/2005:15:37:41 -0500]
"GET /httpTech.visit1.gif HTTP/1.1" 200 1457 SSL_RSA_WITH_RC4_128_SHA
```

For non-secure requests, "-" will be logged for the cipher specification. You can log other SSL environment variables in the same manner as `HTTPS_CIPHER`.

SSL cipher specifications

When an SSL connection is established, the client (web browser) and the web server negotiate the cipher to use for the connection. The web server has an ordered list of ciphers, and the first cipher in the list that is supported by the client is selected.

Introduction

View the list of current of SSL ciphers.

Attention: This list of ciphers could change as a result of updates to industry standards. You can determine the list of ciphers supported in a particular version of IBM HTTP Server by configuring it to load `mod_ibm_ssl` and running `bin/apachectl -t -f path/to/httpd.conf -DDUMP_SSL_CIPHERS`.

The `SSLFIPSEnable` directive enables Federal Information Processing Standards (FIPS). When the `SSLFIPSEnable` directive is enabled, the set of ciphers available is restricted as shown, and SSLv2 and SSLv3 are disabled.

SSL and TLS ciphers

Attention: Note the following SSL and TLS cipher values:

- - = cipher that is not valid for the protocol
- d = cipher is enabled by default
- y = cipher is valid but not enabled by default

Attention: TLS v1.1 and v1.2 are not available on the z/OS operating system.

Note: To improve security, IBM HTTP Server Version 8.0 disables weak SSL ciphers, export SSL ciphers, and the SSL version 2 protocol by default. SSL Version 2, weak ciphers, and export ciphers are generally unsuitable for production SSL workloads on the internet and are flagged by security scanners. To enable ciphers, use the `SSLCipherSpec` directive.

Table 12. SSL and TLS ciphers

Short name	Long name	Key size (bits)	FIPS	SSLV2	SSLV3	TLSv10	TLSv11	TLSv12

Table 12. SSL and TLS ciphers (continued)

2F	TLS_RSA_ WITH_ AES_128_ CBC_SHA	128	Y	-	d	d	d	d
34	SSL_ RSA_ WITH_ RC4_128_ MD5	128	-	-	d	d	d	-
35	SSL_ RSA_ WITH_ RC4_128_ SHA	128	-	-	d	d	d	y
35b	TLS_ RSA_ WITH_ AES_256_ CBC_SHA	256	Y	-	d	d	d	d
3A	SSL_ RSA_ WITH_ 3DES_ EDE_ CBC_SHA	168	Y	-	d	d	d	d
3C	TLS_ RSA_ WITH_ AES_128_ CBC_ SHA256	128	Y	-	-	-	-	d
3D	TLS_ RSA_ WITH_ AES_256_ CBC_ SHA256	256	Y	-	-	-	-	d
9C	TLS_ RSA_ WITH_ AES_128_ GCM_ SHA256	128	Y	-	-	-	-	d
9D	TLS_ RSA_ WITH_ AES_256_ GCM_ SHA384	256	Y	-	-	-	-	d

Weaker ciphers, not enabled by default:

Table 13. Ciphers that are not enabled by default

Short name	Long name	Key size (bits)	FIPS	SSLV2	SSLV3	TLSv10	TLSv11	TLSv12
21	SSL_RC4_128_WITH_MD5	128	-	y	-	-	-	-
22	SSL_RC4_128_EXPORT_40_WITH_MD5	40	-	y	-	-	-	-
23	SSL_RC2_CBC_128_WITH_MD5	128	-	y	-	-	-	-
24	SSL_RC2_CBC_128_EXPORT_40_WITH_MD5	40	-	y	-	-	-	-
26	SSL_DES_64_CBC_WITH_MD5	56	-	y	-	-	-	-
27	SSL_DES_192_EDE3_CBC_WITH_MD5	168	-	y	-	-	-	-
30	SSL_NULL_WITH_NULL_NULL	0	-	-	y	y	y	y
31	SSL_RSA_WITH_NULL_MD5	0	-	-	y	y	y	-
32	SSL_RSA_WITH_NULL_SHA	0	-	-	y	y	y	y

Table 13. Ciphers that are not enabled by default (continued)

33	SSL_ RSA_ EXPORT_ WITH_ RC4_ 40_MD5	40	-	-	y	y	-	-
36	SSL_ RSA_ EXPORT_ WITH_ RC2_ CBC_ 40_MD5	40	-	-	y	y	-	-
39	SSL_ RSA_ WITH_ DES_ CBC_ SHA	56	-	-	y	y	y	-
3B	TLS_ RSA_ WITH_ NULL_ SHA256	0	Y	-	-	-	-	y
62	TLS_ RSA_ EXPORT 1024_ WITH_ DES_ CBC_ SHA	56	-	-	y	y	-	-
64	TLS_ RSA_ EXPORT 1024_ WITH_ RC4_ 56_ SHA	56	-	-	y	y	-	-
FE	SSL_ RSA_ FIPS_ WITH_ DES_ CBC_ SHA	56	-	-	-	-	-	-
FF	SSL_ RSA_ FIPS_ WITH_ 3DES_ EDE_ CBC_ SHA	168	-	-	-	-	-	-

Defining SSL for multiple-IP virtual hosts

You can define different Secure Sockets Layer (SSL) options for various virtual hosts, or multiple servers running on one machine. In the configuration file, define each SSL directive in the stanza for the virtual host to which the directive applies. When you do not define an SSL directive on a virtual host, the server uses the directive default.

About this task

The default disables SSL for each virtual host. To enable SSL:

Procedure

1. Specify the `SSLEnable` directive on the virtual host stanza in the configuration file, to enable SSL for a virtual host.
2. Specify a `Keyfile` directive and any SSL directives you want enabled for that particular virtual host. You can specify any directive, except the cache directives inside a virtual host.
3. Restart the server.

Setting up a reverse proxy configuration with SSL

This topic describes how to set up a site to act as a reverse proxy for a resource that is hosted on a secure site.

About this task

The following steps describe how to set up a reverse proxy configuration for a company (for example, `www.example.com`) which wants to act as a reverse proxy for a resource that is hosted on a secure site (for example, `internal.example.com`).

Procedure

1. Configure `www.example.com` similar to the following example:

```
<VirtualHost *:80>
  ServerName host1
  SSLProxyEngine On
  KeyFile "c:/program files/ibm http server/clientkey.kdb"
  ProxyPass /ssl/password.html https://examplehost/password.html
</VirtualHost>
```

2. Configure `internal.example.com` similar to the following example:

```
<VirtualHost *:443>
  SSLEnable
  KeyFile "c:/program files/ibm http server/serverkey.kdb"
</VirtualHost>
```

Results

When a browser requests `http://www.example.com/ssl/password.html`, IBM HTTP Server makes a connection to `internal.ibm.com` using SSL. If `internal.example.com` requires a client certificate, IBM HTTP Server uses the default certificate of the `KeyFile` for which it is configured.

IBM HTTP Server certificate management

Before you can configure IBM HTTP Server to accept TLS (also known as SSL) connections, you must create a certificate for your web server. An SSL certificate authenticates your web servers identity to clients.

Background information and tools

The primary tool for creating certificates with IBM HTTP Server is iKeyman, a graphical pure Java key management tool.

z/OS On z/OS operating systems, all certificate management is done with the native gskkyman certificate management tool.

Distributed operating systems On Microsoft Windows, you can start iKeyman using the Start Menu. On other platforms, start the tool from the IBM HTTP Server bin/ directory, like all IBM HTTP Server executable files.

Native and Java supplemental command-line certificate management tools are also provided in the IBM HTTP Server bin/ directory as gskcmd (also known as iKeycmd) and gskcapicmd (also known as gsk8capicmd). Both share similar syntax and contain extensive embedded usage information.

Certificate limitations in IBM HTTP Server

- Only RSA certificates (keys) are supported with IBM HTTP Server. DSA and ECC certificates are not supported.
- Certificates with a key length of up to 4096 bits are supported at run time with IBM HTTP Server.
- Ikeyman and gskcmd (ikeycmd) support creating certificates of lengths up to 2048 bits. The gskcapicmd command supports creating certificates of lengths up to 4096 bits.
- Multiple key database files can be used with each instance of IBM HTTP Server, but only one, which can still contain multiple personal certificates, can be used per TLS-enabled virtual host.

Complete documentation for certificate management tools

- Complete documentation of gskkyman is available in the "Cryptographic Services PKI Services Guide and Reference document" (SA22-7693) in the z/OS Internet Library.
- Complete documentation for iKeyman and gskcmd (ikeycmd) are available in the iKeyman v8 Users Guide.
- Complete documentation for gskcapicmd (gsk8capicmd), the native command-line certificate management tool, is available on the IBM HTTP Server library page.

System setup

- Unlike prior releases of IBM HTTP Server, do not move or modify the java/jre/lib/ext/gskikm.jar file.
- Optionally install the Unrestricted JCE policy files from DeveloperWorks to use unlimited strength cryptography in iKeyman and gskcmd (ikeycmd). This step is often required to manipulate PKCS12 keystores.

z/OS Certificate management tasks

Detailed example scenarios for certificate management are documented in the complete documentation for iKeyman (distributed operating systems) and gskkyman (z/OS operating systems).

Distributed operating systems Certificate management tasks

Detailed example scenarios for certificate management are documented in the complete documentation for iKeyman (distributed operating systems) and gskkyman (z/OS operating systems).

Distributed operating systems See the following command-line examples of common tasks. You can view full usage syntax by entering the following commands with only the first two parameters, or you can refer to the

comprehensive documentation for the command. The following table lists the operations that you can perform on CA certificates, the AdminTask object that you can use to perform that operation, and how to navigate to the certificate on the console:

Create a CMS keystore

When creating a keystore to be used with IBM HTTP Server, specify the option to stash the password to a file regardless of the tool used.

```
# Syntax: <ihsroot>/bin/gskcapicmd -keydb -create -db <database> -pw <password> -stash
<ihsroot>/bin/gskcapicmd -keydb -create -db /opt/IBM/HTTPServer/key.kdb -pw password -stash
```

Populate a keystore with a set of default trusted CA certificates

By default, new keystores contain no trusted CA certificates.

```
# The populate operation is supported with Ikeyman and gskcmd (ikeycmd) only, not with gskcapicmd.
# Syntax: <ihsroot>/bin/gskcmd -cert -populate -db <database> -pw <password>
<ihsroot>/bin/gskcmd -cert -populate -db /opt/IBM/HTTPServer/key.kdb -pw password
```

Add additional CA certificates, if wanted (optional)

```
# Syntax: <ihsroot>/bin/gskcapicmd -cert -add -db <database> -pw <password> -file <inputcert> -label <labelname>
<ihsroot>/bin/gskcapicmd -cert -add -db /opt/IBM/HTTPServer/key.kdb -pw password -file cacert.cer -label "CA certificate from example.com"
```

Create a self-signed certificate for test purposes (optional)

```
#Syntax: <ihsroot>/bin/gskcapicmd -cert -create -db <database> -pw <password> \
      -dn <distinguished name> -label <labelname> -size <size>
<ihsroot>/bin/gskcapicmd -cert -create -db /opt/IBM/HTTPServer/key.kdb -pw password \
      -dn "cn=www.example.com" -label "example.com" -size 2048
```

Create a certificate request

Most of the fields and options are optional, including selecting a Signature Algorithm (this signature is used only by your certificate authority, not at runtime). You can also specify other host names for your web server.

```
# Syntax: <ihsroot>/bin/gskcapicmd -certreq -create -db <database> -pw <password> \
      -dn <distinguished name> -label <labelname> -size <size> -file <outputfilename>
<ihsroot>/bin/gskcapicmd -certreq -create -db/opt/IBM/HTTPServer/key.kdb -pw password \
      -dn "cn=www.example.com" -label www.example.com -size 2048 -file example.csr
```

Submit the certificate request to a trusted certificate authority

This task does not include using any local tools. Typically, the certificate request (example.csr) is sent in an email or uploaded to a trusted certificate authority.

Receive the issued certificate

Receiving a certificate associates a signed certificate from your CA with the private key (personal certificate) in your KDB file. A certificate can only be received into the KDB that generated the certificate request.

```
# Syntax: <ihsroot>/bin/gskcapicmd -cert -receive -db <database> -pw <password> -file <inputcertificate>
<ihsroot>/bin/gskcapicmd -cert -receive -db/opt/IBM/HTTPServer/key.kdb -pw password -file certificate.arm
```

List certificates in a keystore.

```
# Syntax <ihsroot>/bin/gskcapicmd -cert -list -db <database> -pw <password>
<ihsroot>/bin/gskcapicmd -cert -list -db /opt/IBM/HTTPServer/key.kdb -pw password
```

Import certificates from JKS or PKCS12 into a key file usable by IBM HTTP Server (optional)

Instead of creating a new private key (personal certificate), you can import an existing private key and certificate created by another tool into an existing key file.

```
# Syntax: <ihsroot>/bin/gskcapicmd -cert -import -db <inputp12file> -pw <pkcs12password>\
        -target <existingkdbfile> -target_pw <existingkdbpassword>
<ihsroot>/bin/gskcapicmd -cert -import -db other.p12 -pw pkcs12password \
        -target key.kdb -target_pw password
```

View certificate expiration data (optional)

The `-expiry` flag causes certificates that will be considered expired "numdays" in the future to be displayed. Use "0" to display already expired certificates, or large numbers to display all certificate expiration dates.

```
# Syntax:<ihsroot>/bin/gskcapicmd -cert -list -db <database> -pw <password> -expiry <numdays>
<ihsroot>/bin/gskcapicmd -cert -list -db key.kdb -password -expiry 365
```

Managing keys with the IKEYMAN graphical interface (Distributed systems)

This section describes topics on how to set up and use the key management utility (IKEYMAN) with IBM HTTP Server. Using the graphical user interface, rather than the command line interface, is recommended.

Linux

Before you begin

Ensure that the required `compat-libstdc++` package exists for your operating system architecture. For more information, see the installation and verification information for Linux packages.

About this task

Global Security Kit (GSKit) certificate management tools are installed in the `<ihsinst>/bin/` directory. These tools should only be run from the installation directory. Examples for the following commands should include the full directory path, such as `<ihsinst>/bin/gskcmd`.

- `gskver`
- `ikeyman`
- `gskcapicmd`
- `gskcmd`

For IKEYMAN, you can run the following command in the installation directory to generate debug information.

```
<ihsinst>/bin/ikeyman -x
```

To have a secure network connection, create a key for secure network communications and receive a certificate from a certificate authority (CA) that is designated as a trusted CA on your server.

Procedure

- Start the Key Management utility user interface. Use IKEYMAN to create key databases, public and private key pairs, and certificate requests.
- Work with key databases. You can use one key database for all your key pairs and certificates, or create multiple databases.
- Change the database password. When you create a new key database, you specify a key database password, which protects the private key. The private key is the only key that can sign documents or decrypt messages that are encrypted with the public key. Changing the key database password frequently is a good practice.
- Create a new key pair and certificate request. You find key pairs and certificate requests stored in a key database.
- Import and export your key into another database or to a PKCS12 file. PKCS12 is a standard for securely storing private keys and certificates.
- List certificate authorities within a key database.

- Display certificate expiration date your key database by viewing the certificate information with the IKEYMAN Key Management utility GUI or using the gskcmd command.
- If you act as your own CA, you can use IKEYMAN to create self-signed certificates.
- Receive a signed certificate from a certificate authority. If you act as your own CA for a private Web network, you have the option to use the server CA utility to generate and issue signed certificates to clients and servers in your private network.
- Display default keys and certificate authorities within a key database.
- Store a certificate from a certificate authority (CA) that is not a trusted CA.
- Store the encrypted database password in a stash file.
- Use IKEYMAN to create key databases, public and private key pairs, and certificate requests.
- If you act as your own CA, you can use IKEYMAN to create self-signed certificates.
- If you act as your own CA for a private Web network, you have the option to use the server CA utility to generate and issue signed certificates to clients and servers in your private network.

What to do next

You may experience a certificate problem when you open a certificate that has a key with a higher level of cryptography than your policy files permit. You can optionally install unlimited strength JCE policy files.

- Download and install the files from the following Web site. <https://www14.software.ibm.com/webapp/iwm/web/preLogin.do?source=jcesdk>.
- Find the maximum key sizes permitted by key type with the default policy files at the following Web site. <http://java.sun.com/j2se/1.4.2/docs/guide/security/jce/JCERefGuide.html#AppE>.

For more information about the IKEYMAN utility, see the IKEYMAN User's Guide on the IHS Library page.

Starting the Key Management utility user interface

This section describes how to start the Key Management (IKEYMAN) utility.

Procedure

- From a command line:

```
<install_root>/bin/ikeyman
```

or change to the `<install_root>/bin` directory and type `ikeyman`

- On Windows operating systems: Click **Start > Programs > IBM HTTP Server > Start Key Management Utility**. If you start IKEYMAN to create a new key database file, the utility stores the file in the directory where you start IKEYMAN.

Working with key databases

This article describes how to create a new key database and open an existing key database.

About this task

A *key database* is a file that the server uses to store one or more key pairs and certificates. You can use one key database for all your key pairs and certificates, or create multiple databases.

Procedure

- Create a new key database as follows:
 1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
 2. Click **key database file** from the main user interface, then click **New**. Select **CMS** for the Key database type. IBM HTTP Server does not support database types other than CMS.

3. Enter your password in the Password Prompt dialog box, and confirm the password. Select **Stash the password to a file**. Click **OK**. The new key database should display in the IKEYMAN utility with default signer certificates. Ensure that there is a functional, non-expiring signer certificate for each of your personal certificates. .
- Open an existing key database as follows:
 1. Start the IKEYMAN user interface.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. In the Open dialog box, enter your key database name, or click the **key.kdb** file, if you use the default. Click **OK**.
 4. Enter your correct password in the Password Prompt dialog box, and click **OK**.
 5. The key database name is displayed in the File Name text box.

What to do next

You can add a default list of signer certificates to your new database by following the instructions below. The version of iKeyman that is provided by the bundled Java Runtime Environment (JRE) does not add a default list of signer certificates to newly-created key databases. Add default signer certificates in iKeyman, as follows:

1. Select **Signer Certificates** from the drop-down menu in the iKeyman window.
2. Click the "Populate" on the right-hand side of the iKeyman window.
3. Click the grey boxes next to the certificate authority names (Entrust, RSA Data Security, Thawte, Verisign) so they display as checked.
4. Click **OK**.

Changing the database password

When you create a new key database, you specify a key database password, which protects the private key. The private key is the only key that can sign documents or decrypt messages that are encrypted with the public key. Changing the key database password frequently is a good practice.

About this task

Complete the following steps to change the database password:

Procedure

1. Start the IKEYMAN user interface.
2. Click **Key Database File** from the main UI, then click **Open**.
3. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if you use the default. Click **OK**.
4. Enter your password in the Password Prompt dialog box, and click **OK**.
5. Click **Key Database File** from the main UI, then click **Change Password**.
6. Enter a new password in the Password Prompt dialog box, and a new confirming password. Click **OK**.

Use the following guidelines when specifying the password:

- The password must come from the U.S. English character set.
- The password must contain at least six characters and contain at least two nonconsecutive numbers. Make sure that the password does not consist of publicly obtainable information about you, such as the initials and birth date for you, your spouse, or children.
- Stash the password or enable secure sockets layer (SSL) password prompting.

Keep track of expiration dates for the password. If the password expires, a message writes to the error log. The server starts, but a secure network connection does not exist, if the password has expired.

Creating a new key pair and certificate request

You find key pairs and certificate requests stored in a key database. This section provides information on how to create a key pair and certificate request.

Before you begin

There are GSKit certificate support limitations that you should remember as you create a new key pair and certificate request:

- You cannot use IKEYMAN to create certificates with key sizes that are larger than 2048 bits.
- You can import certificates with key sizes up to 4096 bits into the key database.

About this task

To create a public and private key pair and certificate request, complete the following steps:

Procedure

1. If you have not created the key database, see [Creating a new key database](#) for instructions.
2. Start the IKEYMAN user interface.
3. Click **Key Database File** from the main user interface, then click **Open**.
4. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if you use the default. Click **OK**.
5. In the Password Prompt dialog box, enter your correct password and click **OK**.
6. Click **Create** from the main user interface, then click **New Certificate Request**.
7. In the New Key and Certificate Request dialog box, complete the following information:
 - Key label: Enter a descriptive comment to identify the key and certificate in the database.
 - Key size: Choose your level of encryptions from the drop-down menu.
 - Organization Name: Enter your organization name.
 - Organization Unit
 - Locality
 - State/Province
 - Zip code
 - Country: Enter a country code. Specify at least two characters. Example: US Certificate request file name, or use the default name.

A checksum of the certificate request is cryptographically signed with the new private key, and contains a copy of the new public key. The public key can then be used by a certificate authority to validate that the certificate signing request (CSR) has not been tampered with. Some certificate authorities might require that the checksum that is signed by the public key be calculated with a stronger algorithm such as SHA-1 or SHA-2 (SHA-256, SHA-384, SHA-512).

This checksum is the "Signature Algorithm" of the CSR

Subject Alternate Name (SAN) extensions are fields in a certificate request that inform SSL Clients of alternate hostnames that correspond to the signed certificate. Normal certificates (issued without a wildcard string in their Distinguished Name) are only valid for a single hostname. For example, a certificate created for example.com is not valid on www.example.com unless a Subject Alternate Name of "www.example.com" is added to the certificate. A certificate authority may charge an additional fee if your certificate contains 1 or more SAN extensions.

8. Click **OK**.
9. Click **OK** in the Information dialog box. A reminder to send the file to a certificate authority displays.
10. Optional: On UNIX-based platforms, remove the end of line characters (^M) from the certificate request. To remove the end of line characters, type the following command:

```
cat certreq.arm |tr -d "\r" > new_certreq.arm
```

11. Send the file to the certificate authority (CA) following the instructions from the CA Web site for requesting a new certificate.

Importing and exporting keys

This article describes how to import and export your key into another database or to a PKCS12 file. PKCS12 is a standard for securely storing private keys and certificates.

About this task

To import and export keys from another database, complete the following steps:

Procedure

- Import keys from another database by completing the following steps:
 1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. Enter your key database name in the Password prompt dialog box, or click **key.kdb** if you are using the default.
 4. Enter your correct password in the Password prompt dialog box, and click **OK**.
 5. Click **Personal Certificates** in the Key Database content frame, then click **Export/Import** on the label.
 6. In the Export/Import Key window:
 - a. Click **Import Key**.
 - b. Click the target database type.
 - c. Enter the file name, or use the Browse option.
 - d. Enter the current location.
 7. Click **OK**.
 8. Click **OK** in the Password prompt dialog box, to import the selected key to another key database.
- Import keys to a PKCS12 file by completing the following steps:
 1. Enter `ikeyman` on a command line on the Linux or UNIX platforms, or start the Key Management utility in the IBM HTTP Server folder on the Windows operating system.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click **key.kdb**, if you use the default. Click **OK**.
 4. Enter your password in the Password prompt dialog box, and click **OK**.
 5. Click **Personal Certificates** in the Key Database content frame, then click **Export/Import** on the label.
 6. In the Export/Import Key window:
 - a. Click **Import Key**.
 - b. Click the PKCS12 database file type.
 - c. Enter the file name, or use the Browse option.
 - d. Enter the correct location.
 7. Click **OK**.
 8. Enter the correct password in the Password prompt dialog box, then click **OK**.
- Export keys from another database by completing the following steps:
 1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
 2. Click **key database file** from the main user interface, then click **Open**.

3. Enter your key database name in the Password Prompt dialog box, or click **key.kdb** if you are using the default.
4. Enter your correct password in the Password Prompt dialog box, and click **OK**.
5. Click **Personal Certificates** in the Key database content frame, then click **Export/Import** on the label.
6. In the Export/Import Key window:
 - a. Click **Export Key**.
 - b. Click the target database type.
 - c. Enter the file name, or use the Browse option.
 - d. Enter the current location.
- Export keys to a PKCS12 file by completing the following steps:
 1. Enter `ikeyman` on a command line on the Linux or UNIX platforms, or start the Key Management utility in the IBM HTTP Server folder on the Windows operating system.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click **key.kdb** if you use the default. Click **OK**.
 4. Enter your password in the Password Prompt dialog box, and click **OK**.
 5. Click **Personal Certificates** in the Key Database content frame, then click **Export/Import** on the label.
 6. In the Export/Import Key window:
 - a. Click **ExportKeyM**.
 - b. Click the PKCS12 database file type.
 - c. Enter the file name, or use the Browse option.
 - d. Enter the correct location.
 7. Click **OK**.
 8. Enter the correct password in the Password prompt dialog box, and enter the password again to confirm. Click **OK** to export the selected key to a PKCS12 file.

Listing certificate authorities

You can display a list of trusted certificate authorities within a key database.

About this task

A trusted certificate authority issues and manages public keys for data encryption. A key database is used to share public keys that are used for secure connections.

To display a list of trusted certificate authorities (CAs) in a key database, complete the following steps:

Procedure

1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
2. Click **Key Database File** from the main UI, then click **Open**.
3. Enter your key database name in the Open dialog box, or click **key.kdb** if you are using the default.
4. Enter your correct password in the Password prompt dialog box, and click **OK**.
5. Click **Signer Certificates** in the Key database content frame.
6. Click **Signer Certificates**, **Personal Certificates**, or **Certificate Requests**, to view the list of CAs in the Key Information window.

What to do next

When the `<ihsinst>/java/jre/lib/ext/gskikm.jar` file has not been removed, the version of iKeyman that is provided by the bundled Java Runtime Environment (JRE) does not add a default list of signer certificates to newly-created key databases. Add default signer certificates in iKeyman, as follows:

1. Select **Signer Certificates** from the drop-down menu in the iKeyman window.
2. Click the "Populate" on the right-hand side of the iKeyman window.
3. Click the grey boxes next to the certificate authority names (Entrust, RSA Data Security, Thawte, Verisign) so they display as checked.
4. Click **OK**.

Certificate expiration dates

You can display expiration dates of certificates in your key database by viewing the certificate information with the IKEYMAN Key Management utility GUI or using the `gskcmd` command.

The following is an example of how to use the `gskcmd` command to display the validity dates on all certificates in the `key.kdb` certificate key file that will expire within 1825 days (5 years):

```
<ihsinst>/bin/gskcmd -cert -list all -expiry 1825 -db key.kdb -pw <password>
```

```
Certificates in database: key.kdb
VeriSign Class 1 CA Individual Subscriber-Persona Not Validated
Validity
Not Before: Mon May 11 20:00:00 EDT 1998
Not After: Mon May 12 19:59:59 EDT 2008
```

where `<password>` is the password you specified when creating the `key.kdb` key database file.

Creating a self-signed certificate

It usually takes two to three weeks to get a certificate from a well known certificate authority (CA). While waiting for a certificate to be issued, use IKEYMAN to create a self-signed server certificate to enable SSL sessions between clients and the server. Use this procedure if you act as your own CA for a private Web network.

About this task

Complete the following steps to create a self-signed certificate:

Procedure

1. If you have not created the key database, see [Creating a new key database](#) for instructions.
2. Start the IKEYMAN user interface.
3. Click **Key Database File** from the main UI, and then click **Open**.
4. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if you use the default. Click **OK**.
5. In the Password Prompt dialog box, enter your correct password and click **OK**.
6. Click **Personal Certificates** in the Key Database content frame, and click the **New Self-Signed** radio button.
7. Enter the following information in the Password Prompt dialog box:
 - Key label: Enter a descriptive comment to identify the key and certificate in the database.
 - Key size: Choose your level of encryptions from the drop-down menu.
 - Common Name: Enter the fully qualified host name of the Web server as the common name. Example: `www.myserver.com`.
 - Organization Name: Enter your organization name.

- **Optional:** Organization Unit
- **Optional:** Locality
- **Optional:** State/Province
- **Optional:** Zip code
- Country: Enter a country code. Specify at least two characters. Example: US Certificate request file name, or use the default name.
- Validity Period

A checksum of the certificate request is cryptographically signed with the new private key, and contains a copy of the new public key. The public key can then be used by a certificate authority to validate that the certificate signing request (CSR) has not been tampered with. Some certificate authorities might require that the checksum that is signed by the public key be calculated with a stronger algorithm such as SHA-1 or SHA-2 (SHA-256, SHA-384, SHA-512).

This checksum is the "Signature Algorithm" of the CSR.

Note: IBM HTTP Server 8.0 ships IKEYMAN version 8.x. When using IKEYMAN version 8.x to create a certificate request, the user is asked to select a signature algorithm from a pull-down list.

Subject Alternate Name (SAN) extensions are fields in a certificate request that inform SSL Clients of alternate hostnames that correspond to the signed certificate. Normal certificates (issued without a wildcard string in their Distinguished Name) are only valid for a single hostname. For example, a certificate created for example.com is not valid on www.example.com unless a Subject Alternate Name of "www.example.com" is added to the certificate. A certificate authority may charge an additional fee if your certificate contains 1 or more SAN extensions.

8. Click **OK**.

Receiving a signed certificate from a certificate authority

This topic describes how to receive an electronically mailed certificate from a certificate authority (CA), that is designated as a trusted CA on your server. A certificate authority is a trusted third-party organization or company that issues digital certificates used to create digital signatures and public-private key pairs.

About this task

The certificate authority can send more than one certificate. In addition to the certificate for your server, the CA can also send additional signing certificates or intermediate CA certificates. For example, Verisign includes an intermediate CA certificate when sending a Global Server ID certificate. Before receiving the server certificate, receive any additional intermediate CA certificates. Follow the instructions in the Storing a CA certificate topic to receive intermediate CA certificates.

Receive the CA-signed certificate into a key database as follows:

Procedure

1. Start the IKEYMAN user interface.
2. Click **Key Database File** from the main UI, then click **Open**.
3. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if you use the default. Click **OK**.
4. Enter your correct password in the Password Prompt dialog box, then click **OK**.
5. Click **Personal Certificates** in the Key database content frame, then click **Receive**.
6. Enter the name of a valid Base64-encoded file in the Certificate file name text field in the Receive certificate from a file dialog box. Click **OK**.

Displaying default keys and certificate authorities

This section describes how to view trusted certificate authorities and display default keys within a key database.

About this task

A trusted certificate authority (CA) issues and manages public keys for data encryption. A key database is used to share public keys that are used for secure connections. The tasks that follow show how to view the certificate authorities that are in your database, along with their expiration dates.

Procedure

- Display the default key entry as follows:
 1. Start the IKEYMAN user interface.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click the **key.kdbfile**, if using the default. Click **OK**.
 4. Enter your password in the Password Prompt dialog box, then click **OK**.
 5. Click **Personal Certificates** in the Key Database content frame, and click the **CA certificate** label name.
 6. Click **View/Edit** and view the certificate default key information in the Key Information window.
- Display a list of trusted certificate authorities (CAs) in a key database as follows:
 1. Start the IKEYMAN user interface.
 2. Click **Key Database File** from the main UI, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click **key.kdb** if you are using the default.
 4. Enter your correct password in the Password prompt dialog box, and click **OK**.
 5. Click **Signer Certificates** in the Key database content frame.
 6. Click **Signer Certificates**, **Personal Certificates**, or **Certificate Requests**, to view the list of CAs in the Key Information window.

What to do next

The version of iKeyman that is provided by the bundled Java Runtime Environment (JRE) does not add a default list of signer certificates to newly-created key databases. Add default signer certificates in iKeyman, as follows:

1. Select **Signer Certificates** from the drop-down menu in the iKeyman window.
2. Click the "Populate" on the right-hand side of the iKeyman window.
3. Click the grey boxes next to the certificate authority names (Entrust, RSA Data Security, Thawte, Verisign) so they display as checked.
4. Click **OK**.

Storing a certificate authority certificate

This topic describes how to store a certificate from a certificate authority (CA) that is not a trusted CA.

About this task

Store a certificate from a certificate authority (CA) who is not a trusted CA as follows:

Procedure

1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.

2. Click **Key Database File** from the main user interface, then click **Open**.
3. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if using the default. Click **OK**.
4. Enter your password in the Password Prompt dialog box, then click **OK**.
5. Click **Signer Certificates** in the Key Database content frame, then click **Add**.
6. In the Add CA Certificate from a File dialog box, click the **Base64-encoded ASCII data certificate file name**, or use the Browse option. Click **OK**.
7. In the Label dialog box, enter a label name and click **OK**.

Storing the encrypted database password in a stash file

This section describes how you would store your database password in a stash file.

About this task

For a secure network connection, you can store the encrypted database password in a stash file.

Important: These stash files should be treated as highly sensitive information.

Procedure

- Store the password while a database creates as follows:
 1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
 2. Click **Key Database File** from the main user interface, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click the **key.kdbfile**, if using the default. Click **OK**.
 4. Enter your password in the Password Prompt dialog box, then enter again to confirm your password.
 5. Select the stash box and click **OK**.
 6. Click **Key Database File > Stash Password**.
 7. Click **OK** in the information dialog box.
- Store the password after creating a database as follows:
 1. Start the IKEYMAN user interface. Refer to Starting the Key Management utility for platform-specific instructions.
 2. Click **Key Database File** from the main user interface, then click **Open**.
 3. Enter your key database name in the Open dialog box, or click the **key.kdb** file, if you use the default. Click **OK**.
 4. Enter your correct password in the Password Prompt dialog box, and click **OK**.
 5. Click **Key Database File**, then click **Stash Password**.
 6. Click **OK** in the Information dialog box.





Managing keys with the gskcmd command line interface (Distributed systems)

The Java command line interface to IKEYMAN, gskcmd, provides the necessary options to create and manage keys, certificates and certificate requests.

About this task

Important: Only use gskcmd, the command line interface, if you are unable to use IKEYMAN, the graphical user interface.

Global Security Kit (GSKit) certificate management tools are installed in the `<ih$inst>/bin/` directory. These tools should only be run from the installation directory. Examples for the following commands should include the full directory path, such as `<ih$inst>/bin/gskcmd`.

-  `gskver.bat`, `ikeyman.bat`, `gskcmd.bat`, `gskcmd`, and `gskcapiCmd`.
-    `gskver`, `ikeyman`, and `gskcmd`.

To have a secure network connection, create a key for secure network communications and receive a certificate from a certificate authority (CA) that is designated as a trusted CA on your server. Use `gskcmd`, the utility command line interface, for configuration tasks that are related to public and private key creation and management.

The `gskcmd` user interface uses Java and native command line invocation, enabling IKEYMAN task scripting.

You cannot use `gskcmd` for configuration options that update the server configuration file, `httpd.conf`. For options that update the server configuration file, use the IBM HTTP Server administration server.

Procedure

- Use `gskcmd` to create key databases, public and private key pairs, and certificate requests using the command-line interface.
- If you act as your own certificate authority (CA), you can use `gskcmd` to create self-signed certificates.
- If you act as your own CA for a private Web network, you have the option to use the server CA utility to generate and issue signed certificates to clients and servers in your private network.
- Manage the database password using the command line.
- Create a public and private key pair and certificate request using the `gskcmd` command-line interface or `GSKCapiCmd`.
- Import and export keys using the command line. If you want to reuse an existing key from another database, you can import that key. Conversely, you can export your key into another database or to a PKCS12 file. PKCS12 is a standard for securely storing private keys and certificates. You can use the `gskcmd` command-line interface or `GSKCapiCmd` tool.
- Display default keys and certificate authorities within a key database.
- Store a certificate authority certificate from a certificate authority (CA) that is not a trusted CA.
- Store the encrypted database password in a stash file.
- Use `gskcmd` to create key databases, public and private key pairs, and certificate requests.
- If you act as your own certificate authority (CA), you can use `gskcmd` to create self-signed certificates.
- If you act as your own CA for a private Web network, you have the option to use the server CA utility to generate and issue signed certificates to clients and servers in your private network.

What to do next

For more information about the `gskcapiCmd` command line interface, see the `GSKCapiCmd` User's Guide on the WebSphere Application Server Library page. For more information about the `gskcmd` (`ikeycmd`) command, see the IBM Developer Kit and Runtime Environment, Java 2 Technology Edition, Version 6.0 `iKeyman 8.0` User's Guide .

Using the `gskcmd` command

The `gskcmd` command provides a command line interface for certificate management tasks that might otherwise be provided by the `ikeyman` command.

Procedure

1. You can invoke the `gskcmd` from the `<ihsinst>/bin/` directory.

- `Windows` `gskcmd.bat`
- `AIX` `Linux` `Solaris` `gskcmd`

2. Perform the certificate management tasks that you want to complete.

Key Management Utility command-line interface (`gskcmd`) syntax

This topic contains a description of the syntax that you can use with the `gskcmd` command.

Syntax

For more information, see “Using the `gskcmd` command” on page 142.

The syntax follows.

```
gskcmd <object> <action> [options]
```

Where:

- The object includes one of the following:
 - `-keydb`: Actions taken on the key database (either a CMS key database file, a WebDB key ring file, or SSLight class)
 - `-cert`: Actions taken on a certificate
 - `-certreq`: Actions taken on a certificate request
 - `-help`: Displays help for the `gskcmd` invocations
 - `-version`: Displays version information for `gskcmd`

The action represents the specific action to take on the object, and options represents the options, both required and optional, specified for the object and action pair.

The object and action keywords are positional and you must specify them in the selected order. However, options are not positional and you can specify them in any order, as an option and operand pair.

Table 14. Actions for `gskcmd` command objects. The table describes each action possible on a specified object that you can use with the `gskcmd` command.

Object	Actions	Description
-keydb	-changepw	Change the password for a key database
	-convert	Convert a key database from one format to another
	-create	Create a key database
	-delete	Delete the key database
	-stashpw	Stash the password of a key database into a file
-cert	-add	Add a CA certificate from a file into a key database
	-create	Create a self-signed certificate
	-delete	Delete a CA certificate

Table 14. Actions for gskcmd command objects (continued). The table describes each action possible on a specified object that you can use with the gskcmd command.

	-export	Export a personal certificate and its associated private key from a key database into a PKCS#12 file, or to another key database
	-extract	Extract a certificate from a key database
	-getdefault	Get the default personal certificate
	-import	Import a certificate from a key database or PKCS#12 file
	-list	List all certificates
	-modify	Modify a certificate. (Currently the only field you can modify is the Certificate trust field)
	-receive	Receive a certificate from a file into a key database
	-setdefault	Set the default personal certificate
	-sign	Sign a certificate stored in a file with a certificate stored in a key database and store the resulting signed certificate in a file
-certreq	-create	Create a certificate request
	-delete	Delete a certificate request from a certificate request database
	-details	List the detailed information of a specific certificate request
	-extract	Extract a certificate request from a certificate request database into a file
	-list	List all certificate requests in the certificate request database
	-recreate	Recreate a certificate request
-help		Display help information for the gskcmd command
-version		Display gskcmd version information

The following table describes the options that you can use with the gskcmd command.

Option	Description
dB	Fully qualified path name of a key database
-default_cert	Sets a certificate to use as the default certificate for client authentication (yes or no). Default is no.
-dn	X.500 distinguished name. Input as a quoted string of the following format (only CN, O, and C are required): "CN=Jane Doe,O=IBM,OU=Java Development,L=Endicott,ST=NY,ZIP=13760,C=country"
encryption	Strength of encryption used in certificate export command (strong or weak). Default is strong.

-expire	Expiration time of either a certificate or a database password (in days). Defaults are: 365 days for a certificate and 60 days for a database password.
-file	File name of a certificate or certificate request (depending on specified object).
-format	Format of a certificate (either ASCII for Base64_encoded ASCII or binary for Binary DER data). Default is ASCII.
-label	Label attached to a certificate or certificate request
-new_format	New format of key database
-new_pw	New database password
-old_format	Old format of key database
-pw	Password for the key database or PKCS#12 file. See Creating a new key database.
-size	Key size (512, 1024, or 2048). Default is 1024. The 2048 key size is available if you are using Global Security Kit (GSKit) Version 7.0.4.14 and later.
-stash	Indicator to stash the key database password to a file. If specified, the password will be stashed in a file.
-target	Destination file or database
-target_pw	Password for the key database if -target specifies a key database. See Creating a new key database.
-target_type	Type of database specified by -target operand (see -type)
-trust	Trust status of a CA certificate (enable or disable). Default is enable.
-type	Type of database. Allowable values are CMS (indicates a CMS key database), webdb (indicates a keyring), sslight (indicates an SSLight .class), or pkcs12 (indicates a PKCS#12 file).
-x509version	Version of X.509 certificate to create (1, 2 or 3). Default is 3.

Creating a new key database using the command-line interface

A *key database* is a file that the server uses to store one or more key pairs and certificates. You can use one key database for all your key pairs and certificates, or create multiple databases.

About this task

You can create multiple databases if you prefer to keep certificates in separate databases.

Procedure

- Create a new key database using the gskcmd command-line interface by entering the following command (as one line):

```
<i>hsinst>/bin/gskcmd -keydb -create -db <filename> -pw <password> -type  
<cms | jks | jceks | pks12> -expire <days> -stash
```

where:

- -db <filename> is the name of the database.
- -expire <days> is the number of days before password expires. This parameter is only valid for CMS key databases.
- -keydb Specifies the command is for the key database.

- `-pw <password>` is the password to access the key database.
- `-type <cms | jks | jceks | pkcsk>` is the database type. Note: IBM HTTP Server only handles a CMS key database.
- `-stash` stashes the password for the key database. When the `-stash` option is specified during the key database creation, the password is stashed in a file with a filename built as follows:
`<filename_of_key_database>.sth`

This parameter is only valid for CMS key databases. For example, if the database being created is named `keydb.kdb`, the stash filename is `keydb.sth`. Note: Stashing the password is required for IBM HTTP Server.

- Create a new key database using the `GSKCapiCmd` tool. `GSKCapiCmd` is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing `GSKit` Java command line tool has, except `GSKCapiCmd` supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use `GSKCapiCmd` to manage all aspects of a CMS key database. `GSKCapiCmd` does not require Java to be installed on the system.

```
<ihsinst>/bin/gskcapiCmd -keydb -create -db <name> [-pw <passwd>] [-type <cms>] [-expire <days>] [-stash] [-fips] [-strong]
```

Managing the database password using the command line

This topic describes passwords for key databases. A key database is used to store public keys that are used for secure connections.

About this task

When you create a new key database, you specify a key database password. This password protects the private key. The private key is the only key that can sign documents or decrypt messages that are encrypted with the public key. Changing the key database password frequently is a good practice.

Use the following guidelines when specifying the password:

- The password must come from the U.S. English character set.
- The password must contain at least six characters and contain at least two nonconsecutive numbers. Make sure that the password does not consist of publicly obtainable information about you, such as the initials and birth date for you, your spouse, or children.
- Stash the password.

Procedure

- Change the password for a key database using the `gskcmd` command-line interface. Enter the following command as one line:

```
<ihsinst>/bin/gskcmd -keydb -changepw -db <filename>.kdb -pw <password> -new_pw <new_password> -expire <days> -stash
```

where:

- `-db <filename>` is the name of the database.
- `-changepw` changes the password.
- `-keydb` specifies the command is for the key database.
- `-new_pw <new_password>` is the new key database password. This password must be different than the old password and cannot be a NULL string.
- `-pw <password>` is the password to access the key database.
- `-expire <days>` is the number of days before password expires. This parameter is only valid for CMS key databases.

- `-stash` stashes the password for the key database. This parameter is only valid for CMS key databases. Stashing the password is required for IBM HTTP Server.
- Change the password using the `GSKCapiCmd` tool. `GSKCapiCmd` is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing `GSKit` Java command line tool has, except `GSKCapiCmd` supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use `GSKCapiCmd` to manage all aspects of a CMS key database. `GSKCapiCmd` does not require Java to be installed on the system.

```
<ihsinst>/bin/gskcapiCmd -keydb -changePW -db <name> [-crypto <module name> -tokenlabel <token label>]
[-pw <passwd>]-new_pw <new passwd> [-expire <days>] [-stash] [-fips] [-strong]
```

Results

The key database now accepts the new password.

Creating a new key pair and certificate request

You find key pairs and certificate requests stored in a key database. This topic provides information on how to create a key pair and certificate request.

About this task

Create a public and private key pair and certificate request using the `gskcmd` command-line interface or `GSKCapiCmd` tool, as follows:

Procedure

1. Use the `gskcmd` command-line interface. Enter the following command (as one line):

```
<ihsinst>/bin/gsk7cmd -certreq -create -db <filename> -pw <password> -label <label> -dn <distinguished_name>
-size <2048 | 1024 | 512> -file <filename> -san_dnsname <DNS name value=[,<DNS name value>] -san_emailaddr <email address>
```

where:

- `-certreq` specifies a certificate request.
- `-create` specifies a create action.
- `-db <filename>` specifies the name of the database.
- `-pw` is the password to access the key database.
- `label` indicates the label attached to the certificate or certificate request.
- `dn <distinguished_name>` indicates an X.500 distinguished name. Input as a quoted string of the following format (only CN, O, and C are required): CN=common_name, O=organization, OU=organization_unit, L=location, ST=state, province, C=country

Note: For example, "CN=weblinux.raleigh.ibm.com,O=IBM,OU=IBM HTTP Server,L=RTP,ST=NC,C=US"

- `-size <2048 | 1024 | 512>` indicates a key size of 2048, 1024, or 512. The default key size is 1024. The 2048 key size is available if you are using Global Security Kit (GSKit) Version 7.0.4.14 and later.
- `-file <filename>` is the name of the file where the certificate request will be stored.
- `-san * <subject alternate name attribute value> | <subject alternate name attribute value>` specifies the subject alternate name extensions in the certificate request that inform SSL clients of alternate hostnames that correspond to the signed certificate.

These options are only valid if the following line is entered in the `ikminit.properties` file.

`DEFAULT_SUBJECT_ALTERNATE_NAME_SUPPORT=true`. The * (asterisk) can have the following values:

dnsname

The value must be formatted using the "preferred name syntax" according to RFC 1034, such as the example, zebra.tek.ibm.com.

emailaddr

The value must be formatted as an "addr-spec" according to RFC 822, such as the example, myname@zebra.tek.ibm.com

ipaddr

The value is a string representing an IP address formatted according to RFC 1338 and RFC 1519, such as the example, 193.168.100.115

The values of these options are accumulated into the subject alternate name extended attribute of the generated certificate. If the options are not used then this extended attribute is not added to the certificate.

- `-ca <true | false>` specifies the basic constraint extension to the self-signed certificate. The extension is added with a `CA:true` and `PathLen:<max int>` if the value passed is true or not added if the value passed is false.

Use the GSKCapiCmd tool. GSKCapiCmd is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing GSKit Java command line tool has, except GSKCapiCmd supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use GSKCapiCmd to manage all aspects of a CMS key database. GSKCapiCmd does not require Java to be installed on the system.

```
<ihsinst>/bin/gskcapiCmd -certreq -create -db <name> [-crypto <module name> [-tokenlabel <token label>]]
[-pw <passwd>] -label <label> -dn <dist name> [-size <2048 | 1024 | 512>] -file <name> [-secondaryDB
<filename> -secondaryDBpw <password>] [-fips] [-sigalg <md5 | sha1|sha224|sha256|sha384|sha512>]
```

gotcha: On Unix type operating systems it is recommended to always encapsulate string values associated with all tags in double quotes (""). You will also need to escape, using a backslash character, the following characters if they appear in the string values: '!', '\', '"', '^'. This will prevent some command line shells from interpreting specific characters within these values. (e.g. `gsk7capiCmd -keydb -create -db "/tmp/key.kdb" -pw "j!jj"`). Note however when prompted by `gsk7capiCmd` for a value (for example a password) quoting the string and adding the escape characters should not be done. This is because the shell is no longer influencing this input.

2. Verify that the certificate was successfully created:
 - a. View the contents of the certificate request file you created.
 - b. Ensure that the key database recorded the certificate request:


```
<ihsinst>/bin/gskcmd -certreq -list -db <filename> -pw <password>
```

 You should see the label listed that you just created.
3. Send the newly-created file to a certificate authority.

Importing and exporting keys using the command line

This topic describes how to import and export keys.

About this task

If you want to reuse an existing key from another database, you can import that key. Conversely, you can export your key into another database or to a PKCS12 file. PKCS12 is a standard for securely storing private keys and certificates. You can use the `gskcmd` command-line interface or GSKCapiCmd tool.

Procedure

- Use the `gskcmd` command-line interface to import certificates from another key database, as follows:

```
<i>hsinst>/bin/gskcmd -cert -import -db <filename> -pw <password>
-label <label> -type <cms | JKS | JCEKS | pkcs12>
-new_label <label> -target <filename> -target_pw <password>
-target_type <cms | JKS | JCEKS | pkcs12>
```

where:

- -cert - specifies a certificate.
- -import - specifies an import action.
- -db <filename> - indicates the name of the database.
- -pw <password> - indicates the password to access the key database.
- -label <label> - indicates the label that is attached to the certificate.
- -new_label <label> - re-labels the certificate in the target key database.
- -type <cms | JKS | JCEKS | pkcs12> - specifies the type of database.
- -target <filename> - indicates the destination database.
- -target_pw <password> - indicates the password for the key database if -target specifies a key database
- -target_type <cms | JKS | JCEKS | pkcs12> - indicates the type of database that is specified by the -target operand.
- pfx - imported file in Microsoft .pfx file format.

Use the GSKCapiCmd tool to import certificates from another key database. GSKCapiCmd is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing GSKit Java command line tool has, except GSKCapiCmd supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use GSKCapiCmd to manage all aspects of a CMS key database. GSKCapiCmd does not require Java to be installed on the system.

```
<i>hsinst>/bin/gskcapiCmd -cert
-import -db <name> [-crypto <module name> [-tokenlabel <token label>] [-pw <passwd>]
[-secondaryDB <filename> -secondaryDBpw <password>]
-label <label> [-type < cms>] -target <name>
[-target_pw<passwd>] [-target_type <cms|pkcs11>] [-new_label < label>] [-fips]
```

- Use the gskcmd command-line interface to export certificates from another key database, as follows:

```
gskcmd -cert -export -db <filename> -pw <password> -label <label>
-type <cms | jks | jceks | pkcs12> -target <filename>
-target_pw <password> -target_type <cms | jks | jceks | pkcs12>
```

where:

- -cert specifies a personal certificate.
- -export specifies an export action.
- -db <filename> is the name of the database.
- -pw <password> is the password to access the key database.
- -label <label> is the label attached to the certificate.
- -target <filename> is the destination file or database. If the **target_type** is JKS, CMS, or JCEKS, the database specified here must exist.
- -target_pw is the password for the target key database.
- -target_type <cms | jks | jceks | pkcs12> is the type of database specified by the -target operand.
- -type <cms | jks | jceks | pkcs12> is the type of database key.

Use the GSKCapiCmd tool to export certificates from another key database. GSKCapiCmd is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing GSKit Java command line tool has, except GSKCapiCmd supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the

existing Java tool. You can use GSKCapiCmd to manage all aspects of a CMS key database. GSKCapiCmd does not require Java to be installed on the system.

```
<ihsinst>/bin/gskcapicmd -cert extract -db <name> |  
-crypto <module name> [-tokenlabel <token label>] -pw <passwd>  
-label <label> -target <name> [-format <ascii | binary>] [-secondaryDB <filename>  
-secondaryDBpw <password> ][-fips]
```

Creating a self-signed certificate

A self-signed certificate provides a certificate to enable SSL sessions between clients and the server, while waiting for the officially-signed certificate to be returned from the certificate authority (CA). A private and public key are created during this process. Creating a self-signed certificate generates a self-signed X509 certificate in the identified key database. A self-signed certificate has the same issuer name as its subject name.

About this task

Use this procedure if you are acting as your own CA for a private Web network. Use the IKEYCMD command-line interface or the GSKCapiCmd tool to create a self-signed certificate.

Procedure

- Create a self-signed certificate using the IKEYCMD command-line interface, as follows:

```
gskcmd -cert -create -db <filename> -pw <password> -size <2048 | 1024 | 512> -dn <distinguished_name>  
-label label> -default_cert <yes | no> - expire <days> -san dnsname <DNS name value>[,<DNS name value>]  
-san emailaddr <email address value>[,<email address value>]  
-san ipaddr <IP address value>[,<IP address value>][-ca <>true | false>]
```

where:

- -cert specifies a self-signed certificate.
- -create specifies a create action.
- -db <filename> is the name of the database.
- -pw <password> is the password to access the key database.
- -dn <distinguished_name> - indicates an X.500 distinguished name. Input as a quoted string of the following format (Only CN, O, and C are required): CN=common_name, O=organization, OU=organization_unit, L=location, ST=state, province, C=country
For example, "CN=weblinux.raleigh.ibm.com,O=IBM,OU=IBM HTTP Server,L=RTP,ST=NC,C=US"
- -label <label> is a descriptive comment used to identify the key and certificate in the database.
- -size <2048 | 1024 | 512> indicates a key size of 2048, 1024, or 512. The default key size is 1024. The 2048 key size is available if you are using Global Security Kit (GSKit) Version 7.0.4.14 and later.
- -default_cert<yes | no> specifies whether this is the default certificate in the key database.
- -expire <days> indicates the default validity period for new self-signed digital certificates is 365 days. The minimum is 1 day. The maximum is 7300 days (twenty years).
- -san * <subject alternate name attribute value> | <subject alternate name attribute value> specifies the subject alternate name extensions in the certificate request that inform SSL clients of alternate hostnames that correspond to the signed certificate.

These options are only valid if the following line is entered in the ikmunit.properties file. DEFAULT_SUBJECT_ALTERNATE_NAME_SUPPORT=true. The * (asterisk) can have the following values:

dnsname

The value must be formatted using the "preferred name syntax" according to RFC 1034, such as the example, zebra,tek.ibm.com.

emailaddr

The value must be formatted as an "addr-spec" according to RFC 822, such as the example, myname@zebra.tek.ibm.com

ipaddr

The value is a string representing an IP address formatted according to RFC 1338 and RFC 1519, such as the example, 193.168.100.115

The values of these options are accumulated into the subject alternate name extended attribute of the generated certificate. If the options are not used then this extended attribute is not added to the certificate.

- `-ca <true | false>` specifies the basic constraint extension to the self-signed certificate. The extension is added with a `CA:true` and `PathLen:<max int>` if the value passed is true or not added if the value passed is false.
- Create a self-signed certificate using the GSKCapiCmd tool. GSKCapiCmd is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing GSKit Java command line tool has, except GSKCapiCmd supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use GSKCapiCmd to manage all aspects of a CMS key database. GSKCapiCmd does not require Java to be installed on the system.

```
gskcapiCmd -cert -create [-db <name>] | [-crypto <module name> -tokenlabel <token label>] [-pw <passwd>]
-label <label> -dn <dist name> [-size <2048|1024|512>] [-x509version <1|2|3>] [-default.cert <yes|no>]
[-expire <days>] [-secondaryDB <filename> -secondaryDBpw <password>] [-ca <true|false>] [-fips]
[-sigalg<md5|sha1|sha224|sha256|sha384|sha512>]
```

Note: On Unix type operating systems it is recommended to always encapsulate string values associated with all tags in double quotes (“”). You will also need to escape, using a ‘\’ character, the following characters if they appear in the string values: ‘!’, ‘\’, ‘”’, ‘\’’. This will prevent some command line shells from interpreting specific characters within these values. (e.g. `gsk7capiCmd -keydb -create -db “/tmp/key.kdb” -pw “j\!j”`). Note however when prompted by `gsk7capiCmd` for a value (for example a password) quoting the string and adding the escape characters should not be done. This is because the shell is no longer influencing this input.

Receiving a signed certificate from a certificate authority

This topic describes how to receive an electronically mailed certificate from a certificate authority (CA) that is designated as a trusted CA on your server. A certificate authority is a trusted third-party organization or company that issues digital certificates that are used to create digital signatures and public-private key pairs.

About this task

The certificate authority can send more than one certificate. In addition to the certificate for your server, the CA can also send additional signing certificates or intermediate CA certificates. For example, Verisign includes an intermediate CA certificate when sending a Global Server ID certificate. Before receiving the server certificate, receive any additional intermediate CA certificates. Follow the instructions in the Storing a CA certificate topic to receive intermediate CA certificates.

If the CA that issuing your CA-signed certificate is not a trusted CA in the key database, store the CA certificate first and designate the CA as a trusted CA. Then you can receive your CA-signed certificate into the database. You cannot receive a CA-signed certificate from a CA that is not a trusted CA. For instructions, see Storing a certificate authority certificate.

Procedure

- Receive the CA-signed certificate into a key database using the `gskcmd` command-line interface, as follows:

```
<ihsinst>/bin/gskcmd -cert -receive -file <filename> -db <filename> -pw <password>
-format <ascii | binary> -label <label> -default_cert <yes | no>
```

where:

- -cert specifies a self-signed certificate.
 - -receive specifies a receive action.
 - -file <filename> is a file containing the CA certificate.
 - -db <filename> is the name of the database.
 - -pw <password> is the password to access the key database.
 - -format <ascii | binary> specifies that the certificate authority might provide the CA certificate in either ASCII or binary format.
 - -default_cert <yes | no> indicates whether this is the default certificate in the key database.
 - -label specifies the label that is attached to a CA certificate.
 - -trust indicates whether this CA can be trusted. Use enable options when receiving a CA certificate.
- Receive the CA-signed certificate into a key database using the GSKCapiCmd tool. GSKCapiCmd is a tool that manages keys, certificates, and certificate requests within a CMS key database. The tool has all of the functionality that the existing GSKit Java command line tool has, except GSKCapiCmd supports CMS and PKCS11 key databases. If you plan to manage key databases other than CMS or PKCS11, use the existing Java tool. You can use GSKCapiCmd to manage all aspects of a CMS key database. GSKCapiCmd does not require Java to be installed on the system.

```
<ihsinst>/bin/gskcapiCmd -cert -receive -file <name>
-db <name> [-crypto <module name> [-tokenlabel <token label>]]
[-pw <passwd>][-default_cert <yes|no>][-fips>
```

Displaying default keys and certificate authorities

This section describes how to view trusted certificate authorities and display default keys within a key database.

About this task

A trusted certificate authority (CA) issues and manages public keys for data encryption. A key database is used to share public keys that are used for secure connections. The tasks that follow show how to view the certificate authorities that are in your database, along with their expiration dates.

Procedure

- Display a list of trusted CAs in a key database by entering the following command as one line:

```
<ihsinst>/bin/gskcmd -cert -list CA -db < dbname > -pw <password> -type <cms | jks |jceks | pkcs12>
```
- Display a list of certificates in a key database and their expiration dates by enter the following command:

```
<ihsinst>/bin/gskcmd -cert -list -expiry < days > -db < filename > -pw < password > - type < type >
```

where:

- -cert indicates the operation applies to a certificate.
- -list <all | personal | CA | site> specifies a list action. The default is to list all certificates.
- -expiry <days> indicates that validity dates should be displayed. Specifying the number of days is optional, though when used will result in displaying all certificates that expire within that amount of days. To list certificates that have already expired, enter the value 0.
- -db <filename> is the name of the key database. It is used when you want to list a certificate for a specific key database.
- -pw <password> specifies the password to access the key database.
- -type <cms | JKS | JCEKS | pkcs12> specifies the type of database.

Storing a certificate authority certificate

This topic describes how to store a certificate from a certificate authority (CA) that is not a trusted CA.

Procedure

To store a certificate from a CA that is not a trusted CA, use the following command:

```
<ihsinst>/bin/gskcmd -cert -add -db <filename>.kdb -pw <password> -label <label>  
-format <ascii | binary> -trust <enable | disable> -file <filename>
```

where:

- -add specifies an add action.
- -cert indicates the operation applies to a certificate.
- -db <filename> is the name of the database.
- -file <filename> specifies the file containing the CA certificate.
- -format <ascii | binary> indicates the certificate authorities might supply a binary or an ASCII file.
- -label <label> is the label attached to a certificate or certificate request.
- -pw <password> is the password to access the key database.
- -trust <enable | disable> indicates whether this CA can be trusted. The default is enable and indicates that the CA can be trusted.

Storing the encrypted database password in a stash file

For a secure network connection, you can store the CMS encrypted database password in a stash file.

About this task

Complete one of the following steps to store the encrypted database password in a stash file.

Procedure

- To store the password using the gskcmd command-line interface, enter the following command on one line.

```
<ihsinst>/bin/gskcmd -keydb -create -db <path_to_db>/<db_name> -pw <password> -type  
cms -expire <days> -stash
```

- If the CMS database has been created, enter the following command on one line to store the password using the gskcmd command line interface.

```
<ihsinst>/bin/gskcmd -keydb -stashpw -db <db_name> -pw <password>
```

Managing keys with the native key database gskkyman (z/OS systems)

Use the native z/OS key management (gskkyman key database) support for key management tasks.

About this task

To have a secure network connection, create a key for secure network communications and receive a certificate from a certificate authority (CA) that is designated as a trusted CA on your server.

IBM HTTP Server on z/OS does not support IKEYMAN or gskcmd.

Use gskkyman to create key databases, public and private key pairs, and certificate requests. If you act as your own CA, you can use gskkyman to create self-signed certificates. If you act as your own CA for a private Web network, you have the option to use the server CA utility to generate and issue signed certificates to clients and servers in your private network.

Procedure

- To use native z/OS key management (gskkyman) tasks, refer to *Cryptographic Services PKI Services Guide and Reference* document (SA22-7693). Link to this document from the *z/OS Internet Library*.
- A typical task that this document contains is using a gskkyman key database for your certificate store. See section "Appendix B. Using a gskkyman key database" for a description of how to use gskkyman.
- **Important:** The certificate requests that gskkyman generates for use with IBM HTTP Server should use RSA keys and not DSA keys.

Getting started with the cryptographic hardware for SSL (Distributed systems)

The IBM 4758 and other cryptographic devices require the PKCS11 support software for the host machine and internal firmware.

About this task

You will need the manual that explains software installation and cryptographic coprocessor microcode loading.

Note: The support software and manual do not come with the IBM 4758 card, but you can download them from <http://www.ibm.com/security/cryptocards/index.shtml>. From the download site, obtain the PKCS#11 Model 002/023 software and the PKCS#11 installation manual.

Procedure

1. After installing the support software on your machine and loading the microcode on the cryptographic device, initialize the card.
2. Configure IBM HTTP Server to pass the module for the PKCS11 device, the token label, the key label of the key created by the PKCS11 device, and the user PIN password of the token to the GSKit for access to the key for the PKCS11 device by modifying the configuration file. The PKCS11 module differs for each platform and PKCS11 device.

AIX For the IBM hardware cryptographic devices (for example, IBM 4758 card and IBM e-business Cryptographic Accelerator) the PKCS11 module ships with the bos.pkcs11 package.

3. Install the devices.pci.14109f00 device for the IBM 4758 and the devices.pci.1410e601 device for the IBM e-business Cryptographic Accelerator.

For the IBM 4758 on Windows, the PKCS11 module comes with the PKCS11 software available for download from: <http://www.ibm.com/security/cryptocards/ordersoftware.shtml>. For nCipher, the PKCS11 module ships with nCipher software and is located in the \$NFAST_HOME/toolkits/pkcs11 directory.

The default locations of the PKCS11 modules for each PKCS11 device follow:

- nCipher:
 - **AIX** **Linux** **Solaris** /opt/nfast/toolkits/pkcs11/libcknfast.so
 - **HP-UX** /opt/nfast/toolkits/pkcs11/libcknfast.sl
 - **Windows** C:\nfast\toolkits\pkcs11\cknfast.dll
- IBM 4758:
 - **AIX** /usr/lib/pkcs11/PKCS11_API.so
 - **Windows** \$PKCS11_HOME\bin\nt\cryptoki.dll
- IBM e-business Cryptographic Accelerator:
 - **AIX** /usr/lib/pkcs11/PKCS11_API.so

Cryptographic hardware for Secure Sockets Layer

IBM HTTP Server supports many types of cryptographic hardware devices.

The following table contains hardware cryptographic devices that have been tested with IBM HTTP Server. However, since device drivers for these devices are frequently upgraded by the hardware vendors to correct customer-reported problems or to provide support for new operating system platforms, check with the hardware vendors for specific applications of these devices.

A list of cryptographic devices tested with GSKit is available at this IBM Web site: IBM Global Security Kit, Version 7 - PKCS#11 Device Integration. If your device is not listed, contact the device vendor to ensure that the device functions correctly when used with IBM HTTP Server.

Device	Key Storage	Acceleration Support	Notes
Rainbow Cryptoswift PCI with BSAFE Interface Model	No	Yes	Use with SSLAcceleratorDisable directive only. Supported on HP, Solaris, and the Windows operating systems.
nCipher nFast Accelerator with BHAPI plug-in under BSAFE 4.0	No	Pure accelerator	Requires either a SCSI or PCI-based nForce unit; use with SSLAcceleratorDisable directive only. Supported on Solaris and Windows operating systems.
nCipher nForce Accelerator, <i>accelerator mode</i>	No	Yes	Uses the BHAPI and BSAFE interface. Supported on Solaris and Windows operating systems.
nCipher nForce Accelerator, Key stored accelerator mode	Yes	Yes	Uses the PKCS#11 interface. Requires either a SCSI, or PCI-based nForce unit. Move to nCipher nForce Accelerator V4.0 or later for better performance. Supported on AIX, HP, Linux, Solaris, and Windows operating systems.
IBM 4758 Model 002/023 PCI Cryptographic Coprocessors	Yes	No	Supported on AIX and Windows operating systems.

AIX operating systems. Support for the following adapters has been tested with WebSphere Application Server V4.0.2 or later:

Device	Key Storage	Acceleration Support	Notes
Rainbow Cryptoswift PCI with BSAFE Interface Model CS/200 and CS/600	No	Yes	Supported on the AIX operating system.
IBM e-business Cryptographic Accelerator	No	Yes	Uses the PKCS11 interface. Because this device uses the PKCS11 interface, the SSLAcceleratorDisable directive does not apply to this device. Supported on the AIX operating system.

Use the Rainbow Cryptoswift, IBM e-business Cryptographic Accelerator, nCipher nFast Accelerator and nCipher nForce Accelerator, for public key operations, and RSA key decryption. These devices store keys on your hard drive. Accelerator devices speed up the public key cryptographic functions of SSL, freeing up your server processor, which increases server throughput and shortens wait time. The Rainbow Cryptoswift, IBM e-business Cryptographic Accelerator, and nCipher accelerators incorporate faster performance and more concurrent secure transactions.

The PKCS#11 protocol either stores RSA keys on cryptographic hardware, or encrypts keys using cryptographic hardware to ensure protection. The nCipher nForce Accelerator can either perform acceleration, or it can perform both acceleration and key storage with PKCS#11 support. The IBM 4758 and nCipher nForce Accelerator with PKCS#11 support ensures inaccessible keys to the outside world. This support never reveals keys in an unencrypted form because the key is either encrypted by the hardware, or stored on the hardware.

nCipher nForce Accelerator V4.0 and later using PKCS11 key storage, has a nonremovable option which can noticeably improve performance. Contact nCipher Technical Support for instructions to turn on this feature.

Initializing IBM 4758 and IBM e-business Cryptographic Accelerator on AIX systems

To initialize IBM cryptographic hardware (IBM 4758 and IBM e-business Cryptographic Accelerator), you must obtain and install the most recent PKCS11 module.

Procedure

1. To initialize the IBM cryptographic hardware (IBM 4758 and IBM e-business Cryptographic Accelerator) on AIX, obtain and install the `bos.pkcs11` software. Obtain the most recent `bos.pkcs11` package from <http://www.ibm.com/servers/eserver/support/pseries/aixfixes.html>. Select your AIX version under **Specific fixes**. The `bos.pkcs11` package installs the PKCS11 module needed for the `SSLPKCSDriver` directive discussed below. You also need the `devices.pci.1410e601` device for the IBM e-business Cryptographic Accelerator and the `devices.pci.14109f00` and `devices.pci.14109f00` for the IBM 4758.
2. Initialize your token. After you install the PKCS11 software, initialize your device. You can access the Manage the PKCS11 subsystem panel from Smitty to initialize your PKCS11 device.
 - a. Select Initialize your token.
 - b. Set a security officer and User PIN, if not already set.
 - c. Initialize your user PIN. See Chapter 5: Token Initialization from the PKCS11 manual for more detailed information.

Initializing IBM 4758 Cryptographic Accelerator on Windows systems

To initialize the IBM 4758 card on Windows operating systems, you will need the PKCS11 software.

Procedure

1. Obtain the PKCS11 software from the following site: <http://www.ibm.com/security/cryptocards/>.
2. Use the `TOKUTIL.EXE` utility that installs with the PKCS11 software to initialize your IBM 4758 card on Windows operating systems.
3. Ensure you have the `cryptoki.dll` module in your path.
4. Refer to *Chapter 5: Token Initialization* from the PKCS11 documentation for more details.

Using IKEYMAN to store keys on a PKCS11 device

For IBM HTTP Server, you can use IKEYMAN for storing keys on a PKCS11 device.

Before you begin

Read about the IBMPKCS11Impl provider at <http://www.ibm.com/developerworks/java/jdk/security/60/secguides/pkcs11implDocs/IBMJavaPKCS11ImplementationProvider.html#ConfigFile>.

Procedure

- Obtain the file name and the path location of the cryptographic driver in order to store the keys on the PKCS11 device. The following are examples of path locations for PKCS11 devices:

- nCipher:

- **AIX** /opt/nfast/toolkits/pkcs11/libcknfast.so
- **HP-UX** /opt/nfast/toolkits/pkcs11/libcknfast.sl
- **Linux** /opt/nfast/toolkits/pkcs11/libcknfast.so
- **Solaris** /opt/nfast/toolkits/pkcs11/libcknfast.so
- **Windows** C:\nfast\toolkits\pkcs11\cknfast.dll

- IBM 4758

- **AIX** /usr/lib/pkcs11/PKCS11_API.so
- **Windows** \$PKCS11_HOME\bin\NT\cryptoki.dll

- IBM e-business Cryptographic Accelerator

- **AIX** /usr/lib/pkcs11/PKCS11_API.so

- If your Web server and Java Development Kit (JDK) are 64-bit, select a 64-bit vendor PKCS11 library. On some platforms, the 64-bit PKCS11 library filename has 64 appended to it.

AIX **HP-UX** **Linux** **Solaris** You can display the architecture of the Web server by running `apachectl -V`.

Windows You can display the architecture of the Web server by running `httpd.exe -V`.

- Determine the token label of the PKCS11 token that you use. Native vendor tools, such as `pkcsconf -its`, often display the token label.
- Create a PKCS11 configuration file describing your PKCS11 device using the following fields:
 - *library*: Full path to the proper architecture PKCS11 driver
 - *name*: Same as the token label from the previous step
 - *description*: Text field with your description
 - *attributes*: A set of attributes that you copy verbatim from the certificates example that the Web server uses

/opt/HTTPServer/conf/pkcs11.cfg example:

```
library = /usr/lib/pkcs11/PKCS11_API.so
name = mytokenlabel
description = description
attributes(*,CKO_PRIVATE_KEY,*) = { CKA_PRIVATE = true }
```

- Update the `java/jre/lib/security/java.security` file under the installation root to add a new security provider.
 - Have the new security provider reference the GSKit PKCS11 classes and the location of your PKCS11 configuration file.
 - Append the provider to the bottom of the provider list as the new highest-numbered provider.
 - Modify the following examples to specify the location of your configuration file.

Some of the lines are split on multiple lines for display purposes. Enter a line as a single line even if it displays in this task as multiple lines.

AIX **HP-UX** **Linux** **Solaris**

```
# The following line is the last pre-existing security provider.
security.provider.12=com...
# Add the following line.
security.provider.13=com.ibm.crypto.pkcs11impl.provider.IBMPKCS11Impl /opt/HTTPServer/conf/pkcs11.cfg
```

Windows

```
# The following line is the last pre-existing security provider.
security.provider.12=com...
# Add the following line.
security.provider.13=com.ibm.crypto.pkcs11impl.provider.IBMPKCS11Impl C:\opt\HTTPServer\conf\pkcs11.cfg
```

- Run **IKEYMAN** to store the keys on the PKCS11 device.

After launching **IKEYMAN**:

1. Select **Key Database File** from the menu, then **Open** to navigate to the Key database information dialog
2. From the drop-down for **Key Database Type**, select **PKCS11Config**.
If **PKCS11Config** is not an option, but **PKCS11Direct** is, you have an error that you must fix. Check your java.security work in prior steps. The PKCS11Direct option is not visible to the Web server. All other fields become locked, as the parameters are provided from the pkcs11.cfg file.
3. Click **OK** to navigate to the **Open Cryptographic Token** dialog.
The **Cryptographic Token Label** label of the PKCS11 device is displayed at the top of the panel. This label comes from the name field of the pkcs11.cfg file, and might be different from the native token label.
4. Complete the following actions on the **Open Cryptographic Token** dialog.
 - a. Enter the cryptographic token password for the PKCS11 device in the **Cryptographic Token Password** field. The password was previously set and is hardware-specific. This password is often called a user PIN in vendor documentation and tools.
 - b. Select the **Create new secondary key database file** option and complete prompts for creating a new secondary key database.
 - c. Click **OK**.

Results

After opening a cryptographic token successfully, **IKEYMAN** will display the certificates stored in the cryptographic token.

What to do next

You can create, import, or receive a personal certificate as you normally would. The private key is stored on your PKCS11 device.

Configuring IBM HTTP Server to use nCipher and Rainbow accelerator devices and PKCS11 devices

The IBM HTTP Server enables nCipher and Rainbow accelerator devices by default. To disable your accelerator device, add the **SSLAcceleratorDisable** directive to your configuration file.

Before you begin

When using the IBM e-business Cryptographic Accelerator, or the IBM 4758, the user ID under which the Web server runs must be a member of the PKCS11 group. You can create the PKCS11 group by installing the bos.pkcs11 package or its updates. Change the **Group** directive in the configuration file to group pkcs11.

About this task

If you want the IBM HTTP Server to use the PKCS11 interface, configure the following:

Procedure

1. Stash your password to the PKCS11 device, or optionally enable password prompting: Syntax: `sslstash [-c] <file> <function> <password>` where:
 - `-c`: Creates a new stash file. If not specified, an existing stash file is updated.
 - `file`: Represents a fully-qualified name of the file to create or update.
 - `function`: Represents the function for which the server uses the password. Valid values include `cr1` or `crypto`.
 - `password`: Indicates the password to stash.
2. Place the following directives in your configuration file:
 - `SSLPKCSDriver <fully qualified name of the PKCS11 driver used to access PKCS11 device>`
See `SSLPKCSDriver` directive for the default locations of the PKCS11 module, for each PKCS11 device.
 - `SSLServerCert <token label: key label of certificate on PKCS11 device>`
 - `SSLStashfile <fully qualified path to the file containing the password for the PKCS11 device>`
 - `Keyfile <fully qualified path to key file with signer certificates>`

Authenticating with LDAP on IBM HTTP Server using mod_ibm_ldap (Distributed systems)

This section describes how to configure LDAP to protect files on IBM HTTP Server.

Before you begin

Best Practice: Distributed operating systems If you are using the `mod_ibm_ldap` module for your LDAP configuration, consider migrating your `mod_ibm_ldap` directives to use the `mod_ldap` module. The `mod_ibm_ldap` module is provided with this release of IBM HTTP Server for compatibility with previous releases, however, you must migrate existing configurations to use the `mod_authnz_ldap` and `mod_ldap` modules to ensure future support for your LDAP configuration.

The LDAP module is not loaded into IBM HTTP Server by default. Without the `LoadModule` directive, the LDAP features are not available for use. In order to enable the LDAP function, add a `LoadModule` directive to the IBM HTTP Server `httpd.conf` file as follows:

- Windows
`LoadModule ibm_ldap_module modules/IBMModuleLDAP.dll`
- AIX HP-UX Linux Solaris
`LoadModule ibm_ldap_module modules/mod_ibm_ldap.so`

If you have the LDAP client installed on your computer, you can use `ldapsearch` as a tool to test the values you intend to use for the various settings.

About this task

See “LDAP directives” on page 165 to obtain detailed descriptions of the LDAP (`mod_ibm_ldap`) directives.

Procedure

1. Edit the `httpd.conf` IBM HTTP Server configuration file.
2. Determine the resource you want to limit access to. For example: `<Directory "/secure_info">`.
3. Add directives in `httpd.conf` to the directory location (container) to be protected with values specific to your environment. For example:
 - `LdapConfigFile path_to_ldap.prop`
 - `AuthType Basic`
 - `AuthName "Title of your protected Realm"`
 - `Require valid-user`
4. There are three options for how to use IBM HTTP Server to authenticate with your existing LDAP installation.
 - **Authorization based on LDAP group membership.**

Use LDAP to check user passwords and verify that the user exists in a group defined in LDAP.

Note: The membership that identifies the user as being able to access the resource is a part of the group, not part of the user's own LDAP entry.

For example, to restrict access to a group, add the following directive:

```
LDAPRequire group grp1
```

For this form of `LDAPRequire`, you must have groups configured in your LDAP repository that conform to the following rules (using the example group name `grp1`):

- There is an entry in your LDAP repository that matches the following search filter, where the values `groupofnames` and `groupofuniquenames` are example values specified in `ldap.group.dnattributes`.

Note: The proper value of `ldap.group.dnattributes` is a list of what objectclasses signify is a group in your LDAP schema.

```
ldapsearch ... "(&(cn=grp1)(|(objectclass=groupofnames)
(objectclass=groupofuniquenames)))"
```

- As part of the LDAP entry for "`grp1`," there are a series of attributes that match the following, where the values `member` and `uniquemember` are example values of `ldap.group.memberAttributes`.

Note: The proper value of `ldap.group.memberAttributes` is a list of what objectclasses signify is a membership in a group. The values of these entries are the Distinguished Names (DN) of your users.

```
ldapsearch ... "(&(cn=grp1)(|(objectclass=groupofnames)
(objectclass=groupofuniquenames)))" member uniquemember
```

Example:

```
ldapsearch -x -h myldapsrv -D cn=root -w rootpw
"(&(cn=grp1)(|(objectclass=groupofnames)(objectclass=groupofuniquenames)))"
member uniquemember
```

```
dn: cn=group1,ou=myunit,o=myorg,c=US
member: cn=user1,ou=otherunit,o=myorg,c=US
member: cn=user12,ou=otherunit,o=myorg,c=US
```

If an object of the type listed in `ldap.group.dnattributes` is a member of the group being searched, then it will be recursively searched in the same fashion, up to a depth of `ldap.group.search.depth`

- First IBM HTTP Server uses the `ldap.group.name.filter` and `ldap.user.cert.filter` to translate the CN provided for the user and the group into distinguished names (DN). Next, IBM HTTP Server searches using the group DN as a base for entries whose value is the user DN.

Example:

```
ldapsearch ... -b "cn=grp1,ou=myunit,o=myorg,c=US"
"|((member=cn=user1,ou=otherunit,o=myorg,c=US)
(uniquemember=cn=user1,ou=otherunit,o=myorg,c=US))"
```

- **Authorization based on LDAP attributes of user.** Use LDAP to check user passwords and verify that the user matches a set of attributes (the attribute that identifies the user as being able to access the resource is a part of the users own LDAP entry).

Example:

```
LDAPRequire filter "(&(jobtitle=accountant)(location=newyork))"
```

To use this form of LDAPRequire, the IBM HTTP Server must use the `ldap.user.cert.filter` to translate the CN provided for the user into a DN. IBM HTTP Server must also search using the user DN as a base and use the search filter provided in the LDAPRequire directive. If any results are returned, authorization succeeds.

Example:

```
ldapsearch ... -b "cn=user1,ou=otherunit,o=myorg,c=US" "(&(jobtitle=accountant)
(location=newyork))"
```

Some attributes (sometimes called Dynamic Roles) in LDAP are calculated dynamically by the LDAP server and might have different semantics that are not valid in a search filter. Such values would fail if used in the preceding example and cannot be used for authorization on IBM HTTP Server.

- **Authentication only: Use LDAP to check user passwords only.**

You can use the Require directive to limit to specific users or maintain a flat group file using AuthGroupFile.

5. Edit the `ldap.prop` configuration file. If you do not have one yet, you can use the `ldap.prop.sample` file that ships with IBM HTTP Server. If you have questions about the correct values, check with your LDAP server administrator. Update the following directives with values that are correct for your environment:

- a. Enter the Web server connection information.
- b. When using SSL, or LDAPS, or LDAP over SSL:
 - Change `ldap.transport` to an SSL value
 - Change `ldap.URL` to include the LDAPS protocol and the proper port value, for example, 636.
 - Configure the SSL key database to be used, for example:

```
ldap.key.fileName=/path_to/key.kdb
ldap.key.file.password.stashFile=/path_to/stashfile
```

Where *stashfile* is created by the `bin/ldapstash` command.

```
ldap.key.label=label
```

Where *label* is the value appearing in IKEYMAN for the referenced *key.kdb*.

Results

Searches that use the `mod_ibm_ldap` directives maintain a pool of server connections that authenticate as the `ldap.application.dn` user. The first connection is created when the first LDAP-protected request is received. Connections will be held open a specified number of seconds (`ldap.idleConnection.timeout`) for subsequent searches on that connection or connections for other requests.

If you are reading logs or looking at an IP trace, the following sequence of events should occur:

- IBM HTTP Server starts.
- If `LDAP_TRACE_FILE` is set, it will have a few entries for `LDAP_obtain_config`
- The first request for LDAP-protected resource is received.

- IBM HTTP Server binds to LDAP using the `ldap.application.dn` username and the password stashed in `ldap.application.password.stashFile` (Application Connection)
- IBM HTTP Server performs a search over this connection to translate the username typed in by the user, or the contents of their client certificate, into a Distinguished Name (DN) using the `user.*.filter` settings.
- IBM HTTP Server binds to the LDAP server as username/password provided by the client to check authentication (This is a "user connection" to the LDAP server)
- If any LDAPRequire directives are in effect for this request, IBM HTTP Server processes them in the manner described in the preceding procedure.
- IBM HTTP Server unbinds the user connection
- The application connection is maintained for the next request

Lightweight Directory Access Protocol

This section addresses questions about what Lightweight Directory Access Protocol (LDAP) is and how it works, and provides high level overviews of X.500 and LDAP.

LDAP is a standard protocol that provides a means of storing and retrieving information about people, groups, or objects on a centralized X.500 or LDAP directory server. X.500 enables that information to be organized and queried, using LDAP, from multiple web servers using a variety of attributes. LDAP queries can be as simple or complex as is required to identify a desired individual entity or group of entities. LDAP reduces required system resources by including only a functional subset of the original X.500 Directory Access Protocol (DAP).

The IBM HTTP Server LDAP module enables the use of an X.500 directory server for authentication and authorization purposes. IBM HTTP Server can use this capability to limit access of a resource to a controlled set of users. This capability reduces the administrative overhead usually required to maintain user and group information for each individual Web server.

You can configure the IBM HTTP Server LDAP module to use TCP/IP or Secure Sockets Layer (SSL) connections to the X.500 directory server. The LDAP module can be configured to reference a single LDAP server or multiple servers.

X.500 overview. X.500 provides a directory service with components that are capable of more efficient retrieval. LDAP uses two of these components: The information model, which determines the form and character, and the namespace, which enables information indexing and referencing.

The X.500 directory structure differs from other directories in information storage and retrieval. This directory service associates information with attributes. A query based on attributes generates and passes to the LDAP server, and the server returns the respective values. LDAP uses a simple, string-based approach for representing directory entries.

An X.500 directory consists of typed entries that are based on the ObjectClass attribute. Each entry consists of attributes. The ObjectClass attribute identifies the type of entry, for example, a person or organization, that determines the required and optional attributes.

You can divide entries, arranged in a tree structure, among servers in geographical and organizational distribution. The directory service names entries, according to their position within the distribution hierarchy, by a distinguished name (DN).

Lightweight Directory Access Protocol overview. Accessing an X.500 directory requires the Directory Access Protocol (DAP). However, DAP requires large amounts of system resources and support mechanisms to handle the complexity of the protocol. To enable desktop workstations to access the X.500 directory service, LDAP was introduced.

LDAP, a client and server-based protocol can handle some of the heavy resources required by DAP clients. An LDAP server can only return results or errors to the client, requiring little from the client. If unable to answer a client request, an LDAP Server must chain the request to another X.500 server. The server must complete the request, or return an error to the LDAP server, which in turn passes the information to the client.

IBM HTTP Server supports the following LDAP servers:

- iPlanet/Netscape Directory Server
- IBM SecureWay™ Directory Server
- Microsoft Active Directory

Querying the Lightweight Directory Access Protocol server

The Lightweight Directory Access Protocol (LDAP) accesses the X.500 directory using text strings called filters. When these query strings pass to the LDAP server, the server returns the requested portions of the specified entity.

About this task

LDAP filters use attributes to simplify queries to the LDAP server. For example, you can use a filter such as "objectclass=person" to limit your query to entities that represent people as opposed to groups or equipment.

Procedure

- To authorize a user as a member of a group, add the following directive to the configuration file:

```
LDAPRequire group "group_name"
```

For example:

```
LDAPRequire group "Administrative Users"
```

- To authorize a user by filter, add the following directive to the configuration file:

```
LDAPRequire filter "ldap_search_filter"
```

For example, to enable access to the resource by a programmer in your department:

```
LDAPRequire filter("&(objectclass=person)(cn=*)(ou=programmer)(o=department))"
```

Or, to enable access for John Doe only:

```
LDAPRequire filter "&(objectclass=person)(cn=John Doe)"
```

Secure Sockets Layer and the Lightweight Directory Access Protocol module

IBM HTTP Server provides the ability to use a secure connection between the LDAP module running in the Web server and the LDAP directory server. If this feature is enabled, any communication between the Web server and the directory server is encrypted.

To enable this feature, edit the `ldap.prop` LDAP configuration file and change the value of `ldap.transport` to `SSL`. Create or obtain a certificate database file (`X.kdb`) and a password stash file (`Y.sth`). You can use `IKEYMAN` to obtain a key database file. You must use the **ldapstash** program to create the stash file. You will also need to change the values for `ldap.URL` and `ldap.group.URL` to use port 636 instead of port 389.

The key database file contains the certificates which establish identity. The LDAP server can require that the Web server provide a certificate before allowing queries. When using a certificate with an SSL connection between the LDAP module and the LDAP server, the user ID that IBM HTTP Server is configured to use must have write permission to the key database file containing the certificate.

Certificates establish identity to prevent other users from stealing or overwriting your certificates (and therefore your identity). If someone has read permission to the key database file, they can retrieve the user's certificates and masquerade as that user. Grant read or write permission only to the owner of the key database file.

SSL certificate revocation list

This section provides information on identifying directives for certificate revocation list (CRL) and those supported in global servers and virtual hosts.

Certificate revocation provides the ability to revoke a client certificate given to IBM HTTP Server by the browser when the key becomes compromised or when access permission to the key gets revoked. CRL represents a database which contains a list of certificates revoked before their scheduled expiration date.

If you want to enable certificate revocation in IBM HTTP Server, publish the CRL on a Lightweight Directory Access Protocol (LDAP) server. Once the CRL is published to an LDAP server, you can access the CRL using the IBM HTTP Server configuration file. The CRL determines the access permission status of the requested client certificate. Be aware, however, that it's not always possible to determine the revocation status of a client certificate if the backend server, the source of revocation data, is not available or not communicating properly with IBM HTTP Server.

Identifying directives needed to set up a certificate revocation list. The SSLClientAuth directive can include two options at once:

- SSLClientAuth 2 crl
- SSLClientAuth 1 crl

The CRL option turns CRL on and off inside an SSL virtual host. If you specify CRL as an option, then you elect to turn CRL on. If you do not specify CRL as an option, then CRL remains off. If the first option for SSLClientAuth equals 0/none, then you cannot use the second option, CRL. If you do not have client authentication on, then CRL processing does not take place.

Identifying directives supported in global or server and virtual host. Global server and virtual host support the following directives:

- SSLCRLHostname: The IP Address and host of the LDAP server, where the CRL database resides. Currently, you must configure any static CRL repositories to allow for checking of other URI forms in the CRLDistributionPoint fields.

z/OS

Only an explicitly configured LDAP server can be queried for CRL, and the SSL handshake fails if the backend server is not reachable.

- SSLCRLPort: The port of the LDAP server where the CRL database resides; the default equals 389.
- SSLCRLUserID: The user ID to send to the LDAP server where the CRL database resides; defaults to anonymous if you do not specify the bind.
- SSLStashfile: The fully qualified path to file where the password for the user name on the LDAP server resides. This directive is not required for an anonymous bind. Use when you specify a user ID.

Use the **sslstash** command, located in the bin directory of IBM HTTP Server, to create your CRL password stash file. The password you specify using the **sslstash** command should equal the one you use to log in to your LDAP server.

Usage:

```
sslstash [-c] &lt;directory_to_password_file_and_file_name>; <function_name> <password>
```

where:

- **-c**: Creates a new stash file. If not specified, an existing file updates.
- **File**: Represents the fully qualified name of the file to create, or update.
- **Function**: Indicates the function for which to use the password. Valid values include crl, or crypto.

- **Password:** Represents the password to stash.
- **Distributed operating systems** SSLUnknownRevocationStatus: This directive allows you to configure how IBM HTTP Server will respond when fresh Certificate Revocation List (CRL) information or OCSP (Online Certificate Status Protocol) information is not available, and the client certificate that is currently offered is not known to be revoked from a previous query. Certificates are presumed not to be revoked, by default, which means they are valid, and a temporary failure to obtain CRL or OCSP information does not automatically result in an SSL handshake failure. This directive is provided to respond to circumstances in which a certificate has been accepted without IBM HTTP Server being able to reliably confirm the revocation status.

This directive has an effect only when all of these conditions are true:

- IBM HTTP Server is configured to accept client certificates with the SSLClientAuth directive.
- IBM HTTP Server is configured with one of the following directives: SSLOCSPEnable, SSLOCSPUrl, or SSLCRLHostname.
- An SSL client certificate is provided.
- IBM HTTP Server does not receive a valid OCSP or CRL response from the configured backend server, and the client certificate does not appear as revoked in a cached, but expired, CRL response. IBM HTTP Server uses a cached CRL that is beyond its published expiration time when a current version is not available. When a certificate has been revoked in such an expired CRL, this will result in a direct SSL handshake failure that is outside the scope of the SSLUnknownRevocationStatus directive.

See the “SSL directives” on page 103 topic for more information.

CRL checking follows the URIDistributionPoint X509 extension in the client certificate as well as trying the DN constructed from the issuer of the client certificate. If the certificate contains a CRL Distribution Point (CDP), then that information is given precedence. The order in which the information is used is as follows:

1. CDP LDAP X.500 name
2. CDP LDAP URI
3. Issuer name combined with the value from the SSLCRLHostname directive

gotcha: If your certificates use the LDAP or HTTP URI forms of the CertificateDistributionPoint or AIA extensions, be sure that the IBM HTTP Server system can establish outgoing connections of this type; you might need to adjust the settings for your firewall.

LDAP directives

These configuration parameters control the Lightweight Directory Access Protocol (LDAP) feature in IBM HTTP Server.

- “LdapCodepageDir directive” on page 166
- “LdapConfigfile directive” on page 166
- “LDAPRequire directive” on page 166
- “Ldap.application.authType directive” on page 167
- “Ldap.application.DN directive” on page 167
- “Ldap.application.password.stashFile directive” on page 167
- “Ldap.cache.timeout directive” on page 168
- “Ldap.group.attribute directive” on page 168
- “Ldap.group.dnattribute directive” on page 168
- “Ldap.group.memberattribute directive” on page 168
- “Ldap.group.memberAttributes directive” on page 169
- “Ldap.group.name.filter directive” on page 169
- “Ldap.group.search.depth directive” on page 169
- “Ldap.group.URL directive” on page 170
- “Ldap.idleConnection.timeout directive” on page 170
- “Ldap.key.file.password.stashfile directive” on page 171

- “Ldap.key.fileName directive” on page 171
- “Ldap.key.label directive” on page 171
- “LdapReferralhoplimit directive” on page 171
- “LdapReferrals directive” on page 172
- “Ldap.realm directive” on page 172
- “Ldap.search.timeout directive” on page 172
- “Ldap.transport directive” on page 172
- “Ldap.url directive” on page 173
- “Ldap.user.authType directive” on page 173
- “Ldap.user.cert.filter directive” on page 173
- “Ldap.user.name.fieldSep directive” on page 174
- “Ldap.user.name.filter directive” on page 174
- “Ldap.version directive” on page 175
- “Ldap.waitForRetryConnection.interval directive” on page 175

Note: **Distributed operating systems** If you are using the mod_ibm_ldap module for your LDAP configuration, consider migrating your mod_ibm_ldap directives to use the mod_ldap module. The mod_ibm_ldap module is provided with this release of IBM HTTP Server for compatibility with previous releases, however, you must migrate existing configurations to use the mod_authnz_ldap and mod_ldap modules to ensure future support for your LDAP configuration.

LdapCodepageDir directive

Codepages are now automatically installed in the IHS installation directory and are referenced relative to the IHS installation directory, as opposed to the configured server root directory as in previous versions.

LdapConfigfile directive

The LdapConfigFile directive indicates the name of the LDAP properties file associated with a group of LDAP parameters.

Syntax	LdapConfigFile <Fully qualified path to configuration file>
Scope	Single instance per directory stanza
Default	c:\program files\ibm http server\conf\ldap.prop.sample
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	Fully qualified path to a single configuration file. Use this directive in the httpd.conf file.

LDAPRequire directive

The LDAPRequire directive is used to restrict access to a resource that is controlled by LDAP authentication to a specified collection of users. It can either use groups that are defined in LDAP by using the group type, or it can use an LDAP filter type to designate a collection of users with a similar set of attribute values.

Syntax	LDAPRequire filter <filter name> or LDAPRequire group <group1 [group2.group3...]>
Scope	Single instance per directory stanza
Default	None
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes

Values

LDAPRequire filter "(
 &(objectclass=person)(cn=*)(ou=IHS)(o=IBM))", or
 LDAPRequire group "sample group".

Use this directive in the httpd.conf file.

If the group type is used, and multiple group values are specified, the group validation is a logical AND of the groups. A user must be a member of *sample Group1* and *sample Group2* if a logical OR of groups is required. For example, if a user is a member of *sample Group1* or *sample Group2*, then a new LDAP group, *our department group*, should be created on the LDAP server that has *sample Group1* and *sample Group2* as its members. You would then use the directive: LDAPRequire group *our Department Group* .

Ldap.application.authType directive

The Ldap.application.authType directive specifies the method for authenticating the Web server to the LDAP server.

Syntax

ldap.application.authType=None

Scope

Single instance per directory stanza

Default

None

Module

mod_ibm_ldap

Multiple instances in the configuration file

yes

Values

- None: If the LDAP server does not require the Web server to authenticate.
- Basic: Uses the distinguished name (DN) of the Web server as the user ID, and the password stored in the stash file, as the password.

Ldap.application.DN directive

The Ldap.application.DN directive indicates the distinguished name (DN) of the Web server. Use this name as the user name when accessing an LDAP server using basic authentication. Use the entry specified in the LDAP server to access the directory server.

Syntax

ldap.application.DN=cn=ldapadm,ou=ihs
 test,o=IBM,c=US

Scope

Single instance per directory stanza

Default

None

Module

mod_ibm_ldap

Multiple instances in the configuration file

yes

Values

Distinguished name

Ldap.application.password.stashFile directive

The Ldap.application.password.stashFile directive indicates the name of the stash file containing the encrypted password for the application to authenticate to the LDAP server when Server Authentication type is Basic.

Syntax

ldap.application.password.stashFile=c:\IHS\ldap.sth

Scope

Single instance per directory stanza

Default

None

Module

mod_ibm_ldap

Multiple instances in the configuration file

yes

Values Fully qualified path to the stash file. You can create this stash file with the **ldapstash** command.

Ldap.cache.timeout directive

The ldap.cache.timeout directive caches responses from the LDAP server. If you configure the Web server to run as multiple processes, each process manages its own copy of the cache.

Syntax	ldap.cache.timeout= <secs>
Scope	Single instance per directory stanza
Default	600
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	The maximum length of time, in seconds, a response returned from the LDAP server remains valid.

Ldap.group.attribute directive

The ldap.group.attributes directive indicates the filter used to determine if a distinguished name (DN) is an actual group through an LDAP search.

Syntax	ldap.group.memberattribute = <attribute>
Scope	Single instance per directory stanza
Default	uniquegroup
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	An ldap attribute - See the ldap.prop.sample directive for more information on the use of this directive.

Ldap.group.dnattribute directive

The ldap.group.dnattributes specifies the filter used to determine, through an LDAP search, if a distinguished name (DN) is an actual group.

Syntax	ldap.group.memberattribute = <ldap filter>
Scope	Single instance per directory stanza
Default	groupofnames groupofuniquenames
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	An ldap filter - See the ldap.prop.sample directive for more information on the use of this directive.

Ldap.group.memberattribute directive

The ldap.group.memberattribute directive specifies the attribute to retrieve unique groups from an existing group.

Syntax	ldap.group.memberattribute = <ldap filter>
Scope	Single instance per directory stanza
Default	groupofnames groupofuniquenames
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes

Values

An ldap filter - See the `ldap.prop.sample` directive for more information on the use of this directive.

Ldap.group.memberAttributes directive

The `ldap.group.memberAttributes` directive serves as a means to extract group members, once the function finds a group entry in an LDAP directory.

Syntax

```
ldap.group.memberAttributes= attribute  
[attribute2...]
```

Scope

Single instance per directory stanza

Default

member and uniquemember

Module

mod_ldap

Multiple instances in the configuration file

yes

Values

Must equal the distinguished names of the group members. You can use more than one attribute to contain member information.

Ldap.group.name.filter directive

The `ldap.group.name.filter` directive indicates the filter LDAP uses to search for group names.

Syntax

```
ldap.group.name.filter = <group name filter>
```

Scope

Single instance per directory stanza

Default

```
(&(cn=%v1) (|(objectclass=groupOfNames)  
(objectclass=groupOfUniqueNames))
```

Module

mod_ldap

Multiple instances in the configuration file

yes

Values

An LDAP filter. See Querying the LDAP server using LDAP search filters.

Ldap.group.search.depth directive

The `ldap.group.search.depth` directive searches subgroups when specifying the `LDAPRequire group <group>` directives. Groups can contain both individual members and other groups.

Syntax

```
ldap.group.search.depth = <integer depth>
```

Scope

Single instance per directory stanza

Default

1

Module

mod_ldap

Multiple instances in the configuration file

yes

Values

An integer. When doing a search for a group, if a member in the process of authentication is not a member of the required group, any subgroups of the required group are also searched. For example:

```
group1 >group2 (group2 is a member of group1)
group2 >group3 (group3 is a member of group2)
group3 >jane (jane is a member of group3)
```

If you search for jane and require her as a member of group1, the search fails with the default `ldap.search.depth` value of 1. If you specify `ldap.group.search.depth>2`, the search succeeds.

Use `ldap.group.search.depth=<depth to search -- number>` to limit the depth of subgroup searches. This type of search can become very intensive on an LDAP server. Where group1 has group2 as a member, and group2 has group1 as a member, this directive limits the depth of the search. In the previous example, group1 has a depth of 1, group2 has a depth of 2 and group3 has a depth of 3.

Ldap.group.URL directive

The `ldap.group.URL` directive specifies a different location for a group on the same LDAP server. You cannot use this directive to specify a different LDAP server from that specified in the `ldap.URL` directive.

Syntax

```
ldap.group.URL = ldap://<hostname:port>/<BaseDN>
```

Scope

Single instance per directory stanza

Default

None

Module

`mod_ibm_ldap`

Multiple instances in the configuration file

yes

Values

- host name: Host name of the LDAP server.
- port number: Optional port number on which the LDAP server listens. The default for TCP connections is 389. If you use SSL, you must specify the port number.
- BaseDN: Provides the root of the LDAP tree in which to perform the search for groups.

Attention:

This property becomes required if the LDAP URL for groups differs from the URL specified by the `ldap.URL` property.

Ldap.idleConnection.timeout directive

The `ldap.idleConnection.timeout` directive caches connections to the LDAP server for performance.

Syntax

```
ldap.idleConnection.timeout = <secs>
```

Scope

Single instance per directory stanza

Default

600

Module

`mod_ibm_ldap`

Multiple instances in the configuration file

yes

Values

Length of time, in seconds, before an idle LDAP server connection closes due to inactivity.

Ldap.key.file.password.stashfile directive

The `Ldap.key.file.password.stashfile` directive indicates the stash file containing the encrypted keyfile password; use the `ldapstash` command to create this stash file.

Syntax	<code>ldap.key.file.password.stashfile =d:\ <Key password file name></code>
Scope	Single instance per directory stanza
Default	None
Module	<code>mod_ibm_ldap</code>
Multiple instances in the configuration file	yes
Values	Fully qualified path to the stash file.

Ldap.key.fileName directive

The `Ldap.key.fileName` directive indicates the file name of the key file database. This option becomes required when you use Secure Sockets Layer (SSL).

Syntax	<code>ldap.key.fileName=d:\<Key file name></code>
Scope	Single instance per directory stanza
Default	None
Module	<code>mod_ibm_ldap</code>
Multiple instances in the configuration file	yes
Values	Fully qualified path to the key file.

Ldap.key.label directive

The `Ldap.key.file.password.stashfile` directive indicates the certificate label name the Web server uses to authenticate to the LDAP server.

Syntax	My Server Certificate
Scope	Single instance per directory stanza
Default	None
Module	<code>mod_ibm_ldap</code>
Multiple instances in the configuration file	yes
Values	A valid label used in the key database file. This label becomes required only when using Secure Sockets Layer (SSL) and the LDAP server requests client authentication from the Web server.

LdapReferralHopLimit directive

The `LdapReferralHopLimit` directive indicates the maximum number of referrals to follow. LDAP authentication will fail if the specified limit is exceeded.

Syntax	<code>LdapReferralHopLimit = <number_of_hops></code>
Scope	Single instance per directory stanza
Default	10
Module	<code>mod_ibm_ldap</code>
Multiple instances in the configuration file	yes
Values	0 to 10

Set the `LdapReferrals` directive on to use the `LdapReferralHopLimit` directive.

Important: An LdapReferralhoplimit value of 0 will cause authentication to fail if any referrals are encountered.

The LdapReferralhoplimit directive is not meaningful when the LdapReferrals directive is off (default).

LdapReferrals directive

The LdapReferrals directive indicates whether referrals (which redirect a client request to another LDAP server) will be chased for searches while performing LDAP queries.

Syntax	LdapReferrals = off on
Scope	Single instance per directory stanza
Default	off
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	On or off

Ldap.realm directive

The ldap.key.realm directive indicates the name of the protected area, as seen by the requesting client.

Syntax	ldap.realm=<Protection Realm>
Scope	Single instance per directory stanza
Default	None
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	A description describing the protected page.

Ldap.search.timeout directive

The ldap.search.timeout directive indicates the maximum time, in seconds, to wait for an LDAP server to complete a search operation.

Syntax	ldap.search.timeout = <secs>
Scope	Single instance per directory stanza
Default	10
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	Length of time, in seconds.

Ldap.transport directive

The ldap.transport directive indicates the transport method used to communicate with the LDAP server.

Syntax	ldap.transport = TCP
Scope	Single instance per directory stanza
Default	TCP
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	TCP or SSL

Ldap.url directive

The `Ldap.url` directive indicates the URL of the LDAP server to authenticate against.

Syntax

```
ldap.url = ldap://<hostname:port>/<BaseDN>
```

where:

- `hostname`: Represents the host name of the LDAP server.
- `port`: Represents the optional port number on which the LDAP server listens. The default for TCP connections is 389. You must specify the port number if you use SSL.
- `BaseDN`: Provides the root of the LDAP tree in which to perform the search for users.

For example: `ldap.url=ldap://ldap.ibm.com:489/o=Ace Industry, c=US>`

Scope

Single instance per directory stanza

Default

None

Module

`mod_ldap`

Multiple instances in the configuration file

yes

Ldap.user.authType directive

The `Ldap.user.authType` directive indicates the method for authenticating the user requesting a Web server. Use this name as the user name when accessing an LDAP server.

Syntax

```
ldap.user.authType = BasicIfNoCert
```

Scope

Single instance per directory stanza

Default

Basic

Module

`mod_ldap`

Multiple instances in the configuration file

yes

Values

Basic, Cert, BasicIfNoCert

Ldap.user.cert.filter directive

The `Ldap.user.cert.filter` directive indicates the filter used to convert the information in the client certificate passed over Secure Sockets Layer (SSL) to a search filter for an LDAP entry.

Syntax

```
ldap.user.cert.filter=(&(objectclass=person)(cn=%v1))
```

Scope

Single instance per directory stanza

Default

```
"(&(objectclass=person)(cn=%v1, ou=%v2, o=%v3,c=%v4))"
```

Module

`mod_ldap`

Multiple instances in the configuration file

yes

Values

An LDAP filter. See *Querying the LDAP server using LDAP search filters*.

Secure Socket Layer (SSL) certificates include the following fields, all of which you can convert to a search filter:

Certificate field	Variable
common name	<code>%v1</code>
organizational unit	<code>%v2</code>

organization	%v3
country	%v4
locality	%v5
state or country	%v6
serial number	%v7

When you generate the search filter, you can find the field values in the matching variable fields (%v1, %v2). The following table shows the conversion:

User certificate	Filter conversion
Certificate	cn=Road Runner, o=Acme Inc, c=US
Filter	(cn=%v1, o=%v3, c=%v4)
Resulting query	(cn=RoadRunner, o=Acme, Inc, c=US)

Ldap.user.name.fieldSep directive

The ldap.usr.name.fieldSep directive indicates characters as valid field separator characters when parsing the user name into fields.

Syntax	ldap.user.name.fieldSep=/ /
Scope	Single instance per directory stanza
Default	The space, comma, and the tab (/t) character.
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	Characters. If '/' represents the only field separator character and the user enters "Joe Smith/Acme", then '%v2' equals "Acme".

Ldap.user.name.filter directive

The ldap.usr.name.filter directive indicates the filter used to convert the user name entered in a search filter for an LDAP entry.

Syntax	ldap.user.name.filter=<user name filter>
Scope	Single instance per directory stanza
Default	"((objectclass=person) (cn=%v1 %v2))", where %v1 and %v2 represent characters entered by the user.

For example, if the user enters "Paul Kelsey", the resulting search filter becomes "((objectclass=person) (cn=Paul Kelsey))". You can find search filter syntax described in Querying the LDAP server using LDAP search filters.

However, because the Web server cannot differentiate between multiple returned entries, authentication fails when the LDAP server returns more than one entry. For example, if the user makes the ldap.user.name.filter="((objectclass=person) (cn=%v1* %v2*))" and enters **Pa Kel**, the resulting search filter becomes "(cn=Pa* Kel*)". The filter finds multiple entries such as (cn=Paul Kelsey) and (cn=Paula Kelly) and authentication fails. You must modify your search filter.

Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	An LDAP filter. See Querying the LDAP server using LDAP search filters.

Ldap.version directive

The ldap.version directive indicates the version of the LDAP protocol used to connect to the LDAP server. the protocol version used by the LDAP server determines the LDAP version.

Attention: This directive is optional.

Syntax	ldap.version=3
Scope	Single instance per directory stanza
Default	ldap.version=3
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	2 or 3

Ldap.waitForRetryConnection.interval directive

The ldap.waitForRetryConnection.interval directive indicates the time the Web server waits between failed attempts to connect.

If an LDAP server goes down, the Web server continues to try to connect.

Syntax	ldap.waitForRetryConnection.interval=<secs>
Scope	Single instance per directory stanza
Default	300
Module	mod_ibm_ldap
Multiple instances in the configuration file	yes
Values	Time (in seconds)

Converting your directives from mod_ibm_ldap to mod_ldap

Convert directives that use the mod_ibm_ldap module to use the mod_ldap Apache module to ensure continued IBM HTTP Server support for your LDAP configuration.

Before you begin

Determine which directives to convert.

Complete these steps to convert your directives.

Procedure

1. Edit the LoadModule directive in the httpd.conf or ldap.prop configuration file to remove mod_ibm_ldap.

```
LoadModule ibm_ldap_module modules/mod_ibm_ldap.so
```
2. Add the mod_ldap LoadModule directive to the httpd.conf configuration file.

```
LoadModule authnz_ldap_module modules/mod_authnz_ldap.so
LoadModule ldap_module modules/mod_ldap.so
```
3. Convert one or more of the following directives. For more information about converting your directives, see the topic about mod_ibm_ldap migration.

Note: A one to one correlation might not exist for some directives.

Table 15. LDAP configuration directives conversion

mod_ibm_ldap	mod_ldap
ldapCodePageDir	None. The codepages directory cannot be moved from its installed location.
LdapConfigFile	include
LdapRequire	require
ldap.application.authType	None. If the mod_ldap directive, AuthLDAPBindDN, is specified, then you will get Basic auth. If no AuthLDAPBindDN is specified, then you get what would have been the None auth type (anonymous). If the mod_ldap configuration specifies an LDAPTrustedClientCert value then you will get the Cert auth type.
ldap.application.DN	AuthLDAPBindDN
ldap.application.password	AuthLDAPBindPassword
ldap.application.password.stashFile	None. The mod_ldap module does not provide a directive for using stashed passwords.
ldap.cache.timeout	LDAPCacheTTL
ldap.group.dnattributes	AuthLDAPSubGroupClass
ldap.group.memberattribute	AuthLDAPSubGroupAttribute
ldap.group.memberattributes	AuthLDAPGroupAttribute
ldap.group.name.filter	None. The mod_ldap module uses the filter provided at the end of the AuthLDAPURL directive.
ldap.group.search.depth	AuthLDAPMaxSubGroupDepth
ldap.group.URL	AuthLDAPURL
ldap.idleConnection.timeout	None. The mod_ldap module does not provide a directive for connection timeouts.
ldap.key.file.password.stashfile	None. The mod_ldap module does not provide a directive for using stashed passwords. Specify the keyfile password, in clear text, at the end of the LDAPTrustedGlobalCert directive. Alternatively, omit the password on the LDAPTrustedGlobalCert directive and the mod_ldap module automatically looks for a /path/to/keyfile.sth file, assuming /path/to/keyfile.kdb was the specified value of the LDAPTrustedGlobalCert directive.
ldap.key.fileName	LDAPTrustedGlobalCert
ldap.key.label	LDAPTrustedClientCert
ldap.ReferralHopLimit	LDAPReferralHopLimit
ldapReferrals	LDAPReferrals
ldap.realm	None. The mod_ibm_ldap value of this directive was only used for logging purposes. No equivalent directive is required in mod_ldap.
ldap.search.timeout	LDAPSearchTimeout
ldap.transport	LDAPTrustedMode
ldap.URL	AuthLDAPURL

Table 15. LDAP configuration directives conversion (continued)

mod_ibm_ldap	mod_ldap
ldap.user.authType	None. The mod_ldap module authenticates users based on the user ID and password credentials provided.
ldap.user.cert.filter	None. The mod_ldap module does not work directly with client certificates. Authorization directives use the environment values set by the SSL module.
ldap.user.name.fieldSep	None. The mod_ldap module does not provide support for parsing the provided credentials into subcomponents.
ldap.user.name.filter	None. The mod_ldap module specifies the user name filter as part of the AuthLDAPURL directive.
ldap.version	None. The mod_ldap module uses only LDAP version 3.
ldap.waitForRetryConnection.interval	None. The mod_ldap module does not have a timed delay between connection retries when a connection attempt fails. The connection attempt is retried for a maximum of 10 times before request fails.

4. Run the Apache control with the verify flag to verify the configuration.

```
<ihsinst>bin/apachectl -t
```

Attention: This configuration check confirms that the syntax is correct, but you must verify any configuration changes for a directive using the documentation for that directive to ensure an optimal configuration.

Attention: All mod_ibm_ldap directives that use the form ldap.* used to optionally display in the LDAPConfigFile configuration file without the ldap prefix.

A mod_ldap SSL configuration

The following configuration directives show a sample SSL-enabled LDAP configuration. Some of the directives specify default values and would not typically need to be specified, but are retained to provide context. Those directives are included, but are commented out with "##" symbols.

```
##LDAPReferrals On
##LDAPReferralHopLimit 5

LDAPTrustedGlobalCert CMS_KEYFILE /full/path/to/ldap_client.kdb clientkdbPassword
#default cert in this kdb is my_cert1

# Alternatively, you can specify a SAF-based keyring, on systems that support it, as follows:
#LDAPTrustedGlobalCert SAF saf_keyring

<VirtualHost *>
  ServerAdmin admin@my.address.com
  DocumentRoot /path/to/htdocs

  # Ignored because LDAP URLs use ldaps:, where needed
  ##LDAPTrustedMode SSL

  <Directory /minimal_ldap_config>
    AuthBasicProvider ldap
    AuthLDAPURL ldap://our_ldap.server.org/o=OurOrg,c=US
    AuthName "Private root access"
    require valid-user
  </Directory>

  <Directory /path/to/htdocs>
    ##AuthzLDAPAuthoritative on
    AuthBasicProvider ldap
```

```

# This LDAPTrustedClientCert is required to use a different certificate
# than the default
LDAPTrustedClientCert CMS_LABEL my_cert2
AuthLDAPURL ldaps://our_ldap.server.org:636/o=OurOrg,c=US?cn?sub? (objectclass=person)
AuthLDAPBindDN "cn=ldapadm,ou=OurDirectory,o=OurCompany,c=US"
AuthLDAPBindPassword mypassword
AuthName "Private root access"
require ldap-group cn=OurDepartment,o=OurOrg,c=us
</Directory>

<Directory "/path/to/htdocs/employee_of_the_month">
##AuthzLDAPAuthoritative on
AuthBasicProvider ldap
#Uses default cert (my_cert1)
##LDAPTrustedClientCert CMS_LABEL my_cert1
AuthLDAPURL ldaps://our_ldap.server.org:636/o=OurOrg,c=US?cn?sub?(objectclass=person)
AuthLDAPBindDN "cn=ldapadm,ou=OurDirectory,o=OurCompany,c=US"
AuthLDAPBindPassword mypassword
AuthName "Employee of the month login"
require ldap-attribute description="Employee of the Month."
</Directory>

<Directory "/path/to/htdocs/development_groups">

#These are the default values for the subgroup-related directives and only need to be
#specified when the LDAP structure differs.
##AuthzLDAPAuthoritative on
AuthBasicProvider ldap
# This LDAPTrustedClientCert is required to use a different certificate
# than the default LDAPTrustedClientCert CMS_LABEL my_cert3
AuthLDAPURL ldaps://groups_ldap.server.org:636/o=OurOrg,c=US?cn?sub?
(|(objectclass=groupofnames)(objectclass=groupofuniqueNames))
AuthLDAPBindDN "cn=ldapadm,ou=OurDirectory,o=OurCompany,c=US"
AuthLDAPBindPassword mypassword
AuthName "Developer Access"
AuthLDAPGroupAttribute member
AuthLDAPMaxSubGroupDepth 2
AuthLDAPSubGroupClass groupOfUniqueNames
##AuthLDAPSubGroupClass groupOfNames
##AuthLDAPSubGroupAttribute uniqueMember
##AuthLDAPSubGroupAttribute member
require ldap-group cn=Developers_group,o=OurOrg,c=us
</Directory>
</VirtualHost>

```

LDAPTrustedMode None

mod_ibm_ldap directives migration

This article contains information to help with migration from existing directives that use the mod_ibm_ldap module to the use of the open source LDAP modules (mod_authnz_ldap and mod_ldap). Migration will ensure future support for your LDAP configuration.

Attention: Although many of the mod_ibm_ldap directives are located in the ldap.prop file, the open source LDAP directives are all located in the httpd.conf file.

The open source LDAP features are provided by two modules. The AuthLDAP directives are provided by the mod_authnz_ldap module and the LDAP directives are provided by the mod_ldap module. Both modules need to be loaded for the LDAP features to be available. Throughout the following section the generic name, mod_ldap, is used to reference the open source LDAP modules.

ldapCodePageDir:

The `mod_ldap` module does not provide a directive for specifying a codepages directory. The codepages directory is automatically installed in the correct directory, and the codepages directory cannot be moved from its installed location.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldapCodePageDir /location/of/codepages
```

LDAPConfigfile:

The `mod_ldap` module does not provide a directive for specifying an LDAP configuration file. Although there is no `mod_ldap` directive for specifying the LDAP configuration file, if you want to put your LDAP configuration in a separate file, you might use the Apache include directive.

Convert this:

```
ldapConfigFile ldap.prop
```

to this:

```
Include /location/of/ldap_conf/apache_ldap.conf
```

Another alternative for migrating the `mod_ibm_ldap LDAPConfigfile` directive is to use the `mod_authn_alias` module `AuthnProviderAlias` container to create one or more groupings of `ldap` directives, and then use them by referencing the alias labels where required

LdapRequire:

The `mod_ldap` module provides the `require` directive, with LDAP extensions, for LDAP authentication security.

If you used `require valid-user` previously for IBM HTTP Server, you may leave this `require` directive in place without modification. For the highest level of LDAP authentication security, you should migrate `require valid-user` to a more specific form. For additional information, see the Apache documentation for these `require` directives: `ldap-user`, `ldap-dn`, `ldap-attribute`, `ldap-group`, `ldap-filter`, and `valid-user`.

Convert this:

```
LdapRequire filter "&(objectclass=person)(cn=*)(ou=OurUnit)(o=OurOrg)"  
LdapRequire group MyDepartment
```

to this:

```
require ldap-filter &(objectclass=person)(cn=*)(ou=OurUnit)(o=OurOrg)  
require ldap-group cn=MyDepartment,o=OurOrg,c=US
```

ldap.application.authType:

The `mod_ldap` module does not provide a directive specifying an authentication type. If a value is specified for the `AuthLDAPBindDN` directive, then basic authentication is enabled. If a value is not specified for the `AuthLDAPBindDN` directive, then what was previously the `None` authentication type for the `mod_ibm_ldap` module, or `anonymous`, is enabled.

If a value is specified for the `LDAPTrustedClientCert` directive, then the certificate authentication type is used automatically.

```
ldap.application.authType=[None | Basic | Cert]
```

ldap.application.DN:

The `mod_ldap` module provides the `AuthLDAPBindDN` directive to determine the application authentication type.

If a value is specified for the AuthLDAPBindDN directive, then the value of the authType directive is Basic. If the AuthLDAPBindDN directive is not enabled, then the value for the authType directive is None. If a value is specified for the LDAPTrustedClientCert directive, then the value for the authType directive is Cert.

Important: AuthLDAPBindDN also takes the place of ldap.application.authType.

Convert this:

```
ldap.application.DN=cn=ldapadm,ou=OurDirectory,o=OurCompany,c=US
```

to this:

```
AuthLDAPBindDN "cn=ldapadm,ou=OurDirectory,o=OurCompany,c=US"
```

ldap.application.password:

The mod_ldap module provides the AuthLDAPBindPassword directive to specify a bind password. The value is stored in the configuration file in plain text. Therefore, you should restrict access to the configuration file

Convert this:

```
ldap.application.password=mypassword
```

to this:

```
AuthLDAPBindPassword mypassword
```

ldap.application.password.stashFile:

The mod_ldap module does not provide a directive for stashing the password. The directive AuthLDAPBindPassword is the only means to specify a password, and the value is stored in the configuration file in plain text. Therefore, you should restrict access to the configuration file.

This mod_ibm_ldap directive has no mod_ldap equivalent:

```
ldap.application.password.stashfile=/path/to/stashfile.sth
```

ldap.cache.timeout:

The mod_ldap module provides the LDAPCacheTTL directive to specify a timeout for the LDAP cache. The LDAPCacheTTL directive is globally scoped and must be located at the top level of the configuration file. This is different from the mod_ibm_ldap module, because the ldap.cache.timeout directive could be located anywhere in the configuration file.

Convert this:

```
ldap.cache.timeout=60
```

to this:

```
LDAPCacheTTL 60
```

The default value is 600 seconds.

ldap.group.dnattributes:

The mod_ldap module provides the AuthLDAPSubGroupClass directive to specify the object classes which identify groups. For the mod_ibm_ldap module all values were specified on a single directive line; but for the mod_ldap module, the values can either be specified all on one line or on multiple lines, with the directive and one value on each line.

Convert this:

```
ldap.group.dnattributes=groupOfNames GroupOfUniqueNames
```

to this:

```
AuthLDAPSubGroupClass groupOfNames  
AuthLDAPSubGroupClass groupOfUniqueNames
```

These are the default values.

ldap.group.memberattribute:

The `mod_ldap` module provides the `AuthLDAPSubGroupAttribute` directive to specify the labels which identify the subgroup members of the current group. For the `mod_ibm_ldap` module, you could only specify one label; but for the `mod_ldap` module, you can specify multiple labels either by listing all of the labels in one directive line or by providing multiple directive lines, with each label on a separate directive line.

Convert this:

```
ldap.group.memberattribute=member
```

to this:

```
AuthLDAPSubGroupAttribute member  
AuthLDAPSubGroupAttribute uniqueMember
```

ldap.group.memberattributes:

The `mod_ldap` module provides the `AuthLDAPGroupAttribute` directive to specify the labels which identify any member of the current group, such as a user or subgroup. For the `mod_ibm_ldap` module, you specified all labels on one directive line; but for the `mod_ldap` module, you may either specify them all on one directive line or specify each label on a separate directive line.

Convert this:

```
ldap.group.memberattributes=member uniqueMember
```

to this:

```
AuthLDAPGroupAttribute member  
AuthLDAPGroupAttribute uniqueMember
```

ldap.group.name.filter:

The `mod_ldap` module does not provide a directive to specify separate user and group filters. The `mod_ldap` module uses the filter that is provided at the end of the `AuthLDAPURL` directive. You can use the `AuthnProviderAlias` container directive, which is provided by the `mod_authn_alias` module, to create separate `my_ldap_user_alias` and `my_ldap_group_alias` aliases containing the required `ldap` directives. You can then use your group alias in locations where authorization is controlled by way of group membership.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.group.name.filter=(&(cn=%v1)(|(objectclass=groupofnames)(objectclass=groupofuniqueNames)))
```

ldap.group.search.depth:

The `mod_ldap` module provides the `AuthLDAPMaxSubGroupDepth` directive to limit the recursive depth pursued before stopping attempts to locate a user within nested groups.

Convert this:

```
ldap.group.search.depth=5
```

to this:

```
AuthLDAPMaxSubGroupDepth 5
```

The default value is 10.

ldap.group.URL:

The mod_ldap module does not provide a directive for specifying an LDAP server for authorizing a group membership that is different from the LDAP server that is used to authenticate users.

You must also specify the LDAP group server in the AuthLDAPURL directive for the container. Ensure that you specify the correct filter for each group.

```
ldap.group.URL=ldap://groups_ldap.server.org:389/o=OurOrg,c=US  
ldap.group.URL=ldaps://groups_ldap.server.org:636/o=OurOrg,c=US
```

ldap.idleConnection.timeout:

The mod_ldap module does not provide a directive for specifying when established connections to the LDAP server, that have gone idle, should timeout. The mod_ldap module automatically detects when the LDAP server expires connections, but does not cause connections to expire.

This mod_ibm_ldap directive has no mod_ldap equivalent:

```
ldap.idleConnection.timeout=60
```

ldap.key.file.password.stashfile:

If no password is specified in the LDAPTrustedGlobalCert directive, the mod_ldap module automatically uses a /path/to/keyfile.sth file (assuming that /path/to/keyfile.kdb is the keyfile that is specified in the LDAPTrustedGlobalCert directive).

For information about how to specify the keyfile password, see the Apache information for the LDAPTrustedGlobalCert directive. The value is stored in the configuration file in plain text. Therefore, you should restrict access to the configuration file.

This mod_ibm_ldap directive has no mod_ldap equivalent:

```
ldap.key.file.password.stashfile=/path/to/ldap.sth
```

ldap.key.fileName:

The mod_ldap module provides the LDAPTrustedGlobalCert directive to specify the keyfile to be used when loading certificates. The mod_ldap module also uses these directives to specify the password in plain text in the configuration file. Therefore, you should restrict access to the configuration file.

Convert this:

```
ldap.key.filename=/path/to/keyfile.kdb
```

to this:

```
LDAPTrustedGlobalCert CMS_KEYFILE /path/to/keyfile.kdb myKDBpassword  
LDAPTrustedGlobalCert SAF saf_keyring
```

ldap.key.label:

The `mod_ldap` module provides the `LDAPTrustedClientCert` directive to specify which certificate to use from the KDB keyfile. If the default certificate is used, then you do not need to specify a value for these directives.

Convert this:

```
ldap.key.label=certname_from_kdb
```

to this:

```
LDAPTrustedClientCert CMS_LABEL certname_from_kdb
```

ldap.ReferralHopLimit:

The `mod_ldap` module provides the `LDAPReferralHopLimit` directive to limit the number of referrals to chase before stopping attempts to locate a user in a distributed directory tree.

Convert this:

```
ldapReferralHopLimit 5
```

to this:

```
LDAPReferralHopLimit 5
```

The default value is 5.

ldapReferrals:

The `mod_ldap` module provides the `LDAPReferrals` directive to enable or disable referral chasing when locating users in a distributed directory tree.

Convert this:

```
ldapReferrals On
```

to this:

```
LDAPReferrals On
```

The default value is On.

ldap.realm:

The `mod_ldap` module provides the `AuthName` directive to specify the authorization realm.

Convert this:

```
ldap.realm=Some identifying text
```

to this:

```
AuthName "Some identifying text"
```

ldap.search.timeout:

The `mod_ldap` module provides the `LDAPSearchTimeout` directive to specify when a search request should be abandoned.

Convert this:

```
ldap.search.timeout=10
```

to

LDAPSearchTimeout 10

The default value is 10 seconds.

ldap.transport:

The mod_ldap module provides the LDAPTrustedMode directive to specify the type of network transport to use when communicating with the LDAP server.

If no port is specified on the AuthLDAPURL directive, then the mod_ldap module ignores the LDAPTrustedMode directive, and specifies a network transport value of SSL. For more information, see the Apache documentation for the LDAPTrustedMode and AuthLDAPURL directives.

You can specify a value for the following network transport types.

- None or TCP, which indicates no encryption. If no port is specified on the AuthLDAPURL directive, then port 389 is used.
- SSL. If a value of None is specified, then port 636 is used.
- TLS or STARTTLS. These open source types are not supported by IBM HTTP Server.

Convert this:

```
ldap.transport=TCP (or SSL)
```

to this:

```
LDAPTrustedMode NONE (or SSL)
```

If an ldaps://URL is specified, the mode becomes SSL and the setting of LDAPTrustedMode is ignored.

ldap.URL:

The mod_ldap module provides the AuthLDAPURL directive for specifying the LDAP server hostname and port as well as the base DN to use when connecting to the server. The mod_ldap module also provides a means for specifying the user attribute, scope, user filter, and transport mode. For more information, see the Apache documentation for the AuthLDAPURL directives.

Convert this:

```
ldap.URL=ldap://our_ldap.server.org:389/o=OurOrg,c=US  
ldap.URL=ldaps://our_ldap.server.org:636/o=OurOrg,c=US
```

to this:

```
AuthLDAPURL ldap://our_ldap.server.org:389/o=OurOrg,c=US?cn?sub?(objectclass=person)  
AuthLDAPURL ldaps://our_ldap.server.org:636/o=OurOrg,c=US?cn?sub?(objectclass=person)
```

ldap.user.authType:

The mod_ldap module does not provide a directive for specifying a user authentication type. The mod_ldap module authenticates users based on the user ID and password credentials provided.

This mod_ibm_ldap directive has no mod_ldap equivalent:

```
ldap.user.authType=Basic [Basic | Cert | BasicIfNoCert]
```

ldap.user.cert.filter:

The mod_ldap module does not provide a directive for filtering client certificates. The mod_ldap module does not work directly with client certificates.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.user.cert.filter=(&(objectclass=person)(cn=%v1)(ou=%v2)(o=%v3)(c=%v4))
```

ldap.user.name.fieldSep:

The `mod_ldap` module does not provide a directive for parsing provided credentials into subcomponents. The `mod_ibm_ldap` module uses the `ldap.user.name.fieldSep` directive to specify the separator characters used to parse the credentials into the `%v1`, `%v2`, ...`%vN` tokens.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.user.name.fieldSep=/ ,
```

ldap.user.name.filter:

The `mod_ldap` module does not provide a directive for specifying the user name filter. The `mod_ldap` module specifies the user name filter as part of the `AuthLDAPURL` directive.

The `AuthLDAPURL` directive combines the user attribute specified in the directive with the provided filter to create the search filter. The provided filter follows the standard search filter specification. The `mod_ldap` module also does not provide the `%vx` token parsing function available for the `mod_ibm_ldap` module.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.user.name.filter=(&(objectclass=person)(cn=%v1 %v2))
```

ldap.version:

The `mod_ldap` module does not provide a directive for specifying the LDAP version. The `mod_ldap` module uses only LDAP version 3.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.version=2 (or 3)
```

ldap.waitToRetryConnection.interval:

The `mod_ldap` module does not provide a directive for specifying an amount of time before retrying a failed connection attempt. The `mod_ldap` module does not have a timed delay between connection retries when a connection attempt fails. The connection attempt is automatically retried for a maximum of 10 times before a request fails.

When a new request needs to access the same LDAP server, the connection is retried for a maximum of 10 times again. The retry throttle is based on the volume of new requests sent to the LDAP server.

This `mod_ibm_ldap` directive has no `mod_ldap` equivalent:

```
ldap.waitToRetryConnection.interval=300
```

Authenticating with LDAP on IBM HTTP Server using `mod_ldap`

You can configure Lightweight Directory Access Protocol (LDAP) to authenticate and protect files on IBM HTTP Server.

Before you begin

Best Practice: Distributed operating systems If you are using the `mod_ibm_ldap` module for your LDAP configuration, consider migrating your `mod_ibm_ldap` directives to use the `mod_ldap` module. The `mod_ibm_ldap` module is provided with this release of IBM HTTP Server for

compatibility with previous releases, however, you must migrate existing configurations to use the `mod_authnz_ldap` and `mod_ldap` modules to ensure future support for your LDAP configuration.

The `LoadModule` directive for LDAP does not load into IBM HTTP Server by default. Without the `LoadModule` directive, the LDAP features are not available for use.

In order to enable the LDAP function, add a `LoadModule` directive to the IBM HTTP Server `httpd.conf` file as follows:

```
LoadModule ldap_module modules/mod_ldap.so
LoadModule authnz_ldap_module modules/mod_authnz_ldap.so
```

About this task

LDAP authentication is provided by the `mod_ldap` and `mod_authnz_ldap` Apache modules.

- The `mod_ldap` module provides LDAP connection pooling and caching.
- The `mod_authnz_ldap` makes use of the LDAP connection pooling and caching services to provide Web client authentication.

See the following Web sites to obtain detailed descriptions of the LDAP (`ldap_module` and `authnz_ldap_module`) directives:

- http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_ldap.html
- http://publib.boulder.ibm.com/htpserv/manual70/mod/mod_authnz_ldap.html

Procedure

1. Edit the `httpd.conf` IBM HTTP Server configuration file.
2. Determine the resource for which you want to limit access. For example: `<Directory "/secure_info">`
3. Add the `LDAPTrustedGlobalCert` directive to `httpd.conf` if the IBM HTTP Server connection to the LDAP server is an SSL connection.

The `LDAPTrustedGlobalCert` directive specifies the directory path and file name of the trusted certificate authority (CA) that `mod_ldap` uses when establishing an SSL connection to an LDAP server.

Certificates can be stored in a `.kdb` file or a SAF key ring. If a `.kdb` file is used, a `.sth` file must be located in the same directory path and have the same filename, but the extension must be `.sth` instead of `.kdb`.

The `LDAPTrustedGlobalCert` directive must be a `CMS_KEYFILE` value type. Use this value if the certificates indicated by the `LDAPTrustedGlobalCert` directive are stored in a `.kdb` file.

The `LDAPTrustedGlobalCert` directive must be a `SAF_KEYRING` value type. Use this value if the certificates indicated by the `LDAPTrustedGlobalCert` directive are stored in a SAF key ring. Example

when the certificate is stored in a `.kdb` file: **Distributed operating systems**

```
LDAPTrustedGlobalCert CMS_KEYFILE /path/to/keyfile.kdb myKDBpassword
```

Example when the certificate is stored in a SAF key ring: **z/OS**

```
LDAPTrustedGlobalCert SAF saf_keyring
```

Important: The user ID that you use to start IBM HTTP Server must have access to the SAF key ring that you name in this directive. If the user ID does not have access to the SAF key ring, SSL initialization fails.

See “Performing required z/OS system configurations” on page 1 for information on accessing SAF key rings defined in RACF.

4. Add the `AuthLDAPUrl` directive, which specifies the LDAP search parameters to use.

The syntax of the URL is:

```
ldap://host:port/basedn?attribute?scope?filter
```

5. Add directives in `httpd.conf` to the directory location (container) to be protected with values specific to your environment, such as:

- `Order deny,allow`
- `Allow from all`
- `AuthName "Title of your protected Realm"`
- `AuthType Basic`
- `AuthBasicProvider ldap`
- `AuthLDAPURL your_ldap_url`
- `Require valid-user`
- `AuthLDAPBindDN "cn=Directory Manager"`
- `AuthLDAPBindPassword auth_password`

For each combination of LDAP server, protection setup, and `protect` directive, code a Location container similar to the following example:

```
<Location /ldapdir>
  AuthName "whatever_LDAP"
  AuthType Basic
  AuthBasicProvider ldap
  AuthLDAPURL ldap://9.27.163.182:389/o=abc.xyz.com?cn?sub?
  Require valid-user
  AuthLDAPBindDN "cn=Directory Manager"
  AuthLDAPBindPassword d44radar
</Location>
```

http://publib.boulder.ibm.com/httperv/manual70/mod/mod_authnz_ldap.html

Authenticating with SAF on IBM HTTP Server (z/OS systems)

You can authenticate to the IBM HTTP Server on z/OS using HTTP basic authentication or client certificates with the System Authorization Facility (SAF) security product. Use SAF authentication for verification of user IDs and passwords or certificates.

Before you begin

The `mod_authz_default` and `mod_auth_basic` directives provide basic authentication and authorization support which is needed in `mod_authnz_saf` configurations. In addition, the `mod_ibm_ssl` directive provides support for SSL client certificates. If you use SAF authentication, ensure that the first three `LoadModule` directives from the following example are activated. If you use SSL client certificates, ensure that the `mod_ibm_ssl.so` `LoadModule` directive is activated as well.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_default_module modules/mod_authz_default.so
# Uncomment mod_ibm_ssl if any type of SSL support is required,
# such as client certificate authentication
#LoadModule ibm_ssl_module modules/mod_ibm_ssl.so
```

If the `mod_authz_default` module is not loaded by your Web server, the server returns a response code 500 instead of 401 if the user is not authorized.

About this task

SAF authentication is provided by the `mod_authnz_saf` module. The `mod_authnz_saf` module allows the use of HTTP basic authentication or client certificates to restrict access by looking up users, groups, and SSL client certificates in SAF. This module also allows you to switch the thread from the server ID to another ID prior to responding to the request by using the `SAFRUNAS` directive. For additional information, see the information center topic about SAF directives. Also, see the topic about migrating and installing

IBM HTTP Server on z/OS systems for information about migrating your SAF directives.

Procedure

1. If you are using SAFRunAs, permit the IBM HTTP Server userid to the BPX.SERVER FACILITY class profile in RACF, and provide the target userid with an OMVS segment.
2. Determine the directory location you want to limit access to. For example: <Location "/admin-bin">.
3. Add directives in the httpd.conf file to the directory or location to be protected with values specific to your environment. If you want to restrict access to files under the /secure directory to only users who provide a valid SAF user ID and password, consider this example.

```
<Directory /secure>
  AuthName protectedrealm_title
  AuthType Basic
  AuthBasicProvider saf
  Require valid-user
</Directory>
```

You can also restrict access based on user ID or SAF group membership by replacing the Require directive in the previous example, as follows:

```
require saf-user USERID
require saf-group GROUPNAME
```

4. Optional: Specify Require saf-user or Require saf-group to restrict access to a specific SAF user or group.

SAF directives

These configuration parameters control the System Authorization Facility (SAF) feature for IBM HTTP Server. Use the SAF directives to provide IBM HTTP Server with user authentication.

- “AuthSAFAuthoritative directive”
- “AuthSAFExpiration directive” on page 189
- “AuthSAFReEnter directive” on page 189
- “SAFRunAs directive” on page 190

AuthSAFAuthoritative directive

The AuthSAFAuthoritative directive sets whether authorization is passed to lower level modules.

Syntax	AuthSAFAuthoritative on off
Default	on
Context	directory, .htaccess
Module	mod_authnz_saf
Values	on or off

Setting the AuthSAFAuthoritative directive off allows for authorization to be passed to lower level modules (as defined in the modules.c files), if there is no user ID or rule matching the supplied user ID. If there is a user ID or rule specified, then the usual password and access checks will be applied and a failure will result in an Authentication Required reply.

If a user ID appears in the database of more than one module, or if a valid Require directive applies to more than one module, then the first module will verify the credentials, and no access is passed on, regardless of the AuthSAFAuthoritative setting.

By default, control is not passed on and an unknown user ID or rule will result in an Authentication Required reply. Not setting it thus keeps the system secure and forces an NCSA compliant behavior.

AuthSAFExpiration directive

The AuthSAFExpiration directive sets the value displayed in the browser prompt. The server sends the value specified for the AuthName directive and this short phrase in an HTTP response header, and then the browser displays them to the user in a password prompt window. The short phrase is subject to the same character limitations as the specified value for the AuthName directive. Therefore, to display a special character in the password prompt window, the server must translate the special character from the EBCDIC CharsetSourceEnc codepage to the ASCII CharsetDefault codepage. For example, if you want to display a lowercase 'a' with umlaut, and the httpd.conf file contains the German language EBCDIC codepage "CharsetSourceEnc IBM-1141" and the ASCII codepage "CharsetDefault ISO08859-1", then you must code the phrase using the hex value '43', which translates to the correct ASCII character.

Syntax	AuthSAFExpiration <i>short_phrase</i>
Default	off
Context	directory, .htaccess
Module	mod_authnz_saf
Values	off or <i>short_phrase</i>

Setting the AuthSAFExpiration directive to a phrase allows IBM HTTP Server to prompt the user to update his SAF password if it expires. When the user enters a valid ID and SAF password but the password has expired, the server will return an Authentication Required reply with a special prompt to allow the user to update the expired password. The prompt consists of the realm (the value from the AuthName directive) followed by the *short_phrase* value from the AuthSAFExpiration directive.

For example, consider the following configuration:

```
<Location /js>
AuthType basic
AuthName "zwasa051_SAF"
AuthBasicProvider saf
Require valid-user
Require saf-group SYS1 WASUSER
AuthSAFExpiration "EXPIRED! oldpw/newpw/newpw"
</Location>
```

If the user attempts to access a file whose URL starts with /js, then the server prompts for a SAF ID and password. The browser will display a prompt containing the realm. The realm is the value from the AuthName directive, which is zwasa051_SAF in this example.

When the user supplies a valid ID and password, if the password has expired, the server will repeat the prompt, but this time with the value zwasa051_SAF EXPIRED! oldpw/newpw/newpw. Whatever the prompt, the user must then re-enter the expired password, followed by a slash, the new password, another slash, and the new password again.

If the password update is successful, the server will send another Authentication Required reply with a distinct special prompt. This last interaction is necessary in order to force the browser to understand which password it should cache. The prompt this time will consist of the realm followed by the prompt Re-enter new password. In this example, it would be zwasa051_SAF Re-enter new password.

AuthSAFReEnter directive

The AuthSAFReEnter directive sets the value appended to realm after a successful password change. For information about coding special characters, see the BAuthSAFExpiration directive.

Syntax	AuthSAFReEnter <i>short_phrase</i>
Default	Re-enter new password
Context	directory, .htaccess

Module	mod_authnz_saf
Values	off or <i>short_phrase</i>

Setting the AuthSAFReEnter directive explicitly to a phrase other than "Re-enter new password" allows the administrator to display an alternative message after an expired password has been updated successfully. If AuthSAFExpiration has been set to off, this directive has no effect.

For example, consider the following configuration:

```
<Location /js>
AuthType basic
AuthName "zwasa051_SAF"
AuthBasicProvider saf
Require saf-user SYSADM USER152 BABAR
AuthSAFExpiration "EXPIRED! oldpw/newpw/newpw"
AuthSAFReEnter "Enter new password one more time"
</Location>
```

In this example, after the expired password is updated successfully, the server will send another Authentication Required reply with the value from the AuthSAFReEnter directive. This last interaction is necessary in order to force the browser to understand which password it should cache. The prompt this time will consist of the realm followed by a special phrase. In this example, it would be zwasa051_SAF Enter new password one more time.

SAFRunAs directive

The SAFRunAs directive sets the SAF user ID under which a request will be served.

Syntax	SAFRunAs <i>value</i>
Default	off
Context	directory, .htaccess
Module	mod_authnz_saf

Values

off | %%CLIENT%% | %%CERTIF%% | %%CERTIF_REQ%% |
<surrogate ID>

Off: The server will run the request under the Web server user ID.

%%CLIENT%%: The server will run the request under the ID supplied in the Authorization request header. Generally, the user supplies the ID and password in a pop-up window on the browser, and the browser creates the header. Requires that SAF is configured to authenticate the URL.

%%CERTIF%%: The server will run the request under the ID associated with the SSL client certificate in SAF. If there is no SSL certificate or if the SSL certificate has not been associated with an ID in SAF, the processing will continue as if %%CLIENT%% had been coded. Does not require SAF authn or authz to be configured.

%%CERTIF_REQ%%: The server will run the request under the ID associated with the SSL client certificate in SAF. If there is no SSL certificate, or if the SSL certificate has not been associated with an ID in SAF, the server will not allow access. Does not require SAF authn or authz to be configured.

<surrogate ID>: The server will run the request under the ID associated with the SAF surrogate ID specified.

IBM HTTP Server can communicate with FastCGI applications using either TCP sockets or UNIX sockets. However, when using SAFRunAs for FastCGI requests, you must use TCP sockets for communication with the application. UNIX sockets that are created for FastCGI applications are accessible by the Web server user ID only. The alternate user ID controlled with the SAFRunAs directive does not have permission to access the UNIX sockets, so requests will fail.

To configure FastCGI to use TCP sockets, define the FastCGI application to the mod_fastcgi module using the FastCGIServer directive with the -port option or using the FastCGIExternalServer directive. Dynamic FastCGI servers that you do not configure with the FastCGIServer or FastCGIExternalServer are not usable with SAFRunAS.

If you do not enable SAFRunAs for FastCGI requests, TCP sockets are not required.

If you want to use SAF for authentication and authorization, consider the following example. This is the most common scenario for SAF users and groups and meets the requirements for web access.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_default_module modules/mod_authz_default.so
...
<Location /saf_protected>
AuthType basic
AuthName x1
AuthBasicProvider saf
# Code "Require valid-user" if you want any valid
# SAF user to be able to access the resource.
Require valid-user
#
# Alternately, you can provide a list of specific SAF users
# who may access the resource.
# Require saf-user USER84 USER85
```

```
#
# Alternatively, you can provide a list of specific SAF groups
# whose members may access the resource.
# Require saf-group WASGRP1 WASGRP2
</Location>
```

If you want to use a SAF file for authentication, but use a non-SAF group file for authorization, consider the following example. In this example, users are authenticated using SAF, but authorized using a different mechanism.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_groupfile_module modules/mod_authz_groupfile.so
LoadModule authz_default_module modules/mod_authz_default.so
...
<Location /saf_password>
AuthType basic
AuthName "SAF auth with hfs groupfile"
AuthBasicProvider saf
AuthGroupFile /www/config/foo.grp
# Code "Require file-group" and a list of groups if you want
# a user in any of the groups in the specified group file to be able
# to access the resource.
# Note: Any authorization module, with its standard configuration, can be used here.
Require group admin1 admin2
</Location>
```

If you want to allow access to a user if the user is authorized by SAF or by a group file, consider the following example.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_groupfile_module modules/mod_authz_groupfile.so
LoadModule authz_default_module modules/mod_authz_default.so
...
<Location /either_group>
AuthType basic
AuthName "SAF auth with SAF groups and hfs groupfile"
AuthBasicProvider saf
AuthGroupFile /www/groupfiles/foo.grp
Require saf-group WASGRP
Require saf-group ADMINS
AuthzGroupFileAuthoritative Off
AuthSAFAuthoritative Off
</Location>
```

If you want to require a request to run using the SAF privileges associated with the authenticated username, consider the following example.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_default_module modules/mod_authz_default.so
...
<Location /runas_admin_bin>
AuthName "SAF RunAs client"
AuthType basic
Require valid-user
AuthBasicProvider saf
SAFRunAs %%CLIENT%%
</Location>
```

If you want to support the changing of expired SAF passwords, consider the following example.

```
LoadModule auth_basic_module modules/mod_auth_basic.so
LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule authz_default_module modules/mod_authz_default.so
...
```

```

<Location /custom_password_change>
AuthType basic
AuthName "Support expired PW"
Require valid-user
AuthBasicProvider saf
AuthSAFExpiration "EXPIRED PW: oldpw/newpw/newpw"
AuthSAFReEnter "New PW again:"
</Location>

```

If you want to require a client certificate before a user can access a resource, use the `mod_ibm_ssl` directive. The `mod_authnz_saf` directive is not needed for this configuration. For additional information, see the documentation for the `SSLClientAuth` and `SSLClientAuthRequire` directives.

If you want to use a client certificate to determine the user for whom request processing is performed, consider the following example. If the user does not have a valid certificate, access is denied.

```

LoadModule authnz_saf_module modules/mod_authnz_saf.so
LoadModule ibm_ssl_module modules/mod_ibm_ssl.so
...
<Location /certificate_required>
SAFRunAs %%CERTIF_REQ%%
</Location>

```

If you want to require a request to run using the SAF privileges associated with a client certificate, but require username and password authentication if the client certificate is not mapped to a SAF user, consider the following example. If the user provides a certificate that SAF can map to a user ID, then the user ID must also pass any `Require` directives.

```

<Location /certificate_or_basic>
AuthName "SAF RunAs certifi"
AuthType basic
Require saf-user USER84 USER103
AuthBasicProvider saf
SAFRunAs %%CERTIF%%
</Location>

```

If you want to require a request to run using the SAF privileges associated with a surrogate ID, consider the following example.

```

<Location /runas_public>
SAFRunAs PUBLIC
# This can be combined with SAF or non-SAF authentication/authorization
</Location>

```

Chapter 6. Troubleshooting and support: IBM HTTP Server

This section provides information about how to troubleshoot a problem with IBM HTTP Server.

Troubleshoot problems with IBM HTTP Server, using the problem determination tools provided with the product. For example, you can perform problem determination with IBM HTTP Server, including platform-specific problems and error messages.

Troubleshooting IBM HTTP Server

This section describes how to start troubleshooting IBM HTTP Server.

Procedure

1. Check the error log to help you determine the type of problem. You can find the error logs in the directory specified by the ErrorLog directive in the configuration file. Depending on the operating system, the default directories are:
 - **AIX** /usr/IBM/HTTPServer/logs/error_log
 - **HP-UX** **Linux** **Solaris** /opt/IBM/HTTPServer/logs/error_log
 - **Windows** <server_root>/logs/error.log
 - **z/OS** <server_root>/logs/error_log
2. Check the IBM HTTP Server Diagnostic Tools and Information package at <http://www.ibm.com/support/docview.wss?uid=swg24008409> for additional diagnostic information, as well as MustGather steps for some problems.
3. Check the IBM HTTP Server support page at <http://www.ibm.com/software/webservers/htpservers/support/> for technotes on a variety of topics.
4. Ensure that you are running with the current level of fixes for your release of IBM HTTP Server. The problem may already be resolved. Find the IBM HTTP Server recommended updates page is at <http://www.ibm.com/support/docview.wss?rs=177&context=SSEQTJ&uid=swg27005198>.

Known problems on Windows platforms

This topic contains troubleshooting information about known problems on Windows platforms.

Problems when the IBM HTTP Server runs on the same system as a Virtual Private Networking Client

A problem occurs when the IBM HTTP Server runs on a system, along with a Virtual Private Networking client, for example, Aventail Connect. You can experience the following problem, or see the following error message:

- The IBM HTTP Server does not start - see Apache HTTP Server - Frequently asked questions.
- The IBM HTTP Server does not start. The error log contains the following message:

```
"[crit] (10045) The attempted operation is not supported for the type of object referenced: Parent: WSADuplicateSocket failed for socket ###"
```

Aventail Connect is a Layered Service Provider (LSP) that inserts itself, as a shim, between the Winsock 2 API and the Windows native Winsock 2 implementation. The Aventail Connect shim does not implement WSADuplicateSocket, the cause of the failure. The shim is not unloaded when Aventail Connect is shut down.

Fix the problem by doing one of the following:

- Explicitly unloading the shim
- Rebooting the machine
- Temporarily removing the Aventail Connect V3.x shim

An application start-up error occurs for some IBM HTTP Server and web server plug-in components on Windows operating systems

A message occurs for some IBM HTTP Server and web server plug-in components on Windows operating systems. The message indicates that an application has failed to start because its side-by-side configuration is incorrect. When Secure Sockets Layer (SSL) is configured in either IBM HTTP Server or the web server plug-in, the web server fails to load.

Additionally, the `<ihsinst>\bin\gskver` and `<ihsinst>\bin\gskcapicmd` programs fail with the same error. These two programs are part of the Global Security Kit (GSKit) certificate management tools.

These programs fail with the following error message: The application has failed to start because its side-by-side configuration is incorrect. Please see the application event log for more detail.

In the application event log, the following event is logged:

```
Activation context generation failed for "_IHS_install_path_gsk8\bin\gsk8ver.exe". Dependent Assembly Microsoft.VC90.CRT,processorArchitecture="x86",publicKeyToken="1fc8b3b9a1e18e3b",type="win32", version="9.0.21022.8" could not be found. Please use sxstrace.exe for detailed diagnosis.
```

Fix the problem by installing the Microsoft Visual C++ 2008 Redistributable Package (x86), available from <http://www.microsoft.com/downloads/details.aspx?familyid=9b2da534-3e03-4391-8a4d-074b9f2bc1bf>. You can also search for the `vcredist_x86.exe` file on the Microsoft website. If you are using a 64-bit web server plug-in, also install the 64-bit redistributable package.

Note: The installation process checks to see if the Microsoft package is installed. If it is not, you receive the following message. The installation package IBM HTTP Server v8.0 requires components supplied by other packages. To fix the issue, either install the required components or deselect the installation package. The required components may be supplied by the following installation packages: Package: Microsoft Visual C++ 2008 Redistributable Package.

Known problems on z/OS platforms

This topic contains troubleshooting information for known problems on z/OS platforms.

MEMLIMIT parameter must be set for the IBM HTTP Server address spaces: The MEMLIMIT parameter can be set on a system-wide basis (in the SMFPRMxx parmlib member) or in the OMVS segment of the server ID for each IBM HTTP Server instance. See the z/OS V1R8.0 MVS Extended Addressability Guide for more information. For recommended MEMLIMIT values, see "Performing required z/OS system configurations" on page 1.

If you do not set the MEMLIMIT parameter, the Web server will not start, and one of the

following console messages might result:

- ABEND=S000 U4093
REASON=00000224
- no output from
bin/apachectl -v
- bin/ab returns
"Killed"

To determine if any 64-bit programs run on this system, run the following command from a shell prompt:
/bin/localedef64.

Expected output:

```
# /bin/localedef64  
EDC4175 40 Missing output locale name.
```

Example of a failure:

```
# /bin/localedef64  
Killed
```

To resolve this problem for IBM HTTP Server, which is an AMODE64 application, the MEMLIMIT must be changed from the system default of 0.

Integrated Cryptographic Services Facility (ICSF) is not enabled for AMODE64: z/OS V1R6 might need ICSF 64-bit Virtual Support to use ICSF cryptographic hardware. To issue messages on ICSF status,
GSK_SSL_HW_DETECT_MESSAGE=1
is set in bin/envvars.

If ICSF is not enabled for AMODE64, the GSK_SSL_HW_DETECT_MESSAGE will result in the following message logged to the error log at startup:
System SSL: ICSF services are not a

Known problems with hardware cryptographic support on AIX

This topic contains troubleshooting information for known problems with the cryptographic hardware on AIX.

You must install the `bos.pkcs11` package to get the PKCS11 module, and to initialize the device on AIX.

An added update to the `bos.pkcs11` package fixed a forking problem. Obtain the most recent copy of the `bos.pkcs11` package from the IBM PSeries Support Site, to ensure you have this fix.

If you are having problems using the IBM eBusiness Cryptographic Accelerator Device with IBM HTTP Server, do the following:

1. Reboot the machine.
2. Kill `pkcsslotd` and the shared memory it created. To determine what shared memory was created, type `ipcs -a` and for a size of 270760. This was the memory created by `pkcsslotd`.
3. Export `EXPSHM=ON`.
4. Start the `pkcs11` process: `/etc/rc.pkcs11`
5. Restart IBM HTTP Server: `./apachectl start`

Symptoms of poor server response time

If you notice that server CPU utilization appears low, but client requests for static pages take a long time to service, your server may be running out of server threads to handle requests.

This situation results when you have more inbound requests than you have Apache threads to handle those requests. New connections queue in the TCP/IP stack listen queue and wait for acceptance from an available thread. As a thread becomes available, it accepts and handles a connection off of the listen queue. Connections can take a long time to reach the top of the listen queue. When this condition occurs, the following error message will appear in the error log:

- **AIX** **HP-UX** **Linux** **Solaris** **z/OS** "Server reached MaxClients setting, consider raising the MaxClients setting"
- **Windows** "Server ran out of threads to serve requests. Consider raising the ThreadsPerChild setting"

Hints and tips for managing IBM HTTP Server using the administrative console

This topic contains helpful tips on using the WebSphere Application Server administrative console for managing the following operations for IBM HTTP Server: Starting, stopping, viewing log files, editing configuration files, and propagating the plug-in configuration file.

Administering IBM HTTP Server with the administrative console using the node agent and deployment manager:

- The following list describes hints and tips on starting, stopping, and obtaining status for IBM HTTP Server using the administrative console.
 - **Windows** The IBM HTTP Server you are managing must be installed as a service. You must install IBM HTTP Server with log on as system rights.
 - **Windows** When defining a Web server using the administrative console, use the actual service name, instead of the display name. The actual service name will not contain spaces. If you do not do this, you will have problems starting and stopping the service.
 - Status is obtained using the Web server host name and port that you have defined. You do not use the remote administration port.
 - If you have problems starting and stopping IBM HTTP Server, check the WebSphere console logs (trace).

- If you have problems starting and stopping IBM HTTP Server using nodeagent, you can try to start and stop the server by setting up the managed profile and issuing the `startserver <IBM HTTP Server> -nowait -trace` command and check the `startServer.log` file for the IBM HTTP Server specified.
- If communication between the administrative console and the Web server is through a firewall, then you must define the Web Server port to the firewall program.
- The following list describes hints and tips for viewing log files, editing configuration files and propagating the plug-in configuration file:
 - Access to files is controlled by `AdminAllowDirective` in the `admin.conf` file. Access is granted to the `conf` and `logs` directory from the IBM HTTP Server installation directory. If you are reading or writing plug-in configuration or trace files, you must add an entry to the `admin.conf` file to allow access there.
 - Always back up the configuration file. It is possible on the upload of the configuration file, information will be lost.

Distributed operating systems Administering IBM HTTP Server with the administrative console using the IBM HTTP Server administration server:

- The following list describes hints and tips on starting, stopping, and obtaining status for IBM HTTP Server using the administrative console.
 - **Windows** The IBM HTTP Server you are managing must be installed as a service.
 - **Windows** When defining a Web server using the administrative console, use the actual service name, instead of the display name. The actual service name will not contain spaces. If you do not do this, you will have problems starting and stopping the service on the Windows 2003 operating system.
 - Status is obtained using the Web server host name and port that you have defined. You do not use the remote administration port.
 - If you have problems starting and stopping IBM HTTP Server, check the WebSphere console logs (trace) and check the `admin_error.log` file.
 - The administration server should be started as root.
 - If communication between the administrative console and the administration server is through a firewall, you must enable the administration server port (default 8008).
 - If communication between the administrative console and the Web server is through a firewall, then you must define the Web Server port to the firewall program.
- The following list describes hints and tips for viewing log files, editing configuration files and propagating the plug-in configuration file:
 - **AIX** **HP-UX** **Linux** **Solaris** File permissions must be correct in order to transfer a file. The **setupadm** script is provided to set appropriate file permissions.
The `setupadm` script should be run prior to starting the administration server. This script will setup file permission and update the User ID and Group ID directives in the `admin.conf` file. The User ID and Group ID created through the `setupadm` script are UNIX IDs that must correspond to the `admin.conf` directives: User and Group.
 - Access to files is controlled by `AdminAllowDirective` in the `admin.conf` file. Access is granted to the `conf` and `logs` directory from the IBM HTTP Server installation directory. If you are reading or writing plug-in configuration or trace files, you must add an entry to the `admin.conf` file to allow access there.
 - Always back up the configuration file. It is possible on the upload of the configuration file, information will be lost.

Could not connect to IBM HTTP Server administration server error

This topic contains troubleshooting information if you receive an error when attempting to connect to the administration server.

If you get the following error:

"Could not connect to IHS Administration server error"

when you are managing an IBM HTTP Server using the WebSphere administrative console, try one of the following:

- Verify that the IBM HTTP Server administration server is running.
- Verify that the Web server hostname and port that is defined in the WebSphere administrative console matches the IBM HTTP Server administration host name and port.
- Verify that the firewall is not preventing you from accessing the IBM HTTP Server administration server from the WebSphere administrative console.
- Verify that the user ID and password that is specified in the WebSphere administrative console, under remote managed, is created in the `admin.passwd` file, using the **htpasswd** command.
- If trying to connect securely, verify that you export the IBM HTTP Server administration server keydb personal certificate into the WebSphere key database as a signer certificate. This key database will be specified by the `com.ibm.ssl.trustStore` in the `sas.client.props` file in the profile your console is running in. This is mainly for self-signed certificates.
- If you still have problems, check the IBM HTTP Server `admin_error.log` file and the WebSphere Application Server logs (`trace.log`) to see if problem can be determined.

Experiencing an IBM HTTP Server Service logon failure on Windows operating systems

When installing the IBM HTTP Server, prompts appear for a login ID and password. The ID you select must have the capability to log on as a service.

About this task

If you get an error when you try to start the IBM HTTP Server Service, indicating a failure to start as a service, try one of the following:

Procedure

1. Click **Start > Programs > Administrative Tools > User Manager**.
2. Select the user from the User Manager list.
3. Click **Policies > User Rights**.
4. Select the **Show Advanced User Rights** check box.
5. Click **Log on as a Service**, from the right drop-down menu.
or
 - a. Click **Start > Settings > Control Panel**.
 - b. Open Administrative Tools.
 - c. Open Services. The local user you select is created in Local Users and Groups, under Computer Management.
 - d. Click **Service > Actions > Properties**.
 - e. Choose the Log on tab.
 - f. Select this account option and click **Browse**, to select the user to associate with the service.

What to do next

If you get the following error when you try to start the IBM HTTP Server Service:

Windows could not start the IBM HTTP Server on Local Computer. For more information, review the Event Log. If this is a non-Microsoft service, contact the service vendor, and refer to service-specific error code 1.

complete the following steps:

1. Check the IBM HTTP Server `<install_root>/logs/error.log` file for a specific error.
2. If there is no error in the `<install_root>/logs/error.log` file, try starting IBM HTTP Server from a command prompt by running the `<install_root>/bin/httpd.exe` command.
3. If the `<install_root>/logs/error.log` file indicates that there was a problem loading the WebSphere Application Server plugin module, check the `http_plugin.log` file for the error.

Viewing error messages for a target server that fails to start

If you encounter an error starting a target server, you can view the error message in the server logs.

About this task

If the target Web server fails to start, a message might appear on the WebSphere Application Server administrative console that indicates that the Web server cannot be started and to view the error messages in the server logs for further details. The types of errors that can result are:

- errors due to caching problems
- errors due to configuration problems
- errors due to SSL handshake failures
- errors due to SSL initialization problems
- errors due to I/O failures
- errors due to Secure Sockets Layer (SSL) stash utility problems

Cache messages

This topic contains error messages that might result due to caching problems and provides a solution to help you troubleshoot the problem.

The following messages are displayed due to caching problems:

- Message: **SSL0600E: Unable to connect to session ID cache**
 - Reason: The server cannot connect to the Session ID caching daemon.
 - Solution: Verify that the daemon successfully started.
- Message: **SSL0601E: Session ID cache daemon process <process-id> exited with exit code <exit-code>; restarting**
 - Reason: If the value of `<exit-code>` is 0, the session ID cache daemon (sidd) received the SIGTERM signal. Other exit codes are not expected. Sidd automatically restarted.
 - Solution: If the value of `<exit-code>` is 0 and IBM HTTP Server did not stop or restart, verify that locally installed CGI scripts, scheduled operating system tasks, or other monitoring software cannot send SIGTERM to sidd.
- Message: **SSL0602E: Session ID cache daemon process <process-id> exited with terminating signal <signal-number>; restarting**
 - Reason: The session ID cache daemon (sidd) received a signal other than SIGTERM was received by the session ID cache daemon (sidd), which caused it to exit. Sidd automatically restarted.
 - Solution: If the value of `<exit-code>` is 0 and IBM HTTP Server did not stop or restart, verify that locally installed CGI scripts, scheduled operating system tasks, or other monitoring software cannot send the signal to sidd.
- Message: **SSL0603E: Session ID cache daemon process <process-id> exited with exit code<exit-code>; not restarting; check sidd configuration or enable sidd error log with SSLCacheErrorLog**
 - Reason: The session ID cache daemon (sidd) did not initialize. The following possible exit code values might be displayed:

Value	Reason
2	Log files could not be opened. The SSLCacheTraceLog or the SSLCacheErrorLog directive is not valid.
3	The AF_UNIX socket cannot be initialized. Use the SSLCachePortFilename directive to specify a different socket for the session ID cache daemon.
4	Sidd cannot switch to the configured user and group. Verify the values for the user and group directives.

- Solution: Provide a valid value for the directives and restart IBM HTTP Server.

Configuration messages

This topic contains error messages that might result due to configuration problems and provides solutions to help you troubleshoot these problems.

The following messages appear due to configuration problems:

- Message: **SSL0300E: Unable to allocate terminal node.**
- Message: **SSL0301E: Unable to allocate string value in node.**
- Message: **SSL0302E: Unable to allocate non terminal node.**
- Message: **SSL0303E: Syntax Error in SSLClientAuthGroup directive.**
- Message: **SSL0304E: Syntax Error in SSLClientAuthRequire directive.**
- Message: **SSL0307E: Invalid token preceding NOT or !**
- Message: **SSL0308E: A group is specified in SSLClientAuthRequire but no groups are specified.**
- Message: **SSL0309E: The group <group> is specified in SSLClientAuthRequire is not defined.**
- Message: **SSL0310I: Access denied to object due to invalid SSL version <version>, expected <version>.**
- Message: **SSL0311E: Unable to get cipher in checkBanCipher.**
- Message: **SSL0312I: Cipher <cipher> is in ban list and client is forbidden to access object.**
- Message: **SSL0313E: Fell through to default return in checkCipherBan.**
- Message: **SSL0314E: Cipher is NULL in checkRequireCipher.**
- Message: **SSL0315E: Cipher <cipher> used is not in the list of required ciphers to access this object.**
- Message: **SSL0316E: Fell through to default return in checkCipherRequire.**
- Message: **SSL0317E: Unable to allocate memory for fake basic authentication username.**
- Message: **SSL0318E: Limit exceeded for specified cipher specs, only 64 total allowed.**
 - Reason: The number of ciphers configured using the SSLCipherSpec directive exceeds the maximum allowed of 64.
 - Solution: Check for duplicate SSLCipherSpec directives.
- Message: **SSL0319E: Cipher Spec <cipher> is not supported by this GSK library.**
 - Reason: The cipher is not a valid cipher for use with the installed SSL libraries.
 - Solution: Check that a valid cipher value was entered with the SSLCipherSpec directive.
- Message: **SSL0320I: Using Version 213 Cipher: <cipher>.**
 - Reason: This is an informational message listing the ciphers used for connections to this virtual host.
 - Solution: None.
- Message: **SSL0321E: Invalid cipher spec <cipher>.**
 - Reason: The cipher is not a valid cipher.
 - Solution: Check the documentation for a list of valid cipher specs.

- Message: **SSL0322E: Cipher Spec <cipher> is not valid.**
 - Reason: The cipher is not a valid cipher.
 - Solution: Check the documentation for a list of valid cipher specs.
- Message: **SSL0323E: Cipher Spec <cipher> has already been added.**
 - Reason: A duplicate SSLCipherSpec directive has been encountered.
 - Solution: This instance of the directive is ignored and should be removed from the configuration file.
- Message: **SSL0324E: Unable to allocate storage for cipher specs.**
 - Reason: The server could not allocate memory needed to complete the operation.
 - Solution: Take action to free up some additional memory. Try reducing the number of threads or processes running, or increasing virtual memory.
- Message: **SSL0325E: Cipher Spec <cipher> has already been added to the v2lv3 banlrequire list.**
 - Reason: A duplicate cipher was specified on the SSLCipherBan directive.
 - Solution: This instance of the directive is ignored and should be removed from the configuration file.
- Message: **SSL0326E: Invalid cipher spec <cipher> set for SSLCipherBan|SSLCipherRequire.**
 - Reason: The cipher is not a valid cipher.
 - Solution: Check the documentation for a list of valid cipher specs.
- Message: **SSL0327E: Invalid value for sslv2timeout|sslv3timeout, using default value of nn seconds.**
 - Reason: The timeout value specified is not in the valid range.
 - Solution: Check the documentation for the proper range of values.
- Message: **SSL0328W: Invalid argument for SSLClientAuth: <args>. CRL can not be turned on unless Client Authentication is on.**
- Message: **SSL0329W: Invalid argument for SSLClientAuth: <args>. If a second argument is entered it must be CRL. CRL cannot be turned on unless client authentication is on.**
- Message: **SSL0330W: Invalid argument for SSLClientAuth: <args>. If a second value is entered it must be crl.**
- Message: **SSL0331W: Invalid argument for SSLClientAuth: <args>. The first value must be 0, 1, 2 none, optional, or required.**
- Message: **SSL0332E: Not enough arguments specified for SSLClientAuthGroup.**
- Message: **SSL0333E: No parse tree created for <parm>.**
 - Reason: An error occurred processing the SSLClientAuthRequire directive.
 - Solution: Check for other error messages. Enable tracing of Client Authentication by adding the directive SSLClientAuthRequireTraceOn to the configuration file.
- Message: **SSL0334E: Function ap_make_table failed processing label <certificate>.**
- Message: **SSL0337E: OCSP is not supported with this level of GSKit**
 - Reason: OCSP support requires GSKit 7.0.4.14 or higher
 - Solution: Upgrade the level of GSKit on the system to 7.0.4.14 or higher

Handshake messages

This topic contains error messages that might result due to SSL handshake failures and provides solutions to help you troubleshoot these problems.

The following messages display due to handshake failures:

- Message: **SSL0192W: IBM HTTP Server is configured to permit client renegotiation which is vulnerable to man-in-the-middle attacks <servername:port>**
 - Reason: IBM HTTP Server is configured to allow client handshake renegotiation using the SSLInsecureRenegotiation directive. This configuration is vulnerable to man-in-the middle attacks.

Use this configuration only if it is necessary for your client and be aware of the risk. For more information about the exposure, refer to the public documentation about CVE-2009-3555.

- Solution: Remove the SSLInsecureRenegotiation directive or set the directive to OFF to avoid the vulnerability. If proprietary clients require SSL renegotiation to function, update these clients to establish new connections.
- Message: **SSL0193W: Error setting GSK_NO_RENEGOTIATION to <GSK_TRUE | GSK_FALSE> <errorcode>**
 - Reason: An error occurred when the server attempted to disable client renegotiation. This setting is the default value. However, this value is also set if you specify the SSLInsecureRenegotiation directive with an OFF value.
 - Solution: Report this problem to IBM Support.
- Message: **SSL0196I: Security library does not support GSK_SESSION_RESET_CALLBACK, rejecting insecure SSL client renegotiation by monitoring SIDs**
 - Reason: When the server attempted to disable client renegotiation, it was determined that the security library on this system does not support GSK_SESSION_RESET_CALLBACK. It will be configured to reject insecure SSL client renegotiation using an alternate mechanism of monitoring SIDs.
 - Solution: This informational message does not indicate a failure, but it reports a configuration condition. An action is not necessary. You can upgrade to a newer z/OS security library that includes support for GSK_SESSION_RESET_CALLBACK or for disabling SSL client renegotiation.
- Message: **SSL0197I: Configured security library to reject insecure SSL client renegotiation.**
 - Reason: The security library has been successfully configured to reject client renegotiation.
 - Solution: This informational message does not indicate a failure, but it reports a particular configuration setting. An action is not necessary.
- Message: **SSL0198I: System is running without a security library capable of directly rejecting insecure SSL client renegotiation. Aborting HTTPS requests that span SSL sessions**
 - Reason: While the server attempted to disable client renegotiation, it was determined that the security library on this system does not support directly rejecting SSL client renegotiation. It will be configured to use an alternate callback mechanism.
 - Solution: This informational message does not indicate a failure, but it reports a configuration condition. An action is not necessary. For z/OS systems, upgrade to a newer security library that includes support for GSK_SESSION_RESET_CALLBACK or for disabling SSL client renegotiation. For distributed systems, upgrade to GSKit Version 7.0.4.27 or later.
- Message: **SSL0200E: Handshake Failed, <code>.**
 - Reason: The handshake failed when the SSL library returned an unknown error.
 - Solution: Report this problem to IBM Support.
- Message: **SSL0201E: Handshake Failed, Internal error - Bad handle.**
 - Reason: An internal error has occurred.
 - Solution: Report this problem to IBM Support.
- Message: **SSL0202E: Handshake Failed, The GSK library unloaded.**
 - Reason: A call to the GSKit function failed because the dynamic link library unloaded (Windows operating systems only).
 - Solution: Shut down the server and restart.
- Message: **SSL0203E: Handshake Failed, GSK internal error.**
 - Reason: The communication between client and the server failed due to an error in the GSKit library.
 - Solution: Retry connection from the client. If the error continues, report the problem to IBM Support.
- Message: **SSL0204E: Handshake Failed, Internal memory allocation failure.**
 - Reason: The server could not allocate memory needed to complete the operation.

- Solution: Take action to free up some additional memory. Try reducing the number of threads or processes running, or increasing virtual memory.
- Message: **SSL0205E: Handshake Failed, GSK handle is in an invalid state for operation.**
 - Reason: The SSL state for the connection is invalid.
 - Solution: Retry connection from the client. If the error continues, report the problem to IBM Support.
- Message: **SSL0206E: Handshake Failed, Key-file label not found**
 - Reason: The label specified for the SSLServerCert directive was not found in the key database (KDB) file specified for the KeyFile directive.
 - Solution: Specify a value for the SSLServerCert directive that corresponds to a personal certificate available in the KDB file specified for the KeyFile directive
- Message: **SSL0207E: Handshake Failed, Certificate is not available.**
 - Reason: The client did not send a certificate.
 - Solution: Set client authentication to optional if a client certificate is not required. Contact the client to determine why it is not sending an acceptable certificate.
- Message: **SSL0208E: Handshake Failed, Certificate validation error.**
 - Reason: The received certificate failed one of the validation checks.
 - Solution: Use another certificate. Contact IBM Support to determine why the certificate failed validation.
- Message: **SSL0209E: Handshake Failed, ERROR processing cryptography.**
 - Reason: A cryptography error occurred.
 - Solution: None. If the problem continues, report it to IBM Support.
- Message: **SSL0210E: Handshake Failed, ERROR validating ASN fields in certificate.**
 - Reason: The server was not able to validate one of the ASN fields in the certificate.
 - Solution: Try another certificate.
- Message: **SSL0211E: Handshake Failed, ERROR connecting to LDAP server.**
 - Reason: The Web server failed to connect to the CRL LDAP server.
 - Solution: Verify that the values entered for the SSLCRLHostname and SSLCRLPort directives are correct. If access to the CRL LDAP server requires authentication, is the SSLCRLUserID directive coded and was the password added to the stash file pointed to by the SSLStashfile directive.
- Message: **SSL0212E: Handshake Failed, Internal unknown error.**
 - Reason: An unknown error has occurred in the SSL library.
 - Solution: Report the problem to IBM Support.
- Message: **SSL0213E: Handshake Failed, Open failed due to cipher error.**
 - Reason: An unknown error has occurred in the SSL library.
 - Solution: Report the problem to IBM Support.
- Message: **SSL0214E: Handshake Failed, I/O error reading key file.**
 - Reason: The server could not read the key database file.
 - Solution: Check file access permissions and verify the Web server user ID is allowed access.
- Message: **SSL0215E: Handshake Failed, Key file has an invalid internal format. Recreate key file.**
 - Reason: Key file has an invalid format.
 - Solution: Recreate key file.
- Message: **SSL0216E: Handshake Failed, Key file has two entries with the same key. Use IKEYMAN to remove the duplicate key.**
 - Reason: Two identical keys exist in key file.
 - Solution: Use IKEYMAN to remove duplicate key.
- Message: **SSL0217E: Handshake Failed, Key file has two entries with the same label. Use IKEYMAN to remove the duplicate label.**

- Reason: A second certificate with the same label was placed in the key database file.
- Solution: Use IKEYMAN to remove duplicate label.
- Message: **SSL0218E: Handshake failed, Either the key file has become corrupted or the password is incorrect.**
 - Reason: The key file password is used as an integrity check and the test failed. Either the key database file is corrupted, or the password is incorrect.
 - Solution: Use IKEYMAN to stash the key database file password again. If that fails, recreate the key database.
- Message: **SSL0219E: SSL Handshake Failed, Either the default key in the keyfile has an expired certificate or the keyfile password expired. Use iKeyman to renew or remove certificates that are expired or to set a new keyfile password.**
 - Reason: Either the default key in the keyfile has an expired certificate or the keyfile password expired.
 - Solution: Use iKeyman to renew or remove certificates that are expired or to set a new keyfile password.
- Message: **SSL0220E: Handshake Failed, There was an error loading one of the GSKdynamic link libraries. Be sure GSK was installed correctly.**
 - Reason: Opening the SSL environment resulted in an error because one of the GSKdynamic link libraries could not load.
 - Solution: Contact Support to make sure the GSKit is installed correctly.
- Message: **SSL0221E: Handshake Failed, Either the certificate has expired or the system clock is incorrect.**
 - Reason: Either the certificate expired or the system clock is incorrect.
 - Solution: Use the key management utility (iKeyman) to recreate or renew your server certificate or change the system date to a valid date.
- Message: **SSL0222W: Handshake failed, no ciphers specified.**
 - Reason: SSLV2 and SSLV3 are disabled.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0223E: Handshake Failed, No certificate.**
 - Reason: The client did not send a certificate.

You can also see this message when your keyfile does not have a default certificate specified and you have not specified an SSLServerCert directive. It will pass initialization but fail at connection (handshake) time.
 - Solution: Set client authentication to optional if a client certificate is not required. Contact the client to determine why it is not sending a certificate.
- Message: **SSL0224E: Handshake failed, Invalid or improperly formatted certificate.**
 - Reason: The client did not specify a valid certificate.
 - Solution: Client problem.
- Message: **SSL0225E: Handshake Failed, Unsupported certificate type.**
 - Reason: The certificate type received from the client is not supported by this version of IBM HTTP Server SSL.
 - Solution: The client must use a different certificate type.
- Message: **SSL0226I: Handshake Failed, I/O error during handshake.**
 - Reason: The communication between the client and the server failed. This is a common error when the client closes the connection before the handshake has completed.
 - Solution: Retry the connection from the client.
- Message: **SSL0227E: Handshake Failed, Specified label could not be found in the key file.**
 - Reason: Specified key label is not present in key file.

- Solution: Check that the SSLServerCert directive is correct, if coded, and that the label is valid for one of the keys in the key database.
- Message: **SSL0228E: Handshake Failed, Invalid password for key file.**
 - Reason: The password retrieved from the stash file could not open the key database file.
 - Solution: Use IKEYMAN to open the key database file and recreate the password stash file. This problem can also result from a corrupted key database file. Creating a new key database file may resolve the problem.
- Message: **SSL0229E: Handshake Failed, Invalid key length for export.**
 - Reason: In a restricted cryptography environment, the key size is too long to be supported.
 - Solution: Select a certificate with a shorter key.
- Message: **SSL0230I: Handshake Failed, An incorrectly formatted SSL message was received.**
- Message: **SSL0231W: Handshake Failed, Could not verify MAC.**
 - Reason: The communication between the client and the server failed.
 - Solution: Retry the connection from the client.
- Message: **SSL0232W: Handshake Failed, Unsupported SSL protocol or unsupported certificate type.**
 - Reason: The communication between the client and the server failed because the client is trying to use a protocol or certificate which the IBM HTTP Server does not support.
 - Solution: Retry the connection from the client using an SSL Version 2 or 3, or TLS 1 protocol. Try another certificate.
- Message: **SSL0233W: Handshake Failed, Invalid certificate signature.**
- Message: **SSL0234W: Handshake Failed, The certificate sent by the peer expired or is invalid.**
 - Reason: The partner did not specify a valid certificate. The server is acting as a reverse proxy to an SSL URL and the `_server_cert` could not be validated.
 - Solution: Partner problem. If this occurs during an SSL Proxy connection, the remote SSL server sent a bad certificate to IBM HTTP Server. Check the certificate and certificate authority chain at the other end of the SSL connection. For more information, see “Securing with SSL communications” on page 90.
- Message: **SSL0235W: Handshake Failed, Invalid peer.**
- Message: **SSL0236W: Handshake Failed, Permission denied.**
- Message: **SSL0237W: Handshake Failed, The self-signed certificate is not valid.**
- Message: **SSL0238E: Handshake Failed, Internal error - read failed.**
 - Reason: The read failed.
 - Solution: None. Report this error to IBM Support.
- Message: **SSL0239E: Handshake Failed, Internal error - write failed.**
 - Reason: The write failed.
 - Solution: None. Report this error to IBM Support.
- Message: **SSL0240I: Handshake Failed, Socket has been closed.**
 - Reason: The client closed the socket before the protocol completed.
 - Solution: Retry connection between client and server.
- Message: **SSL0241E: Handshake Failed, Invalid SSLV2 Cipher Spec.**
 - Reason: The SSL Version 2 cipher specifications passed into the handshake were invalid.
 - Solution: Change the specified Version 2 cipher specs.
- Message: **SSL0242E: Handshake Failed, Invalid SSLV3 Cipher Spec.**
 - Reason: The SSL Version 3 cipher specifications passed into the handshake were invalid.
 - Solution: Change the specified Version 3 cipher specs.
- Message: **SSL0243E: Handshake Failed, Invalid security type.**

- Reason: There was an internal error in the SSL library.
- Solution: Retry the connection from the client. If the error continues, report the problem to IBM Support.
- Message: **SSL0245E: Handshake Failed, Internal error - SSL Handle creation failure.**
 - Reason: There was an internal error in the security libraries.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0246E: Handshake Failed, Internal error - GSK initialization has failed.**
 - Reason: An error in the security library has caused SSL initialization to fail.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0247E: Handshake Failed, LDAP server not available.**
 - Reason: Unable to access the specified LDAP directory when validating a certificate.
 - Solution: Check that the SSLCRLHostname and SSLCRLPort directives are correct. Make sure the LDAP server is available.
- Message: **SSL0248E: Handshake Failed, The specified key did not contain a private key.**
 - Reason: The key does not contain a private key.
 - Solution: Create a new key. If this was an imported key, include the private key when doing the export.
- Message: **SSL0249E: Handshake Failed, A failed attempt was made to load the specified PKCS#11 shared library.**
 - Reason: An error occurred while loading the PKCS#11 shared library.
 - Solution: Verify that the PKCS#11 shared library specified in the SSLPKCSDriver directive is valid.
- Message: **SSL0250E: Handshake Failed, The PKCS#11 driver failed to find the token label specified by the caller.**
 - Reason: The specified token was not found on the PKCS#11 device.
 - Solution: Check that the token label specified on the SSLServerCert directive is valid for your device.
- Message: **SSL0251E: Handshake Failed, A PKCS#11 token is not present for the slot.**
 - Reason: The PKCS#11 device has not been initialized correctly.
 - Solution: Specify a valid slot for the PKCS#11 token or initialize the device.
- Message: **SSL0252E: Handshake Failed, The password/pin to access the PKCS#11 token is either not present, or invalid.**
 - Reason: Specified user password and pin for PKCS#11 token is not present or invalid.
 - Solution: Check that the correct password was stashed using the SSLStash utility and that the SSLStashfile directive is correct.
- Message: **SSL0253E: Handshake Failed, The SSL header received was not a properly SSLV2 formatted header.**
 - Reason: The data received during the handshake does not conform to the SSLV2 protocol.
 - Solution: Retry connection between client and server. Verify that the client is using HTTPS.
- Message: **SSL0254E: Internal error - I/O failed, buffer size invalid.**
 - Reason: The buffer size in the call to the I/O function is zero or negative.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0255E: Handshake Failed, Operation would block.**
 - Reason: The I/O failed because the socket is in non-blocking mode.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0256E: Internal error - SSLV3 is required for reset_cipher, and the connection uses SSLV2.**
 - Reason: A reset_cipher function was attempted on an SSLV2 connection.
 - Solution: None. Report this problem to IBM Support.

- Message: **SSL0257E: Internal error - An invalid ID was specified for the gsk_secure_soc_misc function call.**
 - Reason: An invalid value was passed to the gsk_secure_soc_misc function.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0258E: Handshake Failed, The function call, <function>, has an invalid ID.**
 - Reason: An invalid function ID was passed to the specified function.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0259E: Handshake Failed, Internal error - The attribute has a negative length in: <function>.**
 - Reason: The length value passed to the function is negative, which is invalid.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0260E: Handshake Failed, The enumeration value is invalid for the specified enumeration type in: <function>.**
 - Reason: The function call contains an invalid function ID.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0261E: Handshake Failed, The SID cache is invalid: <function>.**
 - Reason: The function call contains an invalid parameter list for replacing the SID cache routines.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0262E: Handshake Failed, The attribute has an invalid numeric value: <function>.**
 - Reason: The function call contains an invalid value for the attribute being set.
 - Solution: None. Report this problem to IBM Support.
- Message: **SSL0263W: SSL Connection attempted when SSL did not initialize.**
 - Reason: A connection was received on an SSL-enabled virtual host but it could not be completed because there was an error during SSL initialization.
 - Solution: Check for an error message during startup and correct that problem.
- Message: **SSL0264E: Failure obtaining Cert data for label <certificate>.**
 - Reason: A GSKit error prevented the server certificate information from being retrieved.
 - Solution: Check for a previous error message with additional information.
- Message: **SSL0265W: Client did not supply a certificate.**
 - Reason: A client who connected failed to send a client certificate and the server is configured to require a certificate.
 - Solution: Nothing on the server side.
- Message: **SSL0266E: Handshake failed.**
 - Reason: Could not establish SSL proxy connection.
 - Solution: IBM HTTP Server could not establish a proxy connection to a remote server using SSL.
- Message: **SSL0267E: SSL Handshake failed.**
 - Reason: Timeout on network operation during handshake.
 - Solution: Check client connectivity, adjust TimeOuts.
- Message: **SSL0270I: SSL Handshake Failed, Timeout (dd seconds) occurred before any data received.**
 - Reason: A connection was received on an SSL port, but no data was received from the client before the timeout expired.
 - Solution: If the timeout (set by the Timeout directive) has been reduced from the default value, verify that it is reasonable. If the message occurs intermittently, it is probably normal, due to things like users cancelling page loads and browser or system crashes. If the message occurs in bursts, it might indicate a denial of service attack in progress.
- Message: **SSL0271I: SSL Handshake Failed, client closed connection without sending any data.**

- Reason: A connection was received on an SSL port, but the client closed the connection without beginning the handshake.
- Solution: If the timeout (set by the Timeout directive) has been reduced from the default value, verify that it is reasonable. If the message occurs intermittently, it is probably normal, due to things like users cancelling page loads and browser or system crashes. If the message occurs in bursts, it might indicate a denial of service attack in progress.
- Message: **SSL0272I: SSL Handshake Failed, I/O error before any data received.**
 - Reason: A connection was received on an SSL port, but a network error broke the connection before any data was received from the client.
 - Solution: If the message occurs intermittently, it is probably normal, due to things like users cancelling page loads and browser or system crashes. If the message occurs in bursts, it might indicate a denial of service attack in progress.
- Message: **SSL0273I: Non-SSL request received on connection configured for SSL**
 - Reason: A connection was received on an SSL port, but the data received was not SSL, and looked like a normal non-SSL request.
 - Solution: Verify that the port in question is intended to be configured for SSL. Look for bad links to the page in question that should use https:, but instead use http:.
- Message: **SSL0273I: Non-SSL request received on connection configured for SSL**
 - Reason: A connection was received on an SSL port, but the data received was not SSL, and looked like a normal non-SSL request.
 - Solution: Verify that the port in question is intended to be configured for SSL. Look for bad links to the page in question that should use https:, but instead use http:.
- Message: **SSL0276E: SSL: Unexpected SSL client renegotiation detected, aborting SSL connection.**
 - Reason: SSL client renegotiation was attempted, but the configuration does not allow SSL renegotiation. Thus, the SSL connection was stopped.
 - Solution: Retry the connection between the client and the server. Configure the connection to allow SSL renegotiation only if necessary. Be aware of the risk. If proprietary clients require SSL renegotiation to function, update them to establish new connections.

SSL initialization messages

This topic contains error messages that might result due to SSL initialization problems and provides solutions to help you troubleshoot these problems.

The following messages display as a result of initialization problems:

- Message: **SSL0100E: GSK could not initialize, <errorCode>**
 - Reason: Initialization failed when the SSL library returned an unknown error.
 - Solution: None. Report this problem to Service.
- Message: **SSL0101E: GSK could not initialize, Neither the password nor the stash file name was specified. Could not open key file.**
 - Reason: The stash file for the key database could not be found or is corrupted.
 - Solution: Use IKEYMAN to open the key database file and recreate the password stash file.
- Message: **SSL0102E: GSK could not initialize, Could not open key file.**
 - Reason: The server could not open the key database file.
 - Solution: Check that the Keyfile directive is correct and that the file permissions allow the Web server user ID to access the file.
- Message: **SSL0103E: Internal error - GSK could not initialize, Unable to generate a temporary key pair.**
 - Reason: GSK could not initialize; Unable to generate a temporary key pair.

- Solution: Report this problem to Service.
- Message: **SSL0104E: GSK could not initialize, Invalid password for key file.**
 - Reason: The password retrieved from the stash file could not open the key database file.
 - Solution: Use IKEYMAN to open the key database file and recreate the password stash file. This problem could also result from a corrupted key database file. Creating a new key database file may resolve the problem.
- Message: **SSL0105E: GSK could not initialize, Invalid label.**
 - Reason: Specified key label is not present in key file.
 - Solution: Check that the SSLServerCert directive is correct, if coded, and that the label is valid for one of the keys in the key database.
- Message: **SSL0106E: Initialization error, Internal error - Bad handle**
 - Reason: An internal error has occurred.
 - Solution: Report this problem to Service.
- Message: **SSL0107E: Initialization error, The GSK library unloaded.**
 - Reason: A call to the GSKit function failed because the dynamic link library unloaded (Windows only).
 - Solution: Shut down the server and restart.
- Message: **SSL0108E: Initialization error, GSK internal error.**
 - Reason: The communication between client and the server failed due to an error in the GSKit library.
 - Solution: Retry connection from the client. If the error continues, report the problem to Service.
- Message: **SSL0109E: GSK could not initialize, Internal memory allocation failure.**
 - Reason: The server could not allocate memory needed to complete the operation.
 - Solution: Take action to free up some additional memory. Try reducing the number of threads or processes running, or increasing virtual memory.
- Message :**SSL0110E: Initialization error, GSK handle is in an invalid state for operation.**
 - Reason: The SSL state for the connection is invalid.
 - Solution: Retry connection from the client. If the error continues, report the problem to Service.
- Message: **SSL0111E: Initialization error, Key file label not found.**
 - Reason: Certificate or key label specified was not valid.
 - Solution: Verify that the certificate name specified with the SSLServerCert directive is correct or, if no SSLServerCert directive was coded, that a default certificate exists in the key database.
- Message: **SSL0112E: Initialization error, Certificate is not available.**
 - Reason: The client did not send a certificate.
 - Solution: Set Client Authentication to optional if a client certificate is not required. Contact the client to determine why it is not sending an acceptable certificate.
- Message: **SSL0113E: Initialization error, Certificate validation error.**
 - Reason: The received certificate failed one of the validation checks.
 - Solution: Use another certificate. Contact Service to determine why the certificate failed validation.
- Message: **SSL0114E: Initialization error, Error processing cryptography.**
 - Reason: A cryptography error occurred.
 - Solution: None. If the problem continues, report it to Service.
- Message: **SSL0115E: Initialization error, Error validating ASN fields in certificate.**
 - Reason: The server was not able to validate one of the ASN fields in the certificate.
 - Solution: Try another certificate.
- Message: **SSL0116E: Initialization error, Error connecting to LDAP server.**

- Reason: The Web server failed to connect to the CRL LDAP server.
- Solution: Verify that the values entered for the SSLCRLHostname and SSLCRLPort directives are correct. If access to the CRL LDAP server requires authentication, is the SSLCRLUserID directive coded and was the password added to the stash file pointed to by the SSLStashfile directive.
- Message: **SSL0117E: Initialization error, Internal unknown error. Report problem to service.**
 - Reason: Initialization error, Internal unknown error. Report problem to service.
 - Solution: Initialization error, Internal unknown error. Report problem to service.
- Message: **SSL0118E: Initialization error, Open failed due to cipher error.**
 - Reason: Report problem to service.
 - Solution: Report problem to service.
- Message: **SSL0119E: Initialization error, I/O error reading keyfile.**
 - Reason: I/O error trying to read SSL keyfile.
 - Solution: Check the file permissions for keyfile.
- Message: **SSL0120E: Initialization error, Keyfile has and invalid internal format. Recreate keyfile.**
 - Reason: Initialization error, the keyfile has an invalid internal format. Recreate the keyfile.
 - Solution: Verify the keyfile is not corrupted.
- Message: **SSL0121E: Initialization error, Keyfile has two entries with the same key. Use Ikeyman to remove the duplicate key.**
 - Reason: The keyfile has two entries with the same key. Use Ikeyman to remove the duplicate key.
 - Solution: Use Ikeyman to remove the duplicate key.
- Message: **SSL0122E: Initialization error, Keyfile has two entries with the same label. Use Ikeyman to remove the duplicate label.**
 - Reason: The keyfile has two entries with the same label. Use Ikeyman to remove the duplicate label.
 - Solution: Use Ikeyman to remove the duplicate label.
- Message: **SSL0123E: Initialization error, The keyfile password is used as an integrity check. Either the keyfile has become corrupted or the password is incorrect.**
 - Reason: The keyfile password is used as an integrity check. Either the keyfile has become corrupted or the password is incorrect.
 - Solution: Use Ikeyman to verify that the keyfile is valid, check permissions on the stash file, verify passwords.
- Message: **SSL0124E: SSL Handshake Failed, Either the default key in the keyfile has an expired certificate or the keyfile password expired. Use iKeyman to renew or remove certificates that are expired or to set a new keyfile password.**
 - Reason: Either the default key in the keyfile has an expired certificate or the keyfile password expired.
 - Solution: Use iKeyman to renew or remove certificates that are expired or to set a new keyfile password.
- Message: **SSL0125E: Initialization error, There was an error loading one of the GSK dynamic link libraries. Be sure GSK is installed correctly.**
 - Reason: There was an error loading one of the GSK dynamic link libraries. Be sure GSK is installed correctly.
 - Solution: Verify GSK is installed and appropriate level for release of IBM HTTP Server.
- Message: **SSL0126E: Handshake Failed, Either the certificate has expired or the system clock is incorrect.**
 - Reason: Either the certificate expired or the system clock is incorrect.
 - Solution: Use the key management utility (iKeyman) to recreate or renew your server certificate or change the system date to a valid date.
- Message: **SSL0127E: Initialization error, No ciphers specified.**

- Reason: Initialization error, no ciphers specified.
 - Solution: Report problem to service.
 - Message: **SSL0128E: Initialization error, Either the certificate expired or the system clock is incorrect.**
 - Reason: Initialization error, no certificate.
 - Solution: Report problem to service.
 - Message: **SSL0129E: Initialization error, The received certificate was formatted incorrectly.**
 - Reason: The received certificate is formatted incorrectly.
 - Solution: Use Ikeyman to validate certificates used for connection.
 - Message: **SSL0130E: Initialization error, Unsupported certificate type.**
 - Reason: Unsupported certificate type.
 - Solution: Check certificates that are used for this connection in Ikeyman.
 - Message: **SSL0131I: Initialization error, I/O error during handshake.**
 - Reason: I/O error during handshake.
 - Solution: Check network connectivity.
 - Message: **SSL0132E: Initialization error, Invalid key length for export.**
 - Reason: Invalid key length for export.
 - Solution: Report problem to service.
 - Message: **SSL0133W: Initialization error, An incorrectly formatted SSL message was received.**
 - Reason: An incorrectly formatted SSL message was received.
 - Solution: Check client settings.
 - Message: **SSL0134W: Initialization error, Could not verify MAC.**
 - Reason: Could not verify MAC.
 - Solution: Report problem to service.
 - Message: **SSL0135W: Initialization error, Unsupported SSL protocol or unsupported certificate type.**
 - Reason: Unsupported SSL protocol or unsupported certificate type.
 - Solution: Check server ciphers and certificate settings.
 - Message: **SSL0136W: Initialization error, Invalid certificate signature.**
 - Reason: Invalid certificate signature.
 - Solution: Check certificate in Ikeyman.
 - Message: **SSL0137W: Initialization error, Invalid certificate sent by partner.**
 - Reason: Invalid certificate sent by partner.
 - Solution: If this occurs during an SSL Proxy connection, the remote SSL server sent a bad certificate to IBM HTTP Server. Check the certificate and certificate authority chain at the other end of the SSL connection.
 - Message: **SSL0138W: Initialization error, Invalid peer.**
 - Reason: Invalid peer.
 - Solution: Report problem to service.
 - Message: **SSL0139W: Initialization error, Permission denied.** Distributed operating systems
 - Reason: Permission denied.
 - Solution: Report problem to service.
- z/OS
- Reason: If a System Authorization Facility (SAF) SSL keyring is in use, the current user ID is not authorized to read the keyring.

- Solution: See the information about access to SAF keyrings in “Performing required z/OS system configurations” on page 1
- Message: **SSL0140W: Initialization error, The self-signed certificate is not valid.**
 - Reason: The self-signed certificate is not valid.
 - Solution: Check the certificate in Ikeyman.
- Message: **SSL0141E: Initialization error, Internal error - read failed.**
 - Reason: Internal error - read failed.
 - Solution: Report to service.
- Message: **SSL0142E: Initialization error, Internal error - write failed.**
 - Reason: Internal error - write failed.
 - Solution: Report to service.
- Message: **SSL0143I: Initialization error, Socket has been closed.**
 - Reason: Socket has been closed unexpectedly.
 - Solution: Check the client and network. Report problem to service.
- Message: **SSL0144E: Initialization error, Invalid SSLV2 Cipher Spec.**
 - Reason: Invalid SSLV2 cipher spec.
 - Solution: Check the SSLCipherSpec directive.
- Message: **SSL0145E: Initialization error, Invalid SSLV3 Cipher Spec.**
 - Reason: Invalid SSLV3 Cipher Spec.
 - Solution: Check the SSLCipherSpec directive.
- Message: **SSL0146E: Initialization error, Invalid security type.**
 - Reason: Invalid security type.
 - Solution: Report to service.
- Message: **SSL0147E: Initialization error, Invalid security type combination.**
 - Reason: Invalid security type combination.
 - Solution: Report to service.
- Message: **SSL0148E: Initialization error, Internal error - SSL Handle creation failure.**
 - Reason: Internal error - SSL handle creation failure.
 - Solution: Report to service.
- Message: **SSL0149E: Initialization error, Internal error - GSK initialization has failed.**
 - Reason: Internal error - GSK initialization has failed.
 - Solution: Report to service.
- Message: **SSL0150E: Initialization error, LDAP server not available.**
 - Reason: LDAP server not available.
 - Solution: Check CRL directives.
- Message: **SSL0151E: Initialization error, The specified key did not contain a private key.**
 - Reason: The specified key did not contain a private key.
 - Solution: Check the certificate in use in Ikeyman.
- Message: **SSL0152E: Initialization error, A failed attempt was made to load the specified PKCS#11 shared library.**
 - Reason: A failed attempt was made to load the specified PKCS#11 shared library.
 - Solution: Check SSLPKCSDriver directive and file system.
- Message: **SSL0153E: Initialization error, The PKCS#11 driver failed to find the token specified by the caller.**
 - Reason: The PKCS#11 driver failed to find the token specified by the caller.

- Message: **SSL0154E: Initialization error, A PKCS#11 token is not present for the slot.**
 - Reason: A PKCS#11 token is not present for the slot.
 - Solution: Verify PKCS#11 directives.
- Message: **SSL0155E: Initialization error, The password/pin to access the PKCS#11 token is invalid.**
 - Reason: The password and pin to access the PKCS#11 token is invalid.
- Message: **SSL0156E: Initialization error, The SSL header received was not a properly SSLV2 formatted header.**
 - Reason: The SSL header received was not a properly SSLV2 formatted header.
- Message: **SSL0157E: Initialization error, The function call, %s, has an invalid ID.**
 - Reason: The function call, %s, has an invalid ID.
 - Solution: Report problem to service.
- Message: **SSL0158E: Initialization error, Internal error - The attribute has a negative length: %s.**
 - Reason: Internal error - The attribute has a negative length.
 - Solution: Report problem to service.
- Message: **SSL0159E: Initialization error, The enumeration value is invalid for the specified enumeration type: %s.**
 - Reason: The enumeration value is invalid for the specified enumeration type: %s.
 - Solution: Report problem to service.
- Message: **SSL0160E: Initialization error, The SID cache is invalid: %s.**
 - Reason: The SID cache is invalid.
 - Solution: Report problem to service.
- Message: **SSL0161E: Initialization error, The attribute has an invalid numeric value: %s.**
 - Reason: The attribute has an invalid numeric value: %s.
 - Solution: Check SSL directives.
- Message: **SSL0162W: Setting the LD_LIBRARY_PATH or LIBPATH for GSK failed.**
 - Reason: Could not update the environment for GSK libraries.
 - Solution: Report problem to service.
- Message: **SSL0163W: Setting the LIBPATH for GSK failed, could not append /usr/opt/ibm/gskkm/lib.**
 - Reason: Could not append to LD_LIBRARY_PATH or LIBPATH for GSK failed.
 - Solution: Report problem to service.
- Message: **SSL0164W: Error accessing Registry, RegOpenKeyEx/RegQueryValueEx returned [%d].**
 - Reason: Error accessing registry.
 - Solution: Check GSK installation and windows registry.
- Message: **SSL0165W: Storage allocation failed.**
 - Reason: Storage allocation failed.
 - Solution: Check memory usage, report problem to service.
- Message: **SSL0166E: Failure attempting to load GSK library.**
 - Reason: Failure while attempting to load GSK library.
 - Solution: Check the GSK installation.
- Message: **SSL0167E: GSK function address undefined.**
 - Reason: GSK function address is undefined.
 - Solution: Check the GSK installation and level.
- Message: **SSL0168E: SSL initialization for server: %s, port: %u failed due to a configuration error.**

- Reason: Initialization for server: %s, port: %u failed due to a configuration error.
- Solution: Check the SSL configuration.
- Message: **SSL0169E: Keyfile does not exist: %s.**
 - Reason: Keyfile does not exist.
 - Solution: Check to ensure the path that is provided to the KeyFile directive exists, and is readable by the user that IBM HTTP Server is running as.
- Message: **SSL0170E: GSK could not initialize, no keyfile specified.**
 - Reason: Keyfile is not specified.
 - Solution: Specify Keyfile directive.
- Message: **SSL0171E: CRL cannot be specified as an option for the SSLClientAuth directive on HPUX because the IBM HTTP Server does not support CRL on HPUX.**
 - Reason: CRL cannot be specified as an option for the SSLClientAuth directive on HPUX because IBM HTTP Server does not support CRL on HPUX.
 - Solution: Remove CRL directives.
- Message: **SSL0172E: If CRL is turned on, you must specify an LDAP hostname for the SSLCRLHostname directive.**
 - Reason: If CRL is turned on, you must specify an LDAP hostname for the SSLCRLHostname directive.
 - Solution: Specify SSLCRLHostname.
- Message: **SSL0173E: Failure obtaining supported cipher specs from the GSK library.**
 - Reason: Failure obtaining supported cipher specs from the GSK library.
 - Solution: Check the GSK installation, report problem to service.
- Message: **SSL0174I: No CRL password found in the stash file: %s.**
 - Reason: No CRL password is found in the stash file: %s.
 - Solution: Check the stash file permissions, regenerate stash file.
- Message: **SSL0174I: No CRYPTO password found in the stash file: %s.**
 - Reason: No CRYPTO password is found in the stash file: %s.
 - Solution: Check stash file permissions, regenerate stash file.
- Message: **SSL0175E: fopen failed for stash file: %s.**
 - Reason: fopen failed for stash file.
 - Solution: Check stash file permissions, regenerate stash file.
- Message: **SSL0176E: fread failed for the stash file: %s.**
 - Reason: fread failed for the stash file.
 - Solution: Make sure the stash file is readable by user IBM HTTP Server is running as.
- Message: **SSL0179E: Unknown return code from stash_recover(), %d.**
 - Reason: Unknown return code from stash_recover(), %d.
 - Solution: Check the stash file.
- Message: **SSL0181E: Unable to fork for startup of session ID cache.**
 - Reason: Unable to fork for startup of session ID cache.
 - Solution: Check the location of sidd daemon, file permissions.
- Message: **SSL0182E: Error creating file mapped memory for SSL passwords.**
 - Reason: Error creating file mapped memory for SSL passwords.
 - Solution: Report problem to service.
- Message: **SSL0183E: Exceeded map memory limits.**
 - Reason: Exceeded map memory limits.
 - Solution: Report problem to service.

- Message: **SSL0184E: Could not find a password for the resource: %s.**
 - Reason: SSL0184E: Could not find a password for the resource: %s.
 - Solution: Report problem to service, disable password prompting.
- Message: **SSL0185E: ssl_getpwd() failed, unable to obtain memory.**
 - Reason: ssl_getpwd() failed, unable to obtain memory.
 - Solution: Report problem to service, disable password prompting.
- Message: **SSL0186E: Linked list mismatch.**
 - Reason: SSL0186E: Linked list mismatch.
 - Solution: Report problem to service, disable password prompting.
- Message: **SSL0186E: ssl_getpwd() failed, password exceeded maximum size of 4095.**
 - Reason: ssl_getpwd() failed, password exceeded the maximum size of 4095.
 - Solution: The password must be smaller than 4K.
- Message: **SSL0187E: It is invalid to enable password prompting for the SSLServerCert directive without specifying a Crypto Card Token.**
 - Reason: It is invalid to enable password prompting for the SSLServerCert directive without specifying a crypto card token.
 - Solution: Specify a crypto card token or disable password prompting for the SSLServerCert directive.
- Message: **SSL0188E: SSL initialization for server: %s, port: %u failed. SSL timeouts cannot be set in a virtualhost when the SSLCacheDisable directive has not been specified globally.**
 - Reason: When the SSL session cache is being used, only the global timeout settings apply because they are managed by the external session cache daemon. See information about the SSLCacheDisable and SSLCacheEnable directives in the information center topic entitled *SSL directives*.
 - Solution: If separate SSL timeouts are required, disable use of the session ID cache (SSLCacheDisable), otherwise make sure the SSLV3Timeout and SSLV2Timeout directives are only set in the global scope.

I/O error messages

This topic contains error messages that might result due to I/O failures and provides solutions to help you troubleshoot these problems.

The following messages appear due to read failures:

- Message: **SSL0400I: I/O failed, RC <code>.**
 - Reason: The server received an error trying to read on the socket.
 - Solution: Some errors are expected during normal processing, especially a '406' error, which you can ignore. If you are unable to access the server and receive these errors, report this problem to Service.
- Message: **SSL0401E: I/O failed with invalid handle <handle>.**
 - Reason: An internal error has occurred.
 - Solution: Report this problem to Service.
- Message: **SSL0402E: I/O failed, the GSKit library is not available.**
 - Reason: A call to the GSKit function failed because the dynamic link library unloaded (Windows operating systems only).
 - Solution: Shut down the server and restart.
- Message: **SSL0403E: I/O failed, internal error.**
 - Reason: The communication between client and the server failed due to an error in the GSKit library.
 - Solution: Retry connection from the client. If the error continues, report the problem to Service.

- Message: **SSL0404E: I/O failed, insufficient storage.**
 - Reason: The server could not allocate memory needed to complete the operation.
 - Solution: Take action to free up some additional memory. Try reducing the number of threads or processes running, or increasing virtual memory.
- Message: **SSL0405E: I/O failed, SSL handle <handle> is in an invalid state.**
 - Reason: The SSL state for the connection is invalid.
 - Solution: Retry connection from the client. If the error continues, report the problem to Service.
- Message: **SSL0406E: I/O failed, cryptography error.**
 - Reason: A cryptography error occurred.
 - Solution: None. If the problem continues, report it to Service.
- Message: **SSL0407I: I/O failed, Error validating ASN fields in certificate.**
 - Reason: The server was not able to validate one of the ASN fields in the certificate.
 - Solution: Try another certificate.
- Message: **SSL0408E: I/O failed with invalid buffer size. Buffer <address>, size <length>.**
 - Reason: The buffer size in the call to the read function is zero or negative.
 - Solution: None. Report this problem to Service.
- Message: **SSL0409I: I/O error occurred**
 - Reason: An unexpected network error occurred while reading or writing data over an SSL connection, likely a client disconnecting.
 - Solution: This is an informational message that does not indicate any failure in delivering a response, therefore no solution is provided.
- Message: **SSL0410I: Socket was closed**
 - Reason: An SSL client connection was closed by the client.
 - Solution: This is an informational message that does not indicate any failure in delivering a response, therefore a solution is not provided.
- Message: **SSL0411E: Connection aborted due to unexpected client renegotiation or other malformed SSL record <errorcode>**
 - Reason: An unexpected client renegotiation has been detected or an incorrectly formatted SSL message has been received. Thus, the SSL connection has been stopped.
 - Solution: Check the client settings and retry connection between the client and the server.

SSL stash utility messages

This topic contains error messages that might result due to Secure Sockets Layer (SSL) stash utility problems and provides solutions to help you troubleshoot these problems.

The following messages appear due to SSL Stash utility errors:

- Message: **SSL0700S: Invalid function <function>**
 - Reason: An invalid parameter was entered. The valid values are `curl` or `crypto`.
 - Solution: Rerun the command with the proper function.
- Message: **SSL0701S: The password was not entered.**
 - Reason: The password was not entered on the command line.
 - Solution: Rerun the command with the password added.
- Message: **SSL0702S: Password exceeds the allowed length of 512.**
 - Reason: The password that was entered is longer than the allowed maximum of 512 characters.
 - Solution: Use a shorter password.

Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

APACHE INFORMATION. This information may include all or portions of information which IBM obtained under the terms and conditions of the Apache License Version 2.0, January 2004. The information may also consist of voluntary contributions made by many individuals to the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org>. You may obtain a copy of the Apache License at <http://www.apache.org/licenses/LICENSE-2.0>.

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

IBM Director of Intellectual Property & Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

Trademarks and service marks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. For a current list of IBM trademarks, visit the IBM Copyright and trademark information Web site (www.ibm.com/legal/copytrade.shtml).

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

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

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

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

Index

C

- certificate authorities 151
- commands
 - apachectl 6
 - gskcmd 143
 - setupadm 38
- cryptographic hardware
 - SSL 155
 - troubleshooting 198

F

- Fast Response Cache Accelerator 17
- FRCA 17
 - advanced 20
 - AIX 20
 - caching restrictions 18
 - logging 17
 - operational restrictions 19
 - starting
 - Windows services 7

H

- HTTP Server
 - Apache 14

I

- IBM HTTP Server 37, 38
 - administering 198
 - Apache 9, 86
 - changing database passwords 134
 - error messages 201
 - IBM 86
 - installation 51
 - z/OS 71, 72, 78
 - login failure 200
 - migration
 - z/OS 71
 - mounting CD-ROMS 49
 - new functions 85

- IBM HTTP Server (*continued*)
 - starting
 - administrative console 4, 5
 - system configurations
 - z/OS 1
 - third-party plug-ins 34
 - troubleshooting 195
 - Windows 195
 - z/OS 196
 - uninstallation
 - GUI 65
 - z/OS 82
 - updating 55
 - Windows 52
 - IKEYMAN
 - PKCS11 devices 157
 - starting 133
 - IPv4 support
 - HTTP Server 16
 - IPv6 support
 - HTTP Server 16

L

- LDAP 162
 - configuration 185

M

- messages
 - cache errors 201
 - configuration errors 202
 - handshake errors 203
 - I/O errors 217
 - SSL initialization 210
 - SSL stach utility errors 218

U

- utilities
 - htpasswd 38