BEA AquaLogic® Service Bus

SFTP Transport User's Guide

Version: 3.0
Revised: January 2008

Contents

Introduction

Using the SFTP Transport

Introduction

The SFTP transport is a poll-based transport that allows you to transfer files securely over the SSH File Transfer Protocol (SFTP) using SSH version 2. It polls a specified directory at regular intervals based on a predefined polling interval. After authentication, a connection is established between AquaLogic Service Bus (ALSB) services and the SFTP server, enabling file transfer. The SFTP transport supports one-way inbound and outbound connectivity.

Key Features

- **Supported service types**

  The SFTP transport is available for the following service types:

  – Messaging service (with request message type specified)

  – Any XML service

  For more information about configuring service types, see Business Services: Creating and Managing and Proxy Services: Creating and Managing in *Using the AquaLogic Service Bus Console*.

- **Large messages**

  The SFTP transport supports processing of large messages. When you configure a proxy service, you can enable content streaming and specify whether large messages must be buffered in memory or in a disk file. For more information, see Streaming Body Content in *AquaLogic Service Bus User Guide*.

- **Quality of Service (QoS)**

  For inbound message transfer, the QoS is set to **`exactly-once`**, which ensures that every message is processed at least once.

  For outbound message processing, the QoS is **`best-effort`**.

  Note: For messages that are not transferred, you must create the error-handling logic (including any retry logic) in the pipeline error handler. For more information, see Proxy Services: Error Handlers in *Using the AquaLogic Service Bus Console*.

  For more information about QoS in ALSB messaging, see Modeling Message Flow in ALSB in *AquaLogic Service Bus User Guide*.

Environment Values

Environment values are predefined fields in the configuration data and are likely to change when you move the configuration from one domain to another (for example, from test to production). The following table lists the environment values associated with the SFTP transport.

Table 1-1 Environment Values

Environment Value	Description
Archive Directory	The directory to which the files are moved from either the download directory or the remote location.
Download Directory	The directory on your local machine to which files are downloaded during the file transfer.
Error Directory	The location where messages are posted if there is a problem.
Managed Server for Polling	The managed server that is used for polling (in a cluster scenario).

For more information, see:

- Customization in *Using the AquaLogic Service Bus Console*
- “Configuring Proxy Services” on page 2-4
- “Configuring Business Services” on page 2-9.

General Principles of SFTP Authentication

The following principles are applicable to the SFTP authentication process for both proxy and business services:

- **Connection**: The ALSB service (proxy and business) always acts as the SFTP client and connects to the SFTP server.
- **Authentication by the SFTP server**
  - For public key and host-based authentication, the SFTP server authenticates the connection with the public key of the AquaLogic Service Bus service.
  - For username-password authentication, the SFTP server authenticates the connection with the username and password.
- **Authentication by the SFTP client:** The ALSB service always authenticates the SFTP server with the public-key/host/IP combination defined in the **`known_hosts`** file. For more information, see "Creating the Known Hosts File" on page 2-2.
- **Connection establishment:** The connection is established only when both the server and client authentications are successful.
- **Transfer**
  - If the client is a proxy service, the file (message) is downloaded from the SFTP server.
  - If the client is a business service, the file (message) is uploaded to the SFTP server.

Run-Time Behavior

Transferring files by using the SFTP transport involves the following steps:

1. The proxy service polls the input directory at regular intervals.

   Note: A new connection is created for each poll cycle.

2. If the proxy service finds a file in the input directory, it renames the file with a **`.stage`** extension. This renaming ensures that the service does not pick up the same files during the next polling cycle.

   The **`.stage`** file exists in the input directory until it is delivered.

   Note: If the file cannot be retrieved from the input directory (due to network failure, for example), the **`.stage`** file is processed during a clean-up cycle. The clean-up cycle is performed every 15 minutes or after 15 polling cycles, whichever occurs later. If a **`.stage`** file exists during two consecutive clean-up cycles, it is processed again.

3. A JMS task is created for the message and added to the domain-wide JMS queue.
4. A domain-wide MDB receives the task and processes the message.

   Note: The task uses a pooled connection for processing the message. If a connection is not available in the pool, a new connection is created.

5. The message is delivered to the pipeline and the **`.stage`** file is deleted.

   Note: If an SFTP business service is configured, the service puts the message in the outbound directory through a pooled connection.

   In the message is not delivered, the transport attempts to transfer it again and repeats the process up to a predefined number of attempts. If the message cannot be delivered, it is moved to the error directory.

Using the SFTP Transport

You can use the SFTP transport to transfer files securely using SSH File Transfer Protocol (SFTP).

The following sections describe how you can use the SFTP transport to transfer files securely:

- Enabling SFTP Authentication: This section describes the authentication methods that the SFTP transport supports.
- Configuring Proxy Services and Configuring Business Services: These sections describe how you can configure proxy and business services to use the SFTP transport.
- Handling Communication Errors and Troubleshooting: These sections provide information to help you solve problems that may occur while configuring or using the SFTP transport.
- Importing Resources: This section describes the SFTP services-specific policy and configuration details that you can preserve when you import resources.
- Importing and Publishing Services: UDDI Registries: This section lists the properties that are published when SFTP services are published to UDDI registries. It also lists the properties that are imported when SFTP services are imported from UDDI registries.

Enabling SFTP Authentication

The SFTP transport supports the following authentication methods:

- Username-password authentication
- Host-based authentication

- Public key authentication

ALSB services authenticate the SFTP server based on the server details defined in a **`known_hosts`** file. So to enable server authentication, you must create a **`known-hosts`** file on the client machine.

Creating the Known Hosts File

The **`known_hosts`** file must exist in the server on which the ALSB proxy services (inbound requests) or business services (outbound requests) run. The file must contain the host name, IP address, and public key of the remote SFTP servers to which the proxy service or business service can connect.

1. Create a **`known_hosts`** file and enter details in the following format:

   `Hostname,IP algorithm publickey`

   – *`Hostname`* is the host name of the SFTP server.

   – *`IP`* is the IP address of the SFTP server.

   – *`algorithm`* can be either DSA or RSA, based on the SFTP server configuration. Enter **`ssh-rsa`** or **`ssh-dss`** depending on the algorithm that is supported.

   – *`publickey`* is the public key of the SFTP server. It must be in the Open SSH public key format.

   Example:

   ```
   getafix,172.22.52.130 ssh-rsa
   AAAAB3NzaC1yc2EAAAABIwAAAIEAtR+M3Z9HFxnKZTx66fZdnQqAHQcF1vQe1+EjJ/HWYtg
   Anqsn0hMJzqWMatb/u9yFwUpZBirjm3g2I9Qd8VocmeHwoGPhDGfQ5LQ/PPo3esE+CGwdnC
   OyRCktNHeuKxo4kiCCJ/bph5dRpghCQIvsQvRE3sks+XwQ7Wuswz8pv58=
   ```

   The **`known_hosts`** file can contain multiple entries, but each entry must be on a separate line.

2. Move the **`known_hosts`** file to the **`$DOMAIN_HOME\alsb\transports\sftp`** directory.

Enabling Username-Password Authentication

Username-password authentication is the simplest and quickest method of authentication. It is based on the credentials of the user.

To enable username and password authentication for a service:

1. Create a static service account by using the user credentials on the SFTP server. For more information, see Service Accounts in *Using the AquaLogic Service Bus Console*.
2. Create a **`known_hosts`** file. For more information, see "Creating the Known Hosts File" on page 2-2.

Enabling Host-Based Authentication

Host-based authentication uses a private host key. This method can be used when all the users share a private host.

To enable host-based authentication for a service:

1. Configure a service key provider with an SSL client authentication key. For more information, see Service Key Providers.
2. Create a **`known_hosts`** file. For more information, see "Creating the Known Hosts File" on page 2-2.
3. Configure the SFTP server to accept requests from ALSB, which is a client to the SFTP server.

   For example, for an SFTP server on Linux, you must do the following:

   – Edit the **`/etc/ssh/shosts.equiv`** file and add the host name or IP address of the machine on which the AquaLogic Service Bus domain runs.

   – Edit the **`/etc/ssh/ssh_known_hosts`** file and add the host name or IP address of the machine on which the AquaLogic Service Bus domain runs, followed by a space and the public key.

   Note: You can extract the public key from the key store that was used while creating the service key provider. The public key must be in the Open SSH format.

Enabling Public Key Authentication

Public key authentication is performed using your own private key. This method can be used when each user has a private key.

To enable public key authentication:

1. Configure a service key provider with SSL client authentication key. For more information, see Service Key Providers.
2. Configure the SFTP server to accept requests from ALSB (SFTP client).

For example, for an SFTP server on Linux, you must extract the public key from the key store and enter the key in the **`$HOME/.ssh/authorized_keys`** file on the SFTP server. Ensure that the path and file exist.

3. Create a **`known_hosts`** file. For more information, see "Creating the Known Hosts File" on page 2-2.

Configuring Proxy Services

When you create a proxy service in the **Transport Configuration** page of the ALSB console, you must select the transport protocol as **`sftp`** and specify the endpoint configuration in the following format:

`sftp://`*`hostname`***`:`***`port/directory`*

- *`hostname`* is the host name or IP address of the SFTP server.
- *`port`* is the port on which SFTP server is listening. The default port for SFTP is 22.
- *`directory`* is the location that is polled for files at regular intervals. This path to this directory is relative to the home directory of the user.

Note: Since the SFTP transport supports only message and XML service types, you must select **Messaging Service** or **Any XML Service** as the service type in the **General Configuration** page of the ALSB console.

When you select **Messaging Service** as the service type,

- You must specify **Binary**, **Text**, **MFL**, or **XML** as the request message type.
- You must set the response message type to **None** because the SFTP transport supports only one-way messaging.

For more information, see Proxy Services: Creating and Managing in *Using the AquaLogic Service Bus Console*.

Configure the proxy service as described in the following table.

Table 2-1 Configuring SFTP Proxy Service

Field	Description
User Authentication	The proxy service is authenticated by the SFTP server based on the specified user authentication method. Select the required authentication method. • **Username Password Authentication**: Specifies that a static service account is associated with this authentication method and the client is authenticated using the credentials provided in the service account. • **Host-Based Authentication**: Specifies that a user name and service key provider are required. Any user connecting from a known host is authenticated using the private key of the host. • **Public Key Authentication**: Specifies that a user name and service key provider are required. Users have their own private keys.
Service Account	Enter the service account for the user, or click **Browse** and select a service account. For information about using service accounts, see Service Accounts in *Using the AquaLogic Service Bus Console*.
Service Key Provider	This field is available only for the public key and host-based authentication methods. Enter a service key provider, or click **Browse** and select a service key provider. For more information, see Service Key Providers in *Using the AquaLogic Service Bus Console.*
Username	This value is required only when you select either the host-based or public key authentication method. • In host-based authentication, the user name is required for polling the home directory of the user on the SFTP server. • In public key authentication, the user name is required for polling the home directory of the user and for identifying the location of the public key on the SFTP server. Enter the user name.
Pass By Reference	Select this option to stage the file in the archive directory and pass it as a reference in the headers. **Note:** This option is available only when remote streaming is disabled.

Table 2-1 Configuring SFTP Proxy Service (Continued)

Field	Description
Remote Streaming	Select this option to stream the SFTP files directly from the remote server at the time of processing. When you select this option, the archive directory is the remote directory on the SFTP server machine. Therefore, you must specify the archive directory relative to the SFTP user directory.
File Mask	You can use the file mask for transferring files of specific types. Enter a regular expression to select the files that you want to pick from the directory. The default value is `*.*`.
Polling Interval	Polling interval is the frequency at which the input directory is polled. Polling involves creation of an SFTP connection. Enter the interval (in seconds) at which the file must be polled from the specified location. The default value is **`60`**. **Note:** Avoid setting a low polling interval because a low interval causes frequent polls on the directory.
Read Limit	If numerous files exist in the poll directory, you can limit the number of concurrent transfers by selecting an appropriate value in this field. If you do not want to specify a limit, enter **`0`** (zero). The default value is **`10`**. **Note:** In some cases, the SFTP server may limit the number of concurrent connections; make sure that the read limit you define is lower than the server-defined limit.
Post Read Action	Select the action that must be performed on the message after the file is transferred. • **Archive**: The message is archived in the specified archived directory. • **Delete**: The message is deleted.

Table 2-1 Configuring SFTP Proxy Service (Continued)

Field	Description
Archive Directory	If **Post Read Action** is set to **Archive**, then, after the files are transferred, they are moved (from either the download directory or the remote location) to the archive directory. If you selected the **Pass By Reference** option, you must specify the archive directory. If remote streaming is enabled, the archive directory is with respect to the SFTP server. If remote streaming is disabled, the archive directory is available on the ALSB machine. Specify the absolute path of the archive directory. **Note:** If the directory does not exist, it is created automatically. If you specify a relative path, the directory is created at a path that is relative to the Java process that starts the WebLogic Server.
Download Directory	During file transfer, the files are downloaded to this directory. If remote streaming is enabled, this option is disabled. Specify the absolute path of the directory on your local machine to which files are downloaded during the file transfer. **Note:** If the directory does not exist, it is created automatically. If you specify a relative path, the directory is created at a path that is relative to the Java process that starts the WebLogic Server.
Error Directory	If a problem occurs during file transfer, the messages are posted to the error directory. If remote streaming is enabled, the error directory is with respect to the SFTP server. If remote streaming is disabled, the error directory is available on the ALSB machine. Specify the absolute path of the error directory. **Note:** If the directory does not exist, it is created automatically. If you specify a relative path, the directory is created at a path that is relative to the Java process that starts the WebLogic Server.
Request encoding	Accept the default value (`UTF-8`) as the character set encoding for requests in the SFTP transports.

Table 2-1 Configuring SFTP Proxy Service (Continued)

Field	Description
Advanced Settings	
Scan Subdirectories	Select this option if you want all subdirectories within the directory that is specified in the endpoint URI to be scanned recursively. **Note:** Scanning subdirectories requires additional processing overhead; so use this option judiciously.
Sort By Arrival	Select this option to deliver events in the order of arrival. This ensures that message delivery is not random, but based on the time at which the file is downloaded to the destination directory.
Timeout	Enter the socket timeout interval, in seconds, after which the connection must be dropped. If you do not want the connection to time out, enter **`0`**. The default value is 60.
Retry Count	You can use this setting to enable multiple attempts in case of errors such as network failure. Specify the number of retries for SFTP connection failures. The default value is 3.

For more information about configuring proxy services to use the SFTP transport, see Proxy Services: Creating and Managing in *Using the AquaLogic Service Bus Console*.

Configuring Transport Headers and Metadata

When you configure a proxy service, you can use a Transport Header action to set the header values in messages. The following table lists the transport header and metadata related to the SFTP transport.

Table 2-2 Transport Headers and Metadata

Header / Metadata	Description
`FileName`	This value is used as the file name in the destination directory.
`isFilePath`	This is a metadata field. The possible values are true and false. • True: **`FileName`** is interpreted as the absolute path of the file. • False: **`FileName`** is interpreted as the name of the file.
`filePath`	This is a response metadata field that indicates the absolute path at which the file specified in the **`FileName`** header is written.

Configuring Transport Headers in the ALSB Message Flow

You can configure the transport headers only for outbound requests in the ALSB message flow. In the pipeline, use a transport header action to set the header values in messages. For more information, see Proxy Services: Actions in *Using the AquaLogic Service Bus Console*.

Configuring Transports Headers and Metadata in the Test Console

You can configure the **`FileName`** transport header and the **`isFilePath`** metadata values in the ALSB test console when you test the SFTP transport-based services during development. For more information, see Test Console in *Using the AquaLogic Service Bus Console* and Using the Test Console in *AquaLogic Service Bus User Guide*.

Configuring Business Services

When you create a business service in the **Transport Configuration** page of the ALSB console, you must select the transport protocol as **`sftp`** and specify the endpoint URI (location of the service) in the following format:

`sftp:``//`*`hostname`*`:`*`port/directory`*

- *`hostname`* is the host name or IP address of the SFTP server.

- *port* is the port on which SFTP server is listening. The default port for SFTP is 22.
- *directory* is the location in which the outbound message is stored or written. This path is relative to the home directory of the user.

Note: Since the SFTP transport supports only message and XML service types, you must select **Messaging Service** or **Any XML Service** as the service type in the **General Configuration** page of the ALSB console.

When you select **Messaging Service** as the service type,

- You must specify **Binary**, **Text**, **MFL**, or **XML** as the request message type.
- You must set the response message type to **None** because the SFTP transport supports only one-way messaging.

For more information, see Business Services: Creating and Managing in *Using the AquaLogic Service Bus Console*.

Configure the business service as described in the following table.

Table 2-3 Configuring SFTP Business Service

Field	Description
User Authentication	The proxy service is authenticated by the SFTP server based on the specified user authentication method. Select the required authentication method. • **Username Password Authentication**: Specifies that a static service account is associated with this authentication method and the client is authenticated using the credentials provided in the service account. • **Host-Based Authentication**: Specifies that a user name and service key provider are required. Any user connecting from a known host is authenticated using the private key of the host. • **Public Key Authentication**: Specifies that a user name and service key provider are required. Users have their own private keys.
Service Account	Enter the service account for the user, or click **Browse** and select a service account. For information about using service accounts, see Service Accounts in *Using the AquaLogic Service Bus Console*.

Table 2-3 Configuring SFTP Business Service (Continued)

Field	Description
Service Key Provider	This field is available only for the public key and host-based authentication methods. Enter a service key provider, or click **Browse** and select a service key provider. For more information, see Service Key Providers in *Using the AquaLogic Service Bus Console.*
Username	This value is required only when you select either the host-based or public key authentication method. • In host-based authentication, the user name is required for polling the home directory of the user on the SFTP server. • In public key authentication, the user name is required for polling the home directory of the user and for identifying the location of the public key on the SFTP server. Enter the user name.
Timeout	Enter the socket timeout interval, in seconds, after which the connection must be dropped. If you do not want the connection to time out, enter **`0`**. The default value is 60.
Prefix for destination File Name	Enter the prefix for the name of the file that is stored on the remote server.
Suffix for destination File Name	Enter the suffix for the name of the file that is stored on the remote server.
Request encoding	Accept the default value (**`UTF-8`**) as the character set encoding for requests in the SFTP transports.

For more information about configuring business services using the SFTP transport, see Business Services: Creating and Managing in *Using the AquaLogic Service Bus Console*.

Handling Communication Errors

You can configure the SFTP transport-based business services to handle communications errors, which can occur when a connection or user authentication fails while connecting to the remote SFTP server. When you configure the business service, you can enable the business service endpoint URIs to be taken offline after a specified retry interval.

For more information, see the following topics in Monitoring in *Using the AquaLogic Service Bus Console*:

- "Configuring Operational Settings for Business Services"
- "Viewing Business Services Endpoint URIs Metrics"

Troubleshooting

Most of the errors occur while configuring an SFTP proxy or business service. The following are a few tips to help you understand and solve the errors:

- Make sure that you have an appropriately configured **`known_hosts`** file in place.
- For public key authentication, you must store the public key file on the server. For more information, see the documentation accompanying your SFTP server.
- The **`Connection refused`** error message indicates that the SFTP server is not available on the configured host and port.
- The **`Authentication failed`** error message indicates that the username or password is not valid, or that the public key is not configured correctly.
- The **`Connection did not complete`** error message is displayed after the actual error that caused the connection failure (example: **`Key not found`**) is displayed.
- The **`Key not found for IP, host`** error message indicates that the **`known_hosts`** file does not contain an entry that corresponds to the specified IP-host combination. If the entry exists, then try with another algorithm key; for example, if the earlier attempt was with an RSA key, try again with a DSA key.

Importing Resources

When you import a resource that exists in an ALSB domain, you can preserve the existing security and policy configuration details of the resource by selecting the **Preserve Security and Policy Configuration** option.

The following SFTP service-specific details are preserved when you import a resource:

- Client authentication method
- Reference to the service account (for services associated with username-password authentication)

- Reference to the service key provider (for services associated with host-based or public key authentication)
- User name (for services associated with host-based or public key authentication)

For more information about importing resources from the AquaLogic Service Bus Console, see Importing Resources in *Using the AquaLogic Service Bus Console*.

Importing and Publishing Services: UDDI Registries

When an SFTP service is published to the UDDI registry, the following properties are published:

- Authentication mode
- Request encoding
- Sort by arrival
- File mask

Table 2-4 lists the properties that are imported from the registry when an SFTP service is imported from the UDDI registry.

Table 2-4 Properties Imported from UDDI Registry

Property	Description
Prefix for destination file name	The prefix for the name of the file that is stored on the remote server. The default value is “ ” (null).
Suffix for destination file name	The suffix for the name of the file that is stored on the remote server. The default value is “ ” (null).
Authentication mode	The authentication method that is imported from the registry. When an SFTP business service with user authentication is imported from an UDDI registry to ALSB, a conflict is generated. • For username-password authentication, you must create a service account and associate it with the service. • For host-based or public key authentication, you must create a service key provider and associate it with the service.

After the service is import imported, the default value of the load balancing algorithm is **`round-robin`**.

For more information, see UDDI in *Using the AquaLogic Service Bus Console.*