

AMPS Java Development Guide



AMPS Java Development Guide

5.2

Publication date Jun 26, 2017

Copyright © 2017

All rights reserved. 60East, AMPS, and Advanced Message Processing System are trademarks of 60East Technologies, Inc. All other trademarks are the property of their respective owners.

Table of Contents

1. Introduction	1
1.1. Prerequisites	1
1.2. Upgrading from a Previous Version of the AMPS Java Client	1
2. Installing the AMPS Client	2
2.1. Prerequisites	2
2.2. Obtaining the Client	2
2.3. Test Connectivity to AMPS	2
3. Your First AMPS Program	3
3.1. About the Client Library	3
3.2. Connecting to AMPS	3
3.3. Connection Strings	6
3.4. Connection Parameters	7
3.5. Next Steps	8
4. Subscriptions	9
4.1. Subscribing	9
4.2. Asynchronous Message Processing Interface	11
4.3. Unsubscribing	12
4.4. Understanding Messages	13
4.5. Understanding Threading and Message Handlers	13
4.6. Advanced Messaging Support	14
4.7. Next Steps	15
5. Error Handling	16
5.1. Exceptions	16
5.2. Disconnect Handling	18
5.3. Unexpected Messages	21
5.4. Unhandled Exceptions	21
5.5. Detecting Write Failures	22
6. State of the World	24
6.1. Performing SOW Queries	24
6.2. SOW and Subscribe	25
6.3. Setting Batch Size	27
6.4. Client-Side Conflation	28
6.5. Managing SOW Contents	29
7. Using Queues	30
7.1. Backlog and Smart Pipelining	30
8. Delta Publish and Subscribe	33
8.1. Introduction	33
8.2. Delta Subscribe	33
8.3. Delta Publish	33
9. High Availability	35
9.1. Reconnection with HAClient	35
9.2. Choosing Store Durability	36
9.3. Connections and the ServerChooser	37
9.4. Heartbeats and Failure Detection	40
9.5. Considerations for Publishers	40
9.6. Considerations for Subscribers	42
9.7. Conclusion	45
10. AMPS Programming: Working with Commands	46
10.1. Understanding AMPS Messages	46
10.2. Creating and Populating the Command	46
10.3. Using execute	47

10.4. Command Cookbook	48
11. Utilities	69
11.1. Composite Message Types	69
11.2. NVFIX Messages	70
11.3. FIX Messages	72
12. Advanced Topics	75
12.1. C# Client Compatibility	75
12.2. Rebuilding the Client	75
12.3. Transport Filtering	75
12.4. Working with Messages & Byte Buffers	76
12.5. Providing SSL Certificates to the AMPS Java Client	77
13. Performance Tips and Best Practices	81
13.1. Measure Performance and Set Goals	81
13.2. Simplify Message Format and Contents	82
13.3. Use Content Filtering Where Possible	82
13.4. Use Asynchronous Message Processing	82
13.5. Use Hash Indexes Where Possible	83
13.6. Use a Failed Write Handler and Exception Listener	83
13.7. Reduce Bandwidth Requirements	83
13.8. Limit Unnecessary Copies	84
13.9. Manage Publish Stores	85
13.10. Work with 60East as Necessary	85
A. Exceptions	86
Index	88

Chapter 1. Introduction

This document explains how to use the Java client for AMPS. Use this document to learn how to install, configure, develop applications using AMPS.

1.1. Prerequisites

Before reading this book, it is important to have a good understanding of the following topics:

- **Developing in Java.** To be successful using this guide, you will need to possess a working knowledge of the Java language. Visit <http://java.oracle.com> for resources on learning Java.
- **AMPS concepts.** Before reading this guide, you will need to understand the basic concepts of AMPS, such as *topics*, *subscriptions*, *messages*, and *SOW*. Consult the *AMPS Users' Guide* to learn more about these topics before proceeding.

You will need an installed and running AMPS server to use the product as well. You can write and compile programs that use AMPS without a running server, but you will get the most out of this guide by running the programs against a working server.

1.2. Upgrading from a Previous Version of the AMPS Java Client

The AMPS Java client is designed to maintain binary compatibility for hotfix versions, unless otherwise specified in the release notes.

For versions that introduce or change features, as indicated by a change to the major or minor version number, 60East cannot guarantee binary compatibility. An upgrade that changes either a major or minor version number may require a recompile from source. Changes that do not affect source compatibility are not generally noted in upgrade documents when the major or minor version numbers change.

For example, an upgrade from X.Y.1.0 to X.Y.1.1 of the Java client would be a drop-in replacement unless otherwise noted, as would an upgrade from X.Y.1.1 to X.Y.2.7. However, an upgrade from X.Y.2.7 to X.Z.0.0 may require that you recompile your application from source.

Chapter 2. Installing the AMPS Client

2.1. Prerequisites

Before proceeding with this guide, make sure that the following programs are installed and functioning properly on your development machine.

- Java Development Kit version 1.7 or greater.
- Java Runtime Environment version 1.7 or greater.
- *optional* Apache Ant version 1.8 or greater.

2.2. Obtaining the Client

The AMPS Java client source code and pre-compiled JARs are available in the installed AMPS instance. From the directory where AMPS is installed, the Java client can be found in the `api/client/java/` directory.

The pre-compiled JAR files are located in the `api/client/java/dist/lib/` directory, which contains the following JAR files:

- `amps_client.jar` contains the `com.crankuptheamps.client` package, which includes the classes necessary to build an AMPS client. This JAR and its contents will be discussed throughout this Developer Guide.
- `amps_client-sources.jar` contains the source code files used to create the Java implementation of the AMPS client libraries. This file can be included in an IDE to assist in debugging. It can also be used to rebuild the AMPS client libraries if any custom changes are necessary. See Section 12.2 for instructions on how to recompile the AMPS Java client source.
- `amps_client-javadoc.jar`, like the `amps_client-sources.jar`, can be included in an IDE to provide the javadoc annotations for the implemented classes and methods.

2.3. Test Connectivity to AMPS

Before writing programs using AMPS, make sure connectivity to an AMPS server from this computer is working. Launch a terminal window and change the directory to the AMPS directory in your AMPS installation, and use `spark` to test connectivity to your server, for example:

```
./bin/spark ping -type fix -server 192.168.1.2:9004
```

If you receive an error message, verify that your AMPS server is up and running, and work with your systems administrator to determine the cause of the connectivity issues. Without connectivity to AMPS, you will be unable to make the best use of this guide.

Chapter 3. Your First AMPS Program

In this chapter, we will learn more about the structure and features of the AMPS Java client library, and build our first Java program using AMPS.

3.1. About the Client Library

The AMPS client is packaged as a single JAR file, `amps_client.jar`. You can find `amps_client.jar` in the `dist/lib` directory of the AMPS Java client installation. Every Java application you build will need to reference this JAR file, and the JAR file must be deployed along with your application in order for your application to function properly.

3.2. Connecting to AMPS

Let's begin by writing a simple program that connects to an AMPS server and sends a single message to a topic:

```
import com.crankuptheamps.client.Client;
import com.crankuptheamps.client.exception.AMPSEException;

public class TestPublisher
{
    public static void main(String[] args) {
        Client client = new Client("TestPublisher-Client");
        try {
            client.connect("tcp://127.0.0.1:9007/amps/json");
            client.logon();
            client.publish("messages", "{ \"message\" : \"Hello, world!\" }");
        }
        catch (AMPSEException aex) {
            System.err.println("TestListener caught exception.");
        } finally {
            client.close();
        }
    }
}
```

Example 3.1. Connecting to AMPS

In the preceding Example 3.1, we show the entire program; but future examples will isolate one or more specific portions of the code. The next section describes how to build and run the application and explains the code in further detail.

Build and Run with a Java Command Line

To build this program, you can run the following `javac` command, substituting the path to the `amps_client.jar` with the path where you have installed the AMPS Java Client:

```
javac -classpath /opt/AMPS/api/client/java/dist/lib/amps_client.jar \  
./TestPublisher.java
```

This will create the `TestPublisher.class` file. To run the class file and send your first message to AMPS, you can issue the following command:

```
java -cp ./opt/AMPS/api/client/java/dist/lib/amps_client.jar TestPublisher
```

If the message is published successfully, there is no output to the console. We will demonstrate how to create a subscriber to receive messages in Chapter 4.

Build and Run in Eclipse

To include the JAR file in a standalone Eclipse environment:

1. Right-click the project to which you would like to add `amps_client.jar` and select **Properties** to open the **Properties** configuration window.
2. Click the **Java Build Path** item in the left menu, and then select the **Libraries** tab.
3. Click the **Add External JARS** button and choose the location of your `amps_client.jar`.
4. Click **OK**.

Your Eclipse build path should now include the `amps_client.jar`.

To run the project:

1. Select the menu option **Run # Run Configuration**, which opens the **Run Configuration** window.
2. Right-click **Java Application** and select **New** from the menu.
3. Fill in the **Name** field (for this example, you could use “Test Publisher”).
4. Click the **Search...** button for the **Main class** field.
5. Select the `TestPublisher` class.
6. Finally, click the **Run** button to run the application.

If the message is published successfully, there is no output to the console. We will demonstrate how to create a subscriber to receive messages in Chapter 4.

Examining the Code

Let us now revisit the code we listed earlier.


```

❶import com.crankuptheamps.client.Client;
import com.crankuptheamps.client.exception.AMPSEException;

public class TestPublisher
{
    public static void main(String[] args) {
        ❷Client client = new Client("TestPublisher-Client");
        ❸try {
            ❹client.connect("tcp://127.0.0.1:9007/amps/json");
            ❺client.logon();
            ❻client.publish("messages", "{ \"message\" : \"Hello, world!\" }
");
        }
        ❽catch (AMPSEException aex) {
            System.err.println("TestListener caught exception.");
        }
        ❾finally {
            client.close();
        }
    }
}

```

Example 3.2. Connecting to AMPS

- ❶ The import statements add names into reference for convenience in typing later on in the code. These import the names from the AMPS namespaces:
- `com.crankuptheamps.client.Client`—contains the methods for interacting with AMPS
 - `com.crankuptheamps.client.exception.AMPSEException`—the package contains the exception classes thrown by AMPS when errors occur.

When you use AMPS in your programs, you will be using classes from these namespaces.

- ❷ This line creates a new `Client` object. `Client` encapsulates a single connection to an AMPS server. Methods on `Client` allow for connecting, disconnecting, publishing, and subscribing to an AMPS server. The argument to the `Client` constructor, “TestPublisher-Client”, is a name chosen by the client to identify itself to the server. Errors relating to this connection will be logged with reference to this name, and AMPS will use this name to help detect duplicate messages. AMPS enforces uniqueness for client names when a transaction log is configured, and it is good practice to always use unique client names.
- ❸ Here, we open a try block that concludes with `catch (AMPSEException aex)`.
- ❹ This statement declares a connection to AMPS with the provided URI. The URI consists of the transport, the address, and the protocol to use for the AMPS connection. In this case, the transport is `tcp`, the address is `127.0.0.1:9007`, and the protocol is `amps`. This connection will be used for JSON messages. Check with the person who manages the AMPS instance to get the connection string to use for your programs.
- ❺ The AMPS `logon()` command connects to AMPS and creates a named connection. If we had provided logon credentials in the URI, the command would pass those credentials to AMPS. Without credentials, the client logs on to AMPS anonymously. AMPS versions 5.0 and later require a `logon()` command in the default configuration.
- ❻ Here, we publish a single message to AMPS on the `messages` topic, containing the data `{ "message" : "Hello, World!" }`. This data is placed into an XML message and sent to the server. Upon successful completion of this function, the AMPS client has sent the message to the server, and subscribers to the `messages` topic will receive this “Hello world” message.
- ❼ All caught exceptions in AMPS derive from `AMPSEException`.
- ❾ We close out the example with a `finally` block that closes the `Client` connection and releases all accompanying resources, making the connection eligible for garbage collection.

3.3. Connection Strings

The AMPS clients use connection strings to determine the server, port, transport, and protocol to use to connect to AMPS. When the connection point in AMPS accepts multiple message types, the connection string also specifies the precise message type to use for this connection. Connection strings have a number of elements.




Figure 3.1. elements of a connection string

As shown in the figure above, connection strings have the following elements:

- *Transport* defines the network used to send and receive messages from AMPS. In this case, the transport is `tcp`. For connections to transports that use the Secure Sockets Layer (SSL), use `tcps`.
- *Host address* defines the destination on the network where the AMPS instance receives messages. The format of the address is dependent on the transport. For `tcp` and `tcps`, the address consists of a host name and port number. In this case, the host address is `localhost:9007`.
- *Protocol* sets the format in which AMPS receives commands from the client. Most code uses the default `amps` protocol, which sends header information in JSON format. AMPS supports the ability to develop custom protocols as extension modules, and AMPS also supports legacy protocols for backward compatibility.
- *MessageType* specifies the message type that this connection uses. This component of the connection string is required if the protocol accepts multiple message types and the transport is configured to accept multiple message types. If the protocol does not accept multiple message types, this component of the connection string is optional, and defaults to the message type specified in the transport.

Legacy protocols such as `fix`, `nvfix` and `xml` only accept a single message type, and therefore do not require or accept a message type in the connection string.

As an example, a connection string such as

```
tcp://localhost:9007/amps/json
```

would work for programs connecting from the local host to a Transport configured as follows:

```
<AMPSConfig>
...
  <!-- This transport accepts any known
        message type for the instance: the
        client must specify the message type.
  -->
  <Transport>
    <Name>any-tcp</Name>
    <Type>tcp</Type>
    <InetAddr>9007</InetAddr>
    <ReuseAddr>true</ReuseAddr>
    <Protocol>amps</Protocol>
```

```
</Transport>  
...  
</AMPSConfig>
```

See the *AMPS Configuration Guide* for more information on configuring transports.

Providing Credentials in a Connection String

When using the default authenticator, the AMPS clients support the standard format for including a username and password in a URI, as shown below:

```
tcp://user:password@host:port/protocol/message_type
```

When provided in this form, the default authenticator provides the username and password specified in the URI. If you have implemented another authenticator, that authenticator controls how passwords are provided to the AMPS server.

3.4. Connection Parameters

When specifying a URI for connection to an AMPS server, you may specify a number of transport-specific options in the parameters section of the URI. Here is an example:

```
tcp://localhost:9007/amps/json?tcp_nodelay=true&tcp_sndbuf=100000
```

In this example, we have specified the AMPS instance on `localhost`, port `9007`, connecting to a transport that uses the `amps` protocol and sending JSON messages. We have also set two parameters, `tcp_nodelay`, a Boolean (`true/false`) parameter, and `tcp_sndbuf`, an integer parameter. Multiple parameters may be combined to finely tune settings available on the transport. Normally, you'll want to stick with the defaults on your platform, but there may be some cases where experimentation and fine-tuning will yield higher or more efficient performance.

AMPS supports the value of `tcp` in the connection string for TCP/IP connections, and the value of `tcps` in the connection string for SSL encrypted connections.

Transport options

The following transport options are available for TCP connections:

<code>tcp_rcvbuf</code>	(integer) Sets the socket receive buffer size. This defaults to the system default size. (On Linux, you can find the system default size in <code>/proc/sys/net/core/rmem_default</code> .)
<code>tcp_sndbuf</code>	(integer) Sets the socket send buffer size. This defaults to the system default size. (On Linux, you can find the system default size in <code>/proc/sys/net/core/wmem_default</code> .)
<code>tcp_nodelay</code>	(boolean) Enables or disables the <code>TCP_NODELAY</code> setting on the socket. By default <code>TCP_NODELAY</code> is disabled.
<code>tcp_linger</code>	(integer) Enables and sets the <code>SO_LINGER</code> value for the socket. By default, <code>SO_LINGER</code> is enabled with a value of <code>10</code> , which specifies that the socket will linger for 10 seconds.

`tcp_keepalive` (boolean) . Enables or disables the `SO_KEEPALIVE` value for the socket. The default value for this option is true.

3.5. Next Steps

You are now able to develop and deploy an application in Java that publishes messages to AMPS. In the following chapters, you will learn how to subscribe to messages, use content filters, work with SOW caches, and fine-tune messages that you send.

Chapter 4. Subscriptions

Messages published to a topic on an AMPS server are available to other clients via a subscription. Before messages can be received, a client must subscribe to one or more topics on the AMPS server so that the server will begin sending messages to the client. The server will continue sending messages to the client until the client unsubscribes, or the client disconnects. With content filtering, the AMPS server will limit the messages sent to only those messages that match a client-supplied filter. In this chapter, you will learn how to subscribe, unsubscribe, and supply filters for messages using the AMPS Java client.

4.1. Subscribing

You subscribe to an AMPS topic by calling `Client.subscribe()`. Here is a short example showing the simplest way to subscribe to a topic (error handling and connection details are omitted for brevity):

```
class MyApp
{
    public static void main(String[] args)
    {
        ❶ Client client = new Client("subscribe");

        try
        {
            client.connect("tcp://127.0.0.1:9007/amps/json");
            client.logon();

            ❷ try (MessageStream ms = client.subscribe("messages"))
            ❸ for (Message m : ms)
            {
                ❹ System.out.println(m.getData());
            }

        }
        catch(AMPSEException e){;}
        finally{ client.close(); }
    }
}
```

Example 4.1. Subscribing to a Topic

- ❶ We create a `Client`, then `connect()` and `logon()`.
- ❷ Here, we subscribe to the topic messages. We do not provide a filter, so the subscription receives all of the messages published to the topic, regardless of content. We protect the `MessageStream` in a try with resources block, so that the stream is closed when control exits the block.
- ❸ Here, we iterate over the messages returned by the message stream. When we no longer need to subscribe, we can break out of the loop. When the `MessageStream` is cleaned up, the client sends an `unsubscribe` command to AMPS and stops receiving messages.

- ④ Within the loop, we process the message. In this case, we simply print the contents of the message

AMPS creates a background thread that receives messages and copies them into the `MessageStream` that you iterate over. This means that the client application as a whole can continue to receive messages while you are doing processing work.

The simple method described above is provided for convenience. The AMPS Java client provides convenience methods for the most common forms of the commands. AMPS also provides an interface that allows you precise control over the command. Using that interface, the example above becomes:

```
class MyApp
{
    public static void main(String[] args)
    {
        ❶ Client client = new Client("subscribe");

        try
        {
            client.connect("tcp://127.0.0.1:9007/amps");
            client.logon();

            ❷ Command command = new Command("subscribe")
                .setTopic("messages");

            ❸ try (MessageStream ms = client.execute(command))
            for (Message m : ms)
            {
                ❹ System.out.println(m.getData());
            }
        }
        catch (AMPSException e){;}
        finally{ client.close(); }
    }
}
```

Example 4.2. Subscribing to a Topic

- ❶ We create a `Client` that is properly connected to an AMPS server.
- ❷ We create a `Command` object to subscribe to the messages topic.
- ❸ Here we execute the command and subscribe to the topic messages. This works exactly the same way as the command in Example 4.1. If, at any time, we no longer need to subscribe, we can break out of the loop. We use a try with resources to automatically clean up the `MessageStream` when we leave the try block. When the `MessageStream` is cleaned up, the client sends an `unsubscribe` command to AMPS and stops receiving messages.
- ❹ Within the loop, we process the message. In this case, we simply print the contents of the message

The `Command` interface allows you to precisely customize the commands you send to AMPS. For flexibility and ease of maintenance, 60East recommends using the `Command` interface (rather than a named method) for any command that will receive messages from AMPS. For publishing messages, there can be a slight performance advantage to using the named commands where possible.

4.2. Asynchronous Message Processing Interface

The AMPS Java client also supports an interface designed for asynchronous message processing. In this case, you add a message handler to the call to subscribe. The client returns the command ID of the subscribe command once the server has acknowledged that the command has been processed. As messages arrive, AMPS calls your message handler directly on the background thread. This can be an advantage for some applications. For example, if your application is highly multithreaded and copies message data to a work queue processed by multiple threads, there is usually a performance benefit to enqueueing work directly from the background thread. See Section 4.5 for a discussion of threading considerations, including considerations for message handlers.

As with the simple interface, the AMPS client provides both convenience interfaces and interfaces that use a Command object. The following example shows how to use the asynchronous interface.

```
class MyApp {
    public static void main(String[] args)
    {
        ❶ Client client = new Client("subscribe");

        try {
            client.connect("tcp://127.0.0.1:9007/amps/json");
            client.logon();

            ❷ MessagePrinter mp = new MessagePrinter();
            Command command = new Command("subscribe")
                .setTopic("messages");

            ❸ CommandId subscriptionId = client.executeAsync(command, mp);
        }
        catch(AMPSEException e){;}
        finally{ client.close(); }
    }
}
```

Example 4.3. Subscribing to a topic with asynchronous processing

- ❶ We create a `Client` here, then call `connect()` and `logon()` to connect to AMPS.
- ❷ Here, we create a new `MessagePrinter` object. The `MessagePrinter` class implements `MessageHandler`, as described below. The subscription uses this object to handle all messages returned from the subscription.
- ❸ Here, we call the overload of `Client.execute()` that specifies the command and the message handler to invoke with messages received in response to the command.

```
❶ public class MessagePrinter implements MessageHandler {
    public void invoke(Message m) {
        System.out.println(m.getData());
    }
}
```

Example 4.4. Implementing MessageHandler

- ❶ An implementation of `MessageHandler` provides an `invoke()` method that receives a `com.crankuptheamps.client.Message` object. Notice that the same instance of this class is called

for all messages on a given subscription, and that the instance is called asynchronously from the background thread created by the client. Design your message handlers so they have access to any program state that they need to do their work.



In the asynchronous interface, the AMPS client resets and reuses the message provided to message handler functions between calls. This improves performance in the client, but means that if your handler function needs to preserve information contained within the message, you must copy the information rather than just saving the message object. Otherwise, the AMPS client cannot guarantee the state of the object or the contents of the object when your program goes to use it.

4.3. Unsubscribing

With the simple interface, AMPS automatically unsubscribes from the topic when the `MessageStream` is cleaned up by having the `close()` method called. This typically happens as part of the resource management for the `for` loop, although you can also call the method yourself.

In the advanced interface, when the subscription is successfully made, messages will begin flowing to our `MessagePrinter.invoke()` function, and the `Client.subscribe()` call will return a `CommandId` that serves as the identifier for this subscription. A `Client` can have any number of active subscriptions, and this `CommandId` instance is used to refer to the particular subscription we have made here. For example, to unsubscribe, we simply pass in this identifier:

```
Client c = ...;

Command command = new Command("unsubscribe")
    .setTopic("messages");

try (MessageStream ms = client.execute(command))
{
    for (Message msg : ms)
    {
        System.out.println("Response to unsubscribe : " + msg.getAckType());
    }
}
```

Example 4.5. Unsubscribing from a topic

In this example, as in the previous section, we use the `Client.subscribe()` method to create a subscription to the `messages` topic. The `subscribe` method returns an identifier for the subscription created in AMPS. When our application is done listening to this topic, it unsubscribes (on the last line) by passing in the `subscriptionId` returned by `subscribe()`. AMPS deletes the subscription. After the subscription is removed, no more messages will flow into our `MessageHandler.invoke()` function.

AMPS also accepts a comma-delimited list of subscription identifiers to the `unsubscribe` command, or the keyword `all` to unsubscribe all subscriptions for the client.

4.4. Understanding Messages

So far, we have seen that subscribing to a topic involves working with objects of type `com.crankuptheamps.client.Message`. A `Message` represents a single message to or from an AMPS server. Messages are received or sent for every client/server operation in AMPS.

Header properties

There are two parts of each message in AMPS: a set of headers that provide metadata for the message, and the data that the message contains. Every AMPS message has one or more header fields defined. The precise headers present depend on the type and context of the message. There are many possible fields in any given message, but only a few are used for any given message. For each header field, the `Message` class contains a distinct property that allows for retrieval and setting of that field. For example, the `Message.getCommandId()` function corresponds to the `commandId` header field, the `Message.getBatchSize()` function corresponds to the `BatchSize` header field, and so on. For more information on these header fields, consult the *AMPS User Guide* and *AMPS Command Reference*.

To work with header fields, a `Message` contains `getXxx()/setXxx()` methods corresponding to the header fields. 60East does not recommend attempting to parse header fields from the raw data of the message.

In AMPS, fields sometimes need to be set to a unique identifier value. For example, when creating a new subscription, or sending a manually constructed message, you'll need to assign a new unique identifier to multiple fields such as `CommandId` and `SubscriptionId`. For this purpose, `Message` provides `newXxx()` methods for each field that generates a new unique identifier and sets the field to that new value.

getData() method

Access to the data section of a message is provided via the `getData()` method. The data contains the unparsed data in the message. The `getData()` method returns the data as a Java string, which is suitable message formats that can be represented as Unicode text, such as JSON, XML, FIX, or NVFIX. For binary data, the AMPS Java client provides a `getDataRaw()` method to allow you to work with the underlying byte array in the message. See Section 12.4 for details.

The AMPS Java client contains a collection of helper classes for working with message types that are specific to AMPS (for example, FIX, NVFIX, and AMPS composite message types). For message types that are widely used, such as JSON or XML, you can use whichever library you typically use in your environment.

4.5. Understanding Threading and Message Handlers

When you call a subscribe command, the client creates a thread that runs in the background. The command returns, while the thread receives messages. In the simple case, using synchronous message processing, the client provides an internal handler function that populates the `MessageStream`. The `MessageStream` is used on the calling thread, so operations on the `MessageStream` do not block the background thread.

When using asynchronous message processing, AMPS calls the handler function from the background thread. Message handlers provided for asynchronous message processing must be aware of the following considerations.

The client creates one background thread per client object. A message handler that is only provided to a single client will only be called from a single thread. If your message handler will be used by multiple clients, then multiple threads will call your message handler. In this case, you should take care to protect any state that will be shared between threads.

For maximum performance, do as little work in the message handler as possible. For example, if you use the contents of the message to update an external database, a message handler that adds the relevant data to an update queue that is processed by a different thread will typically perform better than a message handler that does this update during the message handler.

While your message handler is running, the thread that calls your message handler is no longer receiving messages. This makes it easier to write a message handler, because you know that no other messages are arriving from the same subscription. However, this also means that you cannot use the same client that called the message handler to send commands to AMPS. Acknowledgements from AMPS cannot be processed, and your application will deadlock waiting for the acknowledgement. Instead, enqueue the command in a work queue to be processed by a separate thread, or use a different client object to submit the commands.

The AMPS client resets and reuses the `Message` provided to this function between calls. This improves performance in the client, but means that if your handler function needs to preserve information contained within the message, you must copy the information rather than just saving the message object. Otherwise, the AMPS client cannot guarantee the state of the object or the contents of the object when your program goes to use it.

4.6. Advanced Messaging Support

Method `Client.subscribe()` provides options for subscribing to topics even when you do not know their exact names. It also provides a filter that works on the server to limit the messages your application must process.

Regex Topics

Regular Expression (Regex) Topics allow a regular expression to be supplied in the place of a topic name. When you supply a regular expression, it is as if a subscription is made to every topic that matches your expression, including topics that do not yet exist at the time of creating the subscription.

To use a regular expression, simply supply the regular expression in place of the topic name in the `subscribe()` call. For example:

```
class MyApp{
...
    CommandId id = client.subscribe(
        mp, // MessageHandler implementation
        "client.*", // Topic with a Regex
        5000); // Timeout
... }
```

Example 4.6. Regex Topic Subscription

In this example, messages on topics “client” and “client1” would match the regular expression, and those messages would all be sent to our `mp` instance, which as before is an implementation of the `MessageHandler` interface. As in the example, you can use the `getTopic()` method to determine the actual topic of the message sent to the `MessageHandler` function.

Content Filtering

One of the most powerful features of AMPS is content filtering. With content filtering, filters based on message content are applied at the server, so that your application and the network are not utilized by messages that are uninteresting for your application. For example, if your application is only displaying messages from a particular user, you can send a content filter to the server so that only messages from that particular user are sent to the client.

To apply a content filter to a subscription, simply pass it into the `Client.subscribe()` call:

```
CommandId subscriptionId = client.subscribe(
    mp, // MessageHandler implementation
    "messages", // Topic
    "/sender = 'mom'", // Content filter
    5000); // Timeout
```

Example 4.7. Using Content Filters

In this example, we have passed in a content filter `"/sender = 'mom' "`. This will cause the server to only send us messages from the `messages` topic that additionally have a `sender` field equal to `mom`.

For example, the AMPS server will send the following message, where `/sender` is `mom`:

```
{ "sender" : "mom",
  "text" : "Happy Birthday!",
  "reminder" : "Call me Thursday!" }
```

The AMPS server will not send a message with a different `/sender` value:

```
{ "sender" : "henry dave",
  "text" : "Things do not change; we change." }
```

Updating the Filter on a Subscription

AMPS allows you to update the filter on a subscription. When you replace a filter on the the subscription, AMPS immediately begins sending only messages that match the updated filter. Notice that if the subscription was entered with a command that includes a SOW query, using the `replace` option can re-issue the SOW query (as described in the *AMPS User Guide*).

To update a the filter on a subscription, you create a `subscribe` command. You set the `SubscriptionId` provided on the `Command` to the identifier of the existing subscription, and include the `replace` option on the `Command`. When you send the `Command`, AMPS atomically replaces the filter and sends messages that match the updated filter from that point forward.

4.7. Next Steps

At this point, you are able to build AMPS programs in Java that publish and subscribe to AMPS topics. For an AMPS application to be truly robust, it needs to be able to handle the errors and disconnections that occur in any distributed system. In the next chapter, we will take a closer look at error handling and recovery, and how you can use it to make your application ready for the real world.

Chapter 5. Error Handling

In every distributed system, the robustness of your application depends on its ability to recover gracefully from unexpected events. The AMPS client provides the building blocks necessary to ensure your application can recover from the kinds of errors and special events that may occur when using AMPS.

5.1. Exceptions

Generally speaking, when an error occurs that prohibits an operation from succeeding, AMPS will throw an exception. AMPS exceptions universally derive from `com.crankuptheamps.client.exception.AMPSEException`, so by catching `AMPSEException`, you will be sure to catch anything AMPS throws. For example:

```
import com.crankuptheamps.client.Client;
import com.crankuptheamps.client.exception.AMPSEException;
...

public void readAndEvaluate(Client client) {
    Scanner scanner = new Scanner(System.in);
    String payload = scanner.nextLine();

    // write a new message to AMPS
    if( payload != null) {
        try{
            client.publish("UserMessage", " { \"data\" : \"" + payload +
"\\" }");
        }
        catch (AMPSEException e) {
            System.err.println("An AMPS exception occurred: " +
                e.toString() );
            e.printStackTrace();
        }
    }
}
```

Example 5.1. Catching an AMPS Exception

In this example, if an error occurs the program writes the error to `stderr`, and the `publish()` command fails. However, `client` is still usable for continued publishing and subscribing. When the error occurs, the exception is written `stderr`, which calls the exception's `toString()` method. As with most Java exceptions, `toString()` will convert the `Exception` into a string that includes a message. The `printStackTrace()` method will write to `stderr` the stack trace and information on any "inner" exceptions (exceptions from outside of AMPS that caused AMPS to throw an exception).

AMPS exception types vary based on the nature of the error that occurs. In your program, if you would like to handle certain kinds of errors differently than others, you can catch the appropriate subclass of `AMPSEException` to detect those specific errors and do something different.

```
public CommandId CreateNewSubscription(Client client)
{
```

```

CommandId id = null;
string topicName;
while (id == null) {
❶ topicName = askUserForTopicName();
    try {
        Command command = new Command("subscribe")
            .setTopic(topicName);

        MessagePrinter mp = new MessagePrinter();

        id = client.executeAsync(command, mp);
    }
❷ catch( BadRegexTopicException ex )
    {
        ❸ System.err.println("Error: bad topic name " +
            "or regular expression " + topicName +
            ". The error was: " + ex.toString());
        // we'll ask the user for another topic
    }
❹ catch(AMPSException ex)
    {
        System.err.println("Error: error setting " +
            "up subscription to topic " +
            topicName + ". The error was: " +
            ex.toString());
        ❺ return null; // give up
    }
}
return id;
}

```

Example 5.2. Catching AMPSException Subclasses

- ❶ In Example 5.2 our program is an interactive program that attempts to retrieve a topic name (or regular expression) from the user.
- ❸ This line indicates that the program catches the `BadRegexTopicException` exception and displays a specific error to the user indicating the topic name or expression was invalid. By not returning from the function in this catch block, the while loop runs again and the user is asked for another topic name.
- ❹ If an AMPS exception of a type other than `BadRegexTopicException` is thrown by AMPS, it is caught here. In that case, the program emits a different error message to the user.
- ❺ At this point the code stops attempting to subscribe to the client by the `return null` statement.

Exception Types

Each method in AMPS documents the kinds of exceptions that it can throw. For reference, Table A.1 contains a list of all of the exception types you may encounter while using AMPS, when they occur, and what they mean.

Exception Handling and Asynchronous Message Processing

When using asynchronous message processing, exceptions thrown from the message handler are silently absorbed by the AMPS Java client by default. The AMPS Java client allows you to register an exception listener to detect and

respond to these exceptions. When an exception listener is registered, AMPS will call the exception listener with the exception. See Example 5.6 for details.

5.2. Disconnect Handling

Every distributed system will experience occasional disconnections between one or more nodes. The reliability of the overall system depends on an application's ability to efficiently detect and recover from these disconnections. Using the AMPS Java client's disconnect handling, you can build powerful applications that are resilient in the face of connection failures and spurious disconnects.

The `HAClient` class, included with the AMPS Java client, contains a disconnect handler and other features for building highly-available applications. 60East recommends using the `HAClient` for automatic reconnection whenever possible, as the `HAClient` disconnect handler has been carefully crafted to handle a wide variety of edge cases and potential failures. This section covers the use of a custom disconnect handler in the event that the behavior of the `HAClient` does not suit the needs of your application.

AMPS disconnect handling gives you the ultimate in control and flexibility regarding how to respond to disconnects. Your application gets to specify exactly what happens when an exception occurs by supplying a function to `Client.setDisconnectHandler()`, which is invoked whenever a disconnect occurs.

Example 5.3 shows the basics:

```
public class TestDisconnect {
    private String _uri = null;

    public TestDisconnect(String uri)
        throws AMPSException
    {
        this._uri = uri;
        Client client = new Client("TestDiscon-Client");
        TestDisconnectHandler dh =
            new TestDisconnectHandler();
        ❶client.setDisconnectHandler(dh);
        client.connect(this._uri);

        for (Message msg : client.subscribe("order") )
        {
            // ...
        }
    }

    ❷class TestDisconnectHandler
        implements ClientDisconnectHandler
    {
        public void invoke(Client client) {
            try{
                client.connect(_uri);
                client.logon();
                Thread.sleep(5000);
            } catch(Exception e){;}
        }
    }
}
```

```

class MsgPrinter implements MessageHandler
{
    public void invoke(Message m) {
        System.out.println(m.getData());
    }
}

```

Example 5.3. Supplying a Disconnect Handler

- ❶ In Example 5.3 the `setDisconnectHandler()` method is called to supply a function for use when AMPS detects a disconnect. At any time, this function may be called by AMPS to indicate that the client has disconnected from the server, and to allow your application to choose what to do about it. The application continues on to connect and subscribe to the `orders` topic.
- ❷ Our disconnect handler's implementation begins here. In this example, we simply try to reconnect to the original server after a 5000 millisecond pause. Errors are likely to occur here—therefore we must have disconnected for a reason—but `Client` takes care of catching errors from our disconnect handler. If an error occurs in our attempt to reconnect and an exception is thrown by `connect()`, `Client` will catch it and absorb it, passing it to the `ExceptionListener` if registered. If the client is not connected by the time the disconnect handler returns, AMPS throws `DisconnectedException`.

By creating a more advanced disconnect handler, you can implement logic to make your application even more robust. For example, imagine you have a group of AMPS servers configured for high availability—you could implement fail-over by simply trying the next server in the list until one is found. Example 5.4 shows a brief example.

```

public class TestDisconnect {
    private String[] _uris;
    private int _currentUri = 0;
    ...

    public TestDisconnect(String[] uris) throws
        AMPSException{
        ❶ _uris = uris;
        Client client = new Client("TestDiscon-Client");
        TestDisconnectHandler dh =
            new TestDisconnectHandler();
        client.setDisconnectHandler(dh);
        ❷ ConnectToNextUri(client);
    }

    private void ConnectToNextUri(Client client) {
        ❸ while(true){
            try {
                MsgPrinter mp = new MsgPrinter();
                client.connect(_uris[_currentUri]);
                client.logon();
                Command command = new
                Command("subscribe").setTopic("orders");
                ❹ client.execute(command, mp);
                return;
            }
            catch (Exception e) {
                _currentUri = (_currentUri + 1) %

```

```

        _uris.length;
        System.err.println("Connection failed " +
            e.toString() + ". Failing over to " +
            _uris[_currentUri]);
    }
}

class TestDisconnectHandler
    implements ClientDisconnectHandler
{
    public void invoke(Client client) {
        try{
            ConnectToNextUri(client);
        }catch(Exception e){;}
    }
}

class MsgPrinter
    implements MessageHandler
{
    public void invoke(Message m) {
        System.out.println(m.getData());
    }
}
}

```

Example 5.4. Simple Client Failover Implementation

- ❶ Here our application is configured with an array of AMPS server URIs to choose from, instead of a single URI. These will be used in the `ConnectToNextUri()` method as explained below.
- ❷ `ConnectToNextUri()` is invoked by our disconnect handler `TestDisconnectHandler`. Since our client is currently disconnected, we manually invoke our disconnect handler to initiate the first connection.
- ❸ In our disconnect handler, we invoke `ConnectToNextUri()`, which loops around the array of URIs, attempting to connect to each one. In the `invoke()` method, it attempts to connect to the current URI and, if it is successful, returns immediately. If the connection attempt fails, the exception handler for `AMPSException` is invoked. In the exception handler, we advance to the next URI, display a warning message, and continue around the loop. This simplistic handler never gives up, but in a typical implementation, you would likely at some point stop attempting to reconnect.
- ❹ At this point the registers a subscription to the server we have connected to. It is important to note that, once a new server is connected, it the responsibility of the application to re-establish any subscriptions placed previously. This behavior provides an important benefit to your application: one reason for disconnect is due to a client's inability to keep up with the rate of message flow. In a more advanced disconnect handler, you could choose to not re-establish subscriptions that are the cause of your application's demise.

Using a Heartbeat to Detect Disconnection

The AMPS client includes a heartbeat feature to help applications detect disconnection from the server within a predictable amount of time. Without using a heartbeat, an application must rely on the operating system to notify the application when a disconnect occurs. For applications that are simply receiving messages, it can be impossible to tell whether a socket is disconnected or whether there are simply no incoming messages for the client.

When you set a heartbeat, the AMPS client sends a heartbeat message to the AMPS server at a regular interval, and waits a specified amount of time for the response. If the operating system reports an error on send, or if the server does not respond within the specified amount of time, the AMPS client considers the server to be disconnected.

The AMPS client processes heartbeat messages on the client receive thread, which is the thread used for asynchronous message processing. If your application uses asynchronous message processing and occupies the thread for longer than the heartbeat interval, the client may fail to respond to heartbeat messages in a timely manner and may be disconnected by the server.

5.3. Unexpected Messages

The AMPS Java client handles most incoming messages and takes appropriate action. Some messages are unexpected or occur only in very rare circumstances. The AMPS Java client provides a way for clients to process these messages. Rather than providing handlers for all of these unusual events, AMPS provides a single handler function for messages that can't be handled during normal processing.

Your application registers this handler by setting the `lastChanceMessageHandler` for the client. This handler is called when the client receives a message that can't be processed by any other handler. This is a rare event, and typically indicates an unexpected condition.

For example, if a client publishes a message that AMPS cannot parse, AMPS returns a failure acknowledgement. This is an unexpected event, so AMPS does not include an explicit handler for this event, and failure acknowledgements are received in the method registered as the `lastChanceMessageHandler`.

Your application is responsible for taking any corrective action needed. For example, if a message publication fails, your application can decide to republish the message, publish a compensating message, log the error, stop publication altogether, or any other action that is appropriate.

5.4. Unhandled Exceptions

In the AMPS Java client, exceptions can occur that are not thrown to the user. For example, when an exception occurs in the process of reading subscription data from the AMPS server, the exception occurs on a thread inside of AMPS. Consider the following example:

```
public class MyApp
{
    ...
    public static void waitToBePoked(Client client)
    {
        Command command = new Command("subscribe")
            .setTopic("pokes")
            .setFilter("/Pokee LIKE'" +
                System.getProperty("user.name") + "-.*")
            .setTimeout(5000);
        client.execute(command, new MsgPrinter());

        Console c = System.console();
        Reader r = c.reader();
        while(r.read() == null) {
            Thread.sleep(10);
        }
    }
}
```

```

    }
}

class MsgPrinter
    implements MessageHandler
{
    public void invoke(Message m) {
        System.out.println(m.getData());
    }
}
}

```

Example 5.5. Where Do Exceptions Go?

In this example, we set up a simple subscription to wait for messages on the “pokes” topic, whose “Pokee” tag begins with our user name. When messages arrive, we print a message out to the console, but otherwise our application waits for a key to be pressed.

Inside of the AMPS client, the client creates a new thread of execution that reads data from the server, and invokes message handlers and disconnect handlers when those events occur. When exceptions occur inside this thread, however, there is no caller for them to be thrown to, and by default they are ignored.

In applications where it is important to deal with every issue that occurs in using AMPS, you can set an `ExceptionHandler` via `Client.setExceptionHandler()` that receives these otherwise-unhandled exceptions. Making the modifications shown in Example 5.6 to our previous example will allow those exceptions to be caught and handled. In this case we are simply printing those caught exceptions out to the console.

```

public class MyApp{
    ...
    client.setExceptionHandler( new
        CustomExceptionHandler() );
    ...
}

class CustomExceptionHandler
    implements ExceptionListener
{
    public void exceptionThrown(Exception ex) {
        System.out.println(ex.toString());
    }
}
}

```

Example 5.6. Exception Listener

In this example we have added a call to `setExceptionHandler()`, registering a simple function that writes the text of the exception out to the console. Even though our application waits for a user to press a key, messages to the console will still be produced, both as incoming “poke” topics arrive, and as issues arise inside of AMPS.

5.5. Detecting Write Failures

The `publish` methods in the Java client deliver the message to be published to AMPS and then return immediately, without waiting for AMPS to return an acknowledgement. Likewise, the `sowDelete` methods request deletion of SOW messages, and return before AMPS processes the message and performs the deletion. This approach provides

high performance for operations that are unlikely to fail in production. However, this means that the methods return before AMPS has processed the command, without the ability to return an error in the event that the command fails.

The AMPS Java client provides a `FailedWriteHandler` that is called when the client receives an acknowledgement that indicates a failure to persist data within AMPS. To use this functionality, you implement the `FailedWriteHandler` interface, construct an instance of your new class, and register that instance with the `setFailedWriteHandler()` function on the client. When an acknowledgement returns that indicates a failed write, AMPS calls the registered handler method with information from the acknowledgement message, supplemented with information from the client publish store if one is available. Your client can log this information, present an error to the user, or take whatever action is appropriate for the failure.

When no `FailedWriteHandler` is registered, acknowledgements that indicate errors in persisting data are treated as unexpected messages and routed to the `LastChanceMessageHandler`. In this case, AMPS provides only the acknowledgement message and does not provide the additional information from the client publish store.

Chapter 6. State of the World

AMPS State of the World (SOW) allows you to automatically keep and query the latest information about a topic on the AMPS server, without building a separate database. Using SOW lets you build impressively high-performance applications that provide rich experiences to users. The AMPS Java client lets you query SOW topics and subscribe to changes with ease.

6.1. Performing SOW Queries

To begin, we will look at a simple example of issuing a SOW query.

```
...
public void executeSOWQuery(Client client) {
    for (Message m : client.sow("messages-sow",
                               "/symbol = 'ROL'"))
    {
        if (m.getCommand() == Message.Command.GroupBegin)
        {
            System.out.println("--- Begin SOW Results ---");
        }
        if (m.getCommand() == Message.Command.GroupEnd)
        {
            System.out.println("--- End SOW Results ---");
        }
        if (m.getCommand() == Message.Command.SOW)
        {
            System.out.println(m.getData());
        }
    }
}
...
```

Example 6.1. Basic SOW Query

In listing Example 6.1 the `executeSOWQuery()` method invokes `Client.sow()` to initiate a SOW query on the orders topic, for all entries that have a symbol of 'ROL'.

As the query executes, the body of the loop processes each matching entry in the topic. Messages containing the data of matching entries have a `Command` of value `sow`; so as those arrive, we write them to the console. AMPS sends a `begin_group` message at the beginning of the results and an `end_group` message at the end of the results. We use those messages to delimit the results of the query.

As with `subscribe`, the `sow` command also provides an asynchronous version, as well as versions that accept a `Command`. For example, the listing below shows an asynchronous SOW command that specifies the *batch size*, or the maximum number of records that AMPS will return at a time.

```
public void executeSOWQuery(Client client) {
    Command command = new Command(Message.Command.SOW)
        .setTopic("orders")
}
```

```

        .setFilter("/symbol = 'ROL'")
        .setBatchSize(100);

    client.executeAsync(command, new MessagePrinter());
}
...
public class MessagePrinter
    implements MessageHandler
{
    public void invoke(Message m) {
        if (m.getCommand() == Message.Command.SOW) {
            System.out.println(m.getData());
        }
    }
}

```

Example 6.2. Asynchronous SOW Query

6.2. SOW and Subscribe

Imagine an application that displays real time information about the position and status of a fleet of delivery vans. When the application starts, it should display the current location of each of the vans along with their current status. As vans move around the city and post other status updates, the application should keep its display up to date. Vans upload information to the system by posting message to a van location topic, configured with a key of `van_id` on the AMPS server.

In this application, it is important to not only stay up-to-date on the latest information about each van, but also to ensure all of the active vans are displayed as soon as the application starts. Combining a SOW with a subscription to the topic is exactly what is needed, and that is accomplished by the AMPS `sow_and_subscribe` command. Now we will look at an example:

```

private void updateVanPosition(Message message)
{
    switch (message.getCommand()) {
        case Message.Command.SOW:
        case Message.Command.Publish:
            ❶ addOrUpdateVan(message);
            break;
        case Message.Command.OOF:
            removeVan(message);
            break;
    }
}

public void subscribeToVanLocation(Client client) {
    try {
        Command command = new Command("sow_and_subscribe")
            .setTopic("van_location")
            .setFilter("/status = 'ACTIVE'")
            .setBatchSize(100)
            ❷ .setOptions("oof");
    }
}

```

```

    ❷ for (Message message : client.execute(command))
    {
        updateVanPosition(message);
    }

    } catch (AMPSEException aex) {
        System.err.println("TestListener caught exception.");
    }
}

public void addOrUpdateVan(message) {

    // use information in the message to add the van or update
    // the van position

    ...

}

public void removeVan(message) {

    // use information in the message to remove information on
    // the van position

    ...

}

```

Example 6.3. Using `sowAndSubscribe`

- ❶ For each of these messages, `addOrUpdateVan()` presumably adds the van to our application's display. As vans send updates to the AMPS server, those are also received by the client because of the subscription placed by `sowAndSubscribe()`. Our application does not need to distinguish between updates and the original set of vans we found via the SOW query, so we use `addOrUpdateVan()` to display the new position of vans as well.
- ❸ In Example 6.3, we issue a `sowAndSubscribe()` to begin receiving information about all of the active delivery vans in the system. All of the vans in the system now are returned as `Messages` whose `getCommand()` returns `SOW`.
- ❷ Notice here that we specified `true` for the `oofEnabled` parameter. Setting this parameter to `true` causes us to receive *Out-of-Focus* ("OOF") messages for topic. OOF messages are sent when an entry that was sent to us in the past no longer matches our query. This happens when an entry is removed from the SOW cache via a `sowDelete()` operation, when the entry expires (as specified by the expiration time on the message or by the configuration of that topic on the AMPS server), or when the entry no longer matches the content filter specified. In our case, if a van's status changes to something other than `ACTIVE`, it no longer matches the content filter, and becomes out of focus. When this occurs, a `Message` is sent with `Command` set to `OOF`. We use OOF messages to remove vans from the display as they become inactive, expire, or are deleted.

Now we will look at an example that uses the asynchronous form of `sow_and_subscribe`:

```

public class VanPositionUpdater {

    public void invoke(Message message) {
        updateVanPosition(message);
    }

}

```

```

private void updateVanPosition(Message message) {
    switch (message.getCommand()) {
        case Message.Command.SOW:
        case Message.Command.Publish:
            addOrUpdateVan(message);
            break;
        case Message.Command.OOF:
            removeVan(message);
            break;
    }
}

public void addOrUpdateVan(message) {
    // use information in the message to add the van or update
    // the van position
    ...
}

public void removeVan(message) {
    // use information in the message to remove information on
    // the van position
    ...
}

public void subscribeToVanLocation(Client client) {
    try {

        VanPositionUpdater vp = new VanPositionUpdater();

        Command command = new Command("sow_and_subscribe")
            .setTopic("van_location")
            .setFilter("/status = 'ACTIVE'")
            .setBatchSize(100)
            .setOptions("oof_enabled");

        client.execute(command, vp);

    } catch (AMPSEException aex) {
        System.err.println("TestListener caught exception.");
    }
}

```

Example 6.4. Asynchronous SOW and Subscribe

6.3. Setting Batch Size

The AMPS clients include a batch size parameter that specifies how many messages the AMPS server will return to the client in a single batch when returning the results of a SOW query. The 60East clients set a batch size of 10 by default. This batch size works well for common message sizes and network configurations.

Adjusting the batch size may produce better network utilization and produce better performance overall for the application. The larger the batch size, the more messages AMPS will send to the network layer at a time. This can result in fewer packets being sent, and therefore less overhead in the network layer. The effect on performance is generally most noticeable for small messages, where setting a larger batch size will allow several messages to fit into a single packet. For larger messages, a batch size may still improve performance, but the improvement is less noticeable.

In general, 60East recommends setting a batch size that is large enough to produce few partially-filled packets. Bear in mind that AMPS holds the messages in memory while batching them, and the client must also hold the messages in memory while receiving the messages. Using batch sizes that require large amounts of memory for these operations can reduce overall application performance, even if network utilization is good.

For smaller message sizes, 60East recommends using the default batch size, and experimenting with tuning the batch size if performance improvements are necessary. For relatively large messages (especially messages with sizes over 1MB), 60East recommends explicitly setting a batch size of 1 as an initial value, and increasing the batch size only if performance testing with a larger batch size shows improved network utilization or faster overall performance.

6.4. Client-Side Conflation

In many cases, applications that use SOW topics only need the current value of a message at the time the message is processed, rather than processing each change that lead to the current value. On the server side, AMPS provides *conflated topics* to meet this need. Conflated topics are described in more detail in the *AMPS User Guide*, and require no special handling on the client side.

In some cases, though, it's important to conflate messages on the client side. This can be particularly useful for applications that do expensive processing on each message, applications that are more efficient when processing batches of messages, or for situations where you cannot provide an appropriate conflation interval for the server to use.

A `MessageStream` has the ability to conflate messages received for a subscription to a SOW topic, view, or conflated topic. When conflation is enabled, for each message received, the client checks to see whether it has already received an unprocessed message with the same `SowKey`. If so, the client replaces the unprocessed message with the new message. The application never receives the message that has been replaced.

To enable client-side conflation, you call `conflate()` on the `MessageStream`, and then use the `MessageStream` as usual:

```
// Query and subscribe
MessageStream results =
    ampsClient.sowAndSubscribe("orders", "/symbol == 'ROL'");

// Turn on conflation
results.conflate();

// Process the results
foreach (Message m : results)
{
    // Process message here
}
```

Notice that if the `MessageStream` is used for a subscription that does not include `SowKeys` (such as a subscription to a topic that does not have a SOW), no conflation will occur.

When using client-side conflation with delta subscriptions, bear in mind that client-side conflation replaces the whole message, and does not attempt to merge deltas. This means that updates can be lost when messages are replaced. For some applications (for example, a ticker application that simply sends delta updates that replace the current price), this causes no problems. For other applications (for example, when several processors may be updating different fields of a message simultaneously), using conflation with deltas could result in lost data, and server-side conflation is a safer alternative.

6.5. Managing SOW Contents

AMPS allows applications to manage the contents of the SOW by explicitly deleting messages that are no longer relevant. For example, if a particular delivery van is retired from service, the application can remove the record for the van by deleting the record for the van.

The client provides the following methods for deleting records from the SOW.

- `sowDelete` accepts a filter, and deletes all messages that match the filter
- `sowDeleteByKeys` accepts a set of SOW keys as a comma-delimited string and deletes messages for those keys, regardless of the contents of the messages. A SOW key is provided in the header of a SOW message, and is the internal identifier AMPS uses for that SOW message
- `sowDeleteByData` accepts a message, and deletes the record that would be updated by that message

Most applications use `sowDelete`, since this is the most useful and flexible method for removing items from the SOW. In some cases, particularly when working with extremely large SOW databases, `sowDeleteByKeys` can provide better performance.

Regardless of the command used, AMPS sends an OOF message to all subscribers who have received updates for the messages removed, as described in the previous section.

The simple form of the `sowDelete` command returns a `MessageStream` that receives the response. The response is an acknowledgement message that contains information on the delete command. For example, the following snippet simply prints informational text with the number of messages deleted:

```
for (Message msg : client.sowDelete("sow_topic",
                                   "/id IN (42, 64, 37)"))
{
    System.out.println("Got an " + msg.getCommand() +
                      " containing " + msg.getAckType() + ": " +
                      " deleted " + msg.getMatches() + " messages.");
}
```

You can also use `client.execute` to send a SOW delete command. As with the other SOW methods, the client provides an asynchronous versions of the SOW delete commands that require a message handler to be invoked.

Acknowledging messages from a queue uses a form of the `sow_delete` command that is only supported for queues. Acknowledgement is discussed in the chapter on queues.

Chapter 7. Using Queues

AMPS message queues provide a high-performance way of distributing messages across a set of workers. The *AMPS User Guide* describes AMPS queues in detail, including the features of AMPS referred to in this chapter. This chapter does not describe message queues in detail, but instead explains how to use the AMPS Java client with message queues.

To publish messages to an message queue, publishers simply publish to any topic that is collected by the queue. There is no difference between publishing to a queue and publishing to any other topic, and a publisher does not need to be aware that the topic will be collected into a queue.

Subscribers must be aware that they are subscribing to a queue, and acknowledge messages from the queue when the message is processed.

7.1. Backlog and Smart Pipelining

AMPS queues are designed for high-volume applications that need minimal latency and overhead. One of the features that helps performance is the *subscription backlog* feature, which allows applications to receive multiple messages at a time. The subscription backlog sets the maximum number of unacknowledged messages that AMPS will provide to the subscription.

When the subscription backlog is larger than 1, AMPS delivers additional messages to a subscriber before the subscriber has acknowledged the first message received. This technique allows subscribers to process messages as fast as possible, without ever having to wait for messages to be delivered. The technique of providing a consistent flow of messages to the application is called *smart pipelining*.

Subscription Backlog

The AMPS server determines the backlog for each subscription. An application can set the maximum backlog that it is willing to accept with the `max_backlog` option. Depending on the configuration of the queue (or queues) specified in the subscription, AMPS may assign a smaller backlog to the subscription. If no `max_backlog` option is specified, AMPS uses a `max_backlog` of 1 for that subscription.

In general, applications that have a constant flow of messages perform better with a `max_backlog` setting higher than 1. The reason for this is that, with a backlog greater than 1, the application can always have a message waiting when the previous message is processed. Setting the optimum `max_backlog` is a matter of understanding the messaging pattern of your application and how quickly your application can process messages.

To request a `max_backlog` for a subscription, you explicitly set the option on the `subscribe` command, as shown below:

```
Command cmd = new Command("subscribe")
    .setTopic("my_queue")
    .setOptions("max_backlog=10");
```

Acknowledging Messages

For each message delivered on a subscription, AMPS counts the message against the subscription backlog until the message is explicitly acknowledged. In addition, when a queue specifies `at-least-once` delivery, AMPS re-

tains the message in the queue until the message expires or until the message has been explicitly acknowledged and removed from the queue. From the point of view of the AMPS server, acknowledgement is implemented as a `sow_delete` from the queue with the bookmarks of the messages to remove. The AMPS Java client provides several ways to make it easier for applications to create and send the appropriate `sow_delete`.

Automatic Acknowledgement

The AMPS client allows you to specify that messages should be automatically acknowledged. When this mode is on, AMPS acknowledges the message automatically in the following cases:

- *Asynchronous message processing interface.* The message handler returns without throwing an exception.
- *Synchronous message processing interface.* The application requests the next message from the `MessageStream`.

AMPS batches acknowledgements created with this method, as described in the following section.

To enable automatic acknowledgement batching, use the `setAutoAck()` method.

```
client.setAutoAck(true); // enable AutoAck
```

Message Convenience Method

The AMPS Java client provides a convenience method, `ack()`, on delivered messages. When the application is finished with the message, the application simply calls `ack()` on the message.

For messages that originated from a queue with `at-least-once` semantics, this adds the bookmark from the message to the batch of messages to acknowledge. For other messages, this method has no effect.

```
message.ack(); // Add this message to the next
               // acknowledgement batch.
```

Manual Acknowledgement

To manually acknowledge processed messages and remove the messages from the queue, applications use the `sow_delete` command with the bookmarks of the messages to remove. Notice that AMPS only supports using a bookmark with `sow_delete` when removing messages from a queue, not when removing records from a SOW.

For example, given a `Message` object to acknowledge and a client, the code below acknowledges the message.

```
void acknowledgeSingle(Client client, Message message)
{
    Message acknowledge = client.allocateMessage();
    acknowledge.setCommand("sow_delete")
                .setTopic(message.getTopic())
                .setBookmark(message.getBookmark());
    client.send(acknowledge);
}
```

Example 7.1. Simple Queue Acknowledgement

In listing Example 7.1 the program creates a `sow_delete` command, specifies the topic and the bookmark, and then sends the command to the server. Because the program does not need or expect a response from AMPS, this function uses the `Message` object rather than the `Command` object.

While this method works, creating and sending an acknowledgement for each individual message can be inefficient if your application is processing a large volume of messages. Rather than acknowledging each message individually, your application can build a comma-delimited list of bookmarks from the processed messages and acknowledge all of the messages at the same time. In this case, it's important to be sure that the number of messages you wait for is less than the maximum backlog -- the number of messages your client can have unacknowledged at a given time. Notice that both automatic acknowledgement and the helper method on the `Message` object take the maximum backlog into account.

Acknowledgement Batching

The AMPS Java client automatically batches acknowledgements when either of the convenience methods is used. Batching acknowledgements reduces the number of round-trips to AMPS, reducing network traffic and improving overall performance. AMPS sends the batch of acknowledgements when the number of acknowledgements exceeds a specified size, or when the amount of time since the last batch was sent exceeds a specified timeout.

You can set the number of messages to batch and the maximum amount of time between batches:

```
client.setAckBatchSize(10); // Send batch after 10 messages
client.setAckTimeout(1000); // ... or 1 second
```

The AMPS Java client is aware of the subscription backlog for a subscription. When AMPS returns the acknowledgement for a subscription that contains queues, AMPS includes information on the subscription backlog for the subscription. If the batch size is larger than the subscription backlog, the AMPS Java client adjusts the requested batch size to match the subscription backlog.

Returning a Message to the Queue

A subscriber can also explicitly release a message back to the queue. AMPS returns the message to the queue, and redelivers the message just as though the lease had expired. To do this, the subscriber sends a `sow_delete` command with the bookmark of the message to release and the `cancel` option.

When using automatic acknowledgements and the asynchronous API, AMPS will cancel a message if an exception is thrown from the message handler.

Chapter 8. Delta Publish and Subscribe

8.1. Introduction

Delta messaging in AMPS has two independent aspects:

- *delta subscribe* allows subscribers to receive just the fields that are updated within a message.
- *delta publish* allows publishers to update and add fields within a message by publishing only the updates into the SOW.

This chapter describes how to create delta publish and delta subscribe commands using the AMPS C++ client. For a discussion of this capability, how it works, and how message types support this capability see the *AMPS User Guide*.

8.2. Delta Subscribe

To delta subscribe, you simply use the `delta_subscribe` command as follows:

```
// assumes that client is connected and logged on

Command cmd("delta_subscribe");
cmd.setTopic("delta_topic");
cmd.setFilter("/thingIWant = 'true'");

try (MessageStream ms = client.execute(cmd))
for (Message m : ms)
{
    // Work with message here
}
```

As described in the *AMPS User Guide*, messages provided to a delta subscription will contain the fields used to generate the SOW key and any changed fields in the message. Your application is responsible for choosing how to handle the changed fields.

8.3. Delta Publish

To delta publish, you use the `delta_publish` command as follows:

```
// assumes that client is connected and logged on

String msg = ... ; // obtain changed fields here

client.deltaPublish("myTopic", msg);
```

The message that you provide to AMPS must include the fields that the topic uses to generate the SOW key. Otherwise, AMPS will not be able to identify the message to update. For SOW topics that use a User-Generated SOW Key, use the Command form of `delta_publish` to set the `SowKey`.

```
// assumes that client is connected and logged on

String msg = ... ; // obtain changed fields here
String key = ... ; // obtain user-generated SOW key

Command cmd("delta_publish");
cmd.setTopic("delta_topic");
cmd.setSowKey(key);
cmd.setData(msg);

// Execute the delta publish. Use null for
// a message handler since any failure acks will
// be routed to the FailedWriteHandler
client.executeAsync(cmd,null);
```

Chapter 9. High Availability

The AMPS Java Client provides an easy way to create highly-available applications using AMPS, via the `HAClient` class. `HAClient` derives from `Client` and offers the same methods, but also adds protection against network, server, and client outages.

Using `HAClient` allows applications to automatically:

- Recover from temporary disconnects between client and server.
- Failover from one server to another when a server becomes unavailable.

Because the `HAClient` automatically manages failover and reconnection, 60East recommends using the `HAClient` for applications that need to:

- Ensure no messages are lost or duplicated after a reconnect or failover.
- Persist messages and bookmarks on disk for protection against client failure.

You can choose how your application uses `HAClient` features. For example, you might need automatic reconnection, but have no need to resume subscriptions or republish messages. The high availability behavior in `HAClient` is provided by implementations of defined interfaces. You can combine different implementations provided by 60East to meet your needs, and implement those interfaces to provide your own policies.

Some of these features require specific configuration settings on your AMPS instance(s). This chapter mentions these features and describes how to use them from the AMPS Java client. You can find full documentation for these settings and server features in the *User Guide*.

9.1. Reconnection with HAClient

The most important difference between `Client` and `HAClient` is that `HAClient` automatically provides a reconnect handler.

This description provides a high-level framework for understanding the components involved in failover with the `HAClient`. The components are described in more detail in the following sections.

The `HAClient` reconnect handler performs the following steps when reconnecting:

- Calls the `ServerChooser` to determine the next URI to connect to and the authenticator to use for that connection.

If the connection fails, calls `get_error` on the `ServerChooser` to get a description of the failure, sends an exception to the exception listener, and stops the reconnection process.

- Calls the `DelayStrategy` to determine how long to wait before attempting to reconnect, and waits for that period of time.
- Connects to the AMPS server. If the connection fails, calls `report_failure` on the `ServerChooser` and begins the process again.
- Logs on to the AMPS server. If the connection fails, calls `report_failure` on the `ServerChooser` and begins the process again.

- Calls `report_success` on the `ServerChooser`.
- Receives the bookmark for the last message that the server has persisted. Discards any older messages from the `PublishStore`.
- Republishes any messages in the `PublishStore` that have not been persisted by the server.
- Re-establishes subscriptions using the `SubscriptionManager` for the client. For bookmark subscriptions, the reconnect handler uses the `BookmarkStore` for the client to determine the most recent bookmark, and resubscribes with that bookmark. For subscriptions that do not use a bookmark, the `SubscriptionManager` simply re-enters the subscription, meaning that it is entered at the point at which the `HAClient` reconnects.

The `ServerChooser`, `DelayStrategy`, `PublishStore`, `SubscriptionManager`, and `BookmarkStore` are all extension points for the `HAClient`. You can adapt the failover and recovery behavior by setting a different object for the behavior you want to customize on the `HAClient` or by providing your own implementation.

For example, the convenience methods in the previous section customize the behavior of the `PublishStore` and `BookmarkStore` by providing either memory-backed or file-backed stores.

9.2. Choosing Store Durability

Most code written with `Client` will also work with `HAClient`, and major differences involve constructing and connecting the `HAClient`.

The `HAClient` provides recovery after disconnection using *Stores*. As the name implies, *stores* hold information about the state of the client. There are two types of store:

- A bookmark store tracks received messages, and is used to resume subscriptions.
- A publish store tracks published messages, and is used to ensure that messages are persisted in AMPS.

The AMPS Java Client provides a memory-backed version of each store and a file-backed version of each store. An `HAClient` can use either a memory backed store or a file backed store for protection. Each method provides resilience to different failures:

- Memory-backed stores provide recovery after disconnection from AMPS by storing messages and bookmarks in your process' address space. This is the highest performance option for working with AMPS in a highly available manner. The trade-off with this method is there is no protection from a crash or failure of your client application. If your application is terminated prematurely or, if the application terminates at the same time as an AMPS instance failure or network outage, then messages may be lost or duplicated.
- File-backed stores provide recovery after client failure and disconnection from AMPS by storing messages and bookmarks on disk. To use this protection method, the `create_file_backed` method requests additional arguments for the two files that will be used for both bookmark storage and message storage. If these files exist and are non-empty (as they would be after a client application is restarted), the `HAClient` loads their contents and ensures synchronization with the AMPS server once connected. The performance of this option depends heavily on the speed of the device on which these files are placed. When the files do not exist (as they would the first time a client starts on a given system), the `HAClient` creates and initializes the files, and in this case the client does not have a point at which to resume the subscription or messages to republish.

The store interface is public, and an application can create and provide a custom store as necessary. While clients provide convenience methods for creating file-backed and memory-backed `HAClient` objects with the appropriate

stores, you can also create and set the stores in your application code. The AMPS Java client also includes default stores, which implement the appropriate interface, but do not actually persist messages.

The `HAClient` provides convenience methods for creating clients and setting stores. You can also construct an `HAClient` and set whichever store implementations you choose.

In this example, we create several clients. The first client uses memory stores for both bookmarks and publishes. The second client uses files for both bookmarks and publishes. The third client uses a file for bookmarks. The third client does not set a store for publishes, which means that AMPS provides the default store (and no outgoing messages are stored). The final client does not specify any stores, and so has no persistence for published messages or bookmark subscriptions, but can take advantage of the automatic failover and reconnection in the `HAClient`.

```
// Memory publish store, memory bookmark store
HAClient memoryClient = HAClient.createMemoryBacked(
    "lessImportantMessages");

// File-backed publish store, file-backed bookmark store
HAClient diskClient = HAClient.createFileBacked(
    "moreImportantMessages",
    "/mnt/fastDisk/moreImportantMessages.outgoing",
    "/mnt/fastDisk/moreImportantMessages.incoming");

// Default publish store, file-backed bookmark store
HAClient subscriberClient = new HAClient("subscriber");
subscriberClient
    .setBookmarkStore(
        new LoggedBookmarkStore("/mnt/fastDisk/bookmark.store"));

// Default publish store, default bookmark store
// Failover behavior only
HAClient streamReader = new HAClient("streamReader");
```

Example 9.1. HAClient creation examples



While this chapter presents the built-in file and memory-based stores, the AMPS Java Client provides open interfaces that allow development of custom persistent message stores. You can implement the `Store` and `BookmarkStore` interfaces in your code, and then pass instances of those to `setPublishStore()` or `setBookmarkStore()` methods in your `Client`. Instructions on developing a custom store are beyond the scope of this document; please refer to the *AMPS Client HA Whitepaper* for more information.

9.3. Connections and the ServerChooser

Unlike `Client`, the `HAClient` attempts to keep itself connected to an AMPS instance at all times, by automatically reconnecting or failing over when it detects that the client is disconnected. When you are using the `Client` directly, your disconnect handler usually takes care of reconnection. `HAClient`, on the other hand, provides a disconnect handler that automatically reconnects to the current server or to the next available server.

To inform the `HAClient` of the addresses of the AMPS instances in your system, you pass a `ServerChooser` instance to the `HAClient`. `ServerChooser` acts as a smart enumerator over the servers available: `HAClient`

calls `ServerChooser` methods to inquire about what server should be connected, and calls methods to indicate whether a given server succeeded or failed.

The AMPS Java Client provides a simple implementation of `ServerChooser`, called `DefaultServerChooser`, that provides very simple logic for reconnecting. This server chooser is most suitable for basic testing, or in cases where an application should simply rotate through a list of servers. For most applications, you implement the `ServerChooser` interface yourself for more advanced logic, such as choosing a backup server based on your network topology, or limiting the number of times your application should try to reconnect to a given address.

To connect to AMPS, you provide a `ServerChooser` to `HAClient` and then invoke `connectAndLogon()` to create the first connection:

```
HAClient myClient = HAClient.createMemoryBacked(
    "myClient");

// primary.amps.xyz.com is the primary AMPS instance, and
// secondary.amps.xyz.com is the secondary
DefaultServerChooser chooser =
    new DefaultServerChooser();
chooser.add("tcp://primary.amps.xyz.com:12345/amps/json");
chooser.add("tcp://secondary.amps.xyz.com:12345/amps/json");
myClient.setServerChooser(chooser);
myClient.connectAndLogon();
...
myClient.disconnect();
```

Example 9.2. HAClient logon

Similar to `Client`, `HAClient` remains connected to the server until `disconnect()` is called. Unlike `Client`, `HAClient` automatically attempts to reconnect to your server if it detects a disconnect and, if that server cannot be connected, fails over to the next server provided by the `ServerChooser`. In this example, the call to `connectAndLogon()` attempts to connect and login to `primary.amps.xyz.com`, and returns if that is successful. If it cannot connect, it tries `secondary.amps.xyz.com`, and continues trying servers from the `ServerChooser` until a connection is established. Likewise, if it detects a disconnection while the client is in use, then `HAClient` attempts to reconnect to the server with which it was most recently connected; if that is not possible, then it moves on to the next server provided by the `ServerChooser`.



While this chapter presents the built-in file and memory-based stores, the AMPS Java Client provides open interfaces that allow development of custom persistent message stores. You can implement the `Store` and `BookmarkStore` interfaces in your code, and then pass instances of those to `setPublishStore()` or `setBookmarkStore()` methods in your `Client`. Instructions on developing a custom store are beyond the scope of this document; please refer to the *AMPS Client HA Whitepaper* for more information.

Setting a Reconnect Delay and Timeout

You can control the amount of time between reconnection attempts and set a total amount of time for the `HAClient` to attempt to reconnect.

The AMPS Java Client includes an interface for managing this behavior called the `ReconnectDelayStrategy`.

Two implementations of this interface are provided with the client:

- `FixedDelayStrategy` provides the same delay each time the `HAClient` tries to reconnect.
- `ExponentialDelayStrategy` provides an exponential backoff until a connection attempt succeeds.

To use either of these classes, you simply create an instance, set the appropriate parameters, and install that instance as the delay strategy for the `HAClient`. For example, the following code sets up a reconnect delay that starts at 200ms and increases the delay by 1.5 times after each failure. The strategy allows a maximum delay between connection attempts of 5 seconds, and will not retry longer than 60 seconds.

```
HAClient theClient = HAClient.createMemoryBacked("demo");

ExponentialDelayStrategy theStrategy = new ExponentialDelayStrategy();
theStrategy.setInitialDelay(200);
theStrategy.setBackoffExponent(1.5);
theStrategy.setMaximumDelay(5000);
theStrategy.setMaximumRetryTime(60000);

theClient.setDelayStrategy(theStrategy);
```

Implementing a Server Chooser

As described above, you provide the `HAClient` with connection strings to one or more AMPS servers using a `ServerChooser`. The purpose of a `ServerChooser` is to provide information to the `HAClient`. A `ServerChooser` does not manage the reconnection process, and should not call methods on the `HAClient`.

A `ServerChooser` has two required responsibilities to the `HAClient`:

- Tells the `HAClient` the connection string for the server to connect to. If there are no servers, or the `ServerChooser` wants the connection to fail, the `ServerChooser` returns an empty string.

To provide this information, the `ServerChooser` implements the `getCurrentURI()` method.

- Provides an `Authenticator` for the current connection string. This is especially important for installations where different servers require different credentials or authentication tokens must be reset after each connection attempt.

To provide the authenticator, the `ServerChooser` implements the `getCurrentAuthenticator()` method.

The `HAClient` calls the `getCurrentURI()` and `getCurrentAuthenticator()` methods each time it needs to make a connection.

Each time a connection succeeds, the `HAClient` calls the `reportSuccess()` method of the `ServerChooser`. Each time a connection fails, the `HAClient` calls the `reportFailure()` method of the `ServerChooser`. The `HAClient` does not require the `ServerChooser` to take any particular action when it calls these methods. These methods are provided for the `HAClient` to do internal maintenance, logging, or record keeping. For example, an `HAClient` might keep a list of available URIs with a current failure count, and skip over URIs that have failed more than 5 consecutive times until all URIs in the list have failed more than 5 consecutive times.

When the `ServerChooser` returns an empty string from `getCurrentURI()`, indicating that no servers are available for connection, the `HAClient` calls `getError()` method on the `ServerChooser` and includes the string returned by `getError()` in the generated exception.

9.4. Heartbeats and Failure Detection

Use of the `HAClient` allows your application to quickly recover from detected connection failures. By default, connection failure detection occurs when AMPS receives an operating system error on the connection. This system may result in unpredictable delays in detecting a connection failure on the client, particularly when failures in network routing hardware occur, and the client primarily acts as a subscriber.

The heartbeat feature of the AMPS client allows connection failure to be detected quickly. Heartbeats ensure that regular messages are sent between the AMPS client and server on a predictable schedule. The AMPS client and server both assume disconnection has occurred if these regular heartbeats cease, ensuring disconnection is detected in a timely manner. To use heartbeating, call the `setHeartbeat` method on `Client` or `HAClient`:

```
HAClient client = HAClient.createMemoryBacked(
    "importantStuff");
...
client.setHeartbeat(3);
client.connectAndLogon();
...
```

Method `setHeartbeat` takes one parameter: the heartbeat interval. The heartbeat interval specifies the periodicity of heartbeat messages sent by the server: the value 3 indicates messages are sent on a three-second interval. If the client receives no messages in a six-second window (two heartbeat intervals), the connection is assumed to be dead, and the `HAClient` attempts reconnection. An additional variant of `setHeartbeat` allows the idle period to be set to a value other than two heartbeat intervals.

Notice that, for `HAClient`, `setHeartbeat` must be called before the client is connected. For `Client`, `setHeartbeat` may be called either before or after the client is connected.

Heartbeats are serviced on the receive thread created by the AMPS client. Your application must not block the receive thread for longer than the heartbeat interval, or the application is subject to being disconnected.

9.5. Considerations for Publishers

Publishing with an `HAClient` is nearly identical to regular publishing; you simply call the `publish()` method with your message's topic and data. The AMPS client sends the message to AMPS, and then returns from the `publish()` call. For maximum performance, the client does not wait for the AMPS server to acknowledge that the message has been received.

When an `HAClient` uses a publish store (other than the `DefaultPublishStore`), the publish store retains a copy of each outgoing message and requests that AMPS acknowledge that the message has been persisted. The AMPS server acknowledges messages back to the publisher. Acknowledgements can be delivered for multiple messages at periodic intervals (for topics recorded in the transaction log) or after each message (for topics that are not recorded in the transaction log). When an acknowledgement for a message is received, the `HAClient` removes that message from the bookmark store. When a connection to a server is made, the `HAClient` automatically determines which

messages from the publish store (if any) the server has not processed, and replays those messages to the server once the connection is established.

For reliable publishers, the application must choose how best to handle application shutdown. For example, it is possible for the network to fail immediately after the publisher sends the message, while the message is still in transit. In this case, the publisher has sent the message, but the server has not processed it and acknowledged it. During normal operation, the `HAClient` will automatically connect and retry the message. On shutdown, however, the application must decide whether to wait for messages to be acknowledged, or whether to exit.

Publish store implementations provide an `unpersistedCount()` method that reports the number of messages that have not yet been acknowledged by the AMPS server. When the `unpersistedCount()` reaches 0, there are no unpersisted messages in the local publish store.

For the highest level of safety, an application can wait until the `unpersistedCount()` reaches 0, which indicates that all of the messages have been persisted to the instance that the application is connected to, and the synchronous replication destinations configured for that instance. When a synchronous replication destination goes offline, this approach will cause the publisher to wait to exit until the destination comes back online or until the destination is downgraded to asynchronous replication.

For applications that are shut down periodically for short periods of time (for example, applications that are only offline during a weekly maintenance window), another approach is to use the `publishFlush()` method to ensure that messages are delivered to AMPS, and then rely on the connection logic to replay messages as necessary when the application restarts.

For example, the following code flushes messages to AMPS, then warns if not all messages have been acknowledged:

```
HAClient pub = HAClient.createMemoryBacked(
    "importantStuff");
...
pub.connectAndLogon();
String topic = "loggedTopic";
String data = ...;
for(int i = 0; i < MESSAGE_COUNT; i++)
{
    pub.publish(topic, data);
}

// We think we are done, but the server may not
// have acknowledged us yet.
while(pub.getPublishStore().unpersistedCount() > 0)
{
    System.out.println("waiting for final ack from "+
        "the server...");
    Thread.sleep(1000);
}
pub.disconnect();
```

Example 9.3. HA Publisher

In this example, the client sends each message immediately when `publish()` is called. If AMPS becomes unavailable between the final `publish()` and the `disconnect()`, or one of the servers that the AMPS instance

replicates to is offline, the client may not have received a persisted acknowledgement for all of the published messages. For example, if a message has not yet been persisted by all of the servers in the replication fabric that are connected with synchronous replication, AMPS will not have acknowledged the message.

This code first flushes messages to the server to ensure that all messages have been delivered to AMPS.

The code next checks to see if all of the messages in the publish store have been acknowledged as persisted by AMPS. If the messages have not been acknowledged, they will remain in the publish store file and will be published to AMPS, if necessary, the next time the application connects. An application may choose to loop until `unpersistedCount()` returns 0, or (as we do in this case) simply warn that AMPS has not confirmed that the messages are fully persisted. The behavior you choose in your application should be consistent with the high-availability guarantees your application needs to provide.



AMPS uses the name of the `HAClient` to determine the origin of messages. For the AMPS server to correctly identify duplicate messages, each instance of an application that publishes messages must use a distinct name. That name must be consistent across different runs of the application.

If your application crashes or is terminated, some published messages may not have been persisted in the AMPS server. If you use the file-based store—in other words, the store created by using `HAClient.createFileBacked()`—the `HAClient` will recover the messages, and once logged on, will correlate the message store to what the AMPS server has received, re-publishing any missing messages. This occurs automatically when `HAClient` connects, without any explicit consideration in your code, other than ensuring that the same file name is passed to `createFileBacked()` if recovery is desired.



AMPS provides persisted acknowledgement messages for topics that do not have a transaction log enabled. However, the level of durability provided for topics with no transaction log is minimal. Learn more about transaction logs in the *User Guide*.

9.6. Considerations for Subscribers

`HAClient` provides two important features for applications that subscribe to one or more topics: re-subscription, and a bookmark store to track the correct point at which to resume a bookmark subscription.

Resubscription With Asynchronous Message Processing

Any asynchronous subscription placed using an `HAClient` is automatically reinstated after a disconnect or a failover. These subscriptions are placed in an in-memory `SubscriptionManager`, which is created automatically when the `HAClient` is instantiated. Most applications will use this built-in subscription manager, but for applications that create a varying number of subscriptions, you may wish to implement `SubscriptionManager` to store subscriptions in a more durable place. Note that these subscriptions contain no message data, but rather simply contain the the parameters of the subscription itself (for instance, the command, topic, message handler, options, and filter).

When a re-subscription occurs, the AMPS Java Client re-executes the command as originally submitted, including the original topic, options, and so on. AMPS sends the subscriber any messages for the specified topic (or topic expression) that are published after the subscription is placed. For a `sow_and_subscribe` command, this means that the client reissues the full command, including the `SOW` query as well as the subscription.

Resubscription With Synchronous Message Processing

The `HAClient` (starting with the AMPS Java Client version 4.3.1.2) does not track synchronous message processing subscriptions in the `SubscriptionManager`. The reason for this is to preserve the expected behavior of an `Iterator`. That is, once the `MessageStream` indicates that there are no more elements in the stream, the `MessageStream` does not suddenly produce more elements.

To resubscribe when the `HAClient` fails over, you can simply reissue the subscription. For example, the snippet below re-issues the subscribe command when the message stream ends:

```
boolean still_need_to_process = true;

while (still_need_to_process == true)
{
    MessageStream ms = client.subscribe("topic");
    try
    {
        for (Message m : ms)
        {
            // process message

            // check condition on still_need_to_process
            if (still_need_to_process == false) break;
        }
        // end of stream, for a subscribe this means
        // that the connection is likely closed.
    }
    finally
    {
        if (ms != null) ms.close();
    }
}
```

Bookmark Stores

In cases where it is critical not to miss a single message, it is important to be able to resume a subscription at the exact point that a failure occurred. In this case, simply recreating a subscription isn't sufficient. Even though the subscription is recreated, the subscriber may have been disconnected at precisely the wrong time, and will not see the message.

To ensure delivery of every message from a topic or set of topics, the AMPS `HAClient` includes a `BookmarkStore` that, combined with the bookmark subscription and transaction log functionality in the AMPS server, ensures that clients receive any messages that might have been missed. The client stores the bookmark associated with each message received, and tracks whether the application has processed that message; if a disconnect occurs, the client uses the `BookmarkStore` to determine the correct resubscription point, and sends that bookmark to AMPS when it re-subscribes. AMPS then replays messages from its transaction log from the point after the specified bookmark, thus ensuring the client is completely up-to-date.

`HAClient` helps you to take advantage of this bookmark mechanism through the `BookmarkStore` interface and `bookmarkSubscribe()` method on `Client`. Whenever a disconnection or failover occurs for subscriptions created with `bookmarkSubscribe()`, your application will automatically resubscribe to the message after the

last message it processed. HA Clients created by `createFileBacked()` additionally store these bookmarks on disk, so that the application can restart with the appropriate message if the client application fails and restarts.

To take advantage of the `BookmarkStore` and bookmark subscriptions, do the following:

- Ensure the topic(s) to be subscribed are included in a transaction log. See the *User Guide* for information on how to specify the contents of a transaction log.
- Use `bookmarkSubscribe()` instead of `subscribe()` when creating a subscription(), and decide how the application will manage subscription identifiers (SubIds).
- Use the `BookmarkStore.discard()` method in message handlers to indicate when a message has been fully processed by the application.

The following example creates a bookmark subscription against a transaction-logged topic, and fully processes each message as soon as it is delivered:

```
final HAClient client = HAClient.createFileBacked(
    "aClient",
    "/logs/aClient.publishLog",
    "/logs/aClient.subscribeLog");

class MyMessageHandler implements MessageHandler
{
    public void invoke(Message message)
    {
        ...
        client.getBookmarkStore().discard(
            message.getSubIdRaw(),
            message.getBookmarkSeqNo());
        ...
    }
}

Command command = new Command("subscribe")
    .setTopic("myTopic")
    .setSubscriptionId(new CommandId("MySubId"))
    .setBookmark(Client.Bookmarks.MOST_RECENT);

client.executeAsync(command, new myMessageHandler());
```

Example 9.4. HA Client Subscription

In this example, the client is a file-backed client, meaning that arriving bookmarks will be stored in a file (`Client.subscribeLog`). Storing these bookmarks in a file allows the application to restart the subscription from the last message processed, in the event of either server or client failure.



For optimum performance, it is critical to discard every message once its processing is complete. If a message is never discarded, it remains in the bookmark store. During re-subscription, HA Client always restarts the bookmark subscription with the oldest undiscarded message, and then filters out any more recent messages that have been discarded. If an old message remains in the store, but is no longer important for the application's functioning, then the client and the AMPS server will incur unnecessary network, disk, and CPU activity.

The `SubscriptionId` field specifies an identifier to be used for this subscription. Passing null, or leaving the field unset, causes `HAClient` to generate a new subscription ID, like most other `Client` functions. However, if you wish to resume a subscription from a previous point after the application has terminated and restarted, the application must pass the same subscription ID as during its previous run. Passing a different subscription ID bypasses any recovery mechanisms, creating an entirely new subscription. When you use an existing subscription ID, the `HAClient` locates the last-used bookmark for that subscription in the local store, and attempts to re-subscribe from that point.

- `Client.Bookmarks.NOW` specifies that the subscription should begin from the moment the server receives the subscription request. This results in the same messages being delivered as if you had invoked `subscribe()` instead, except that the messages will be accompanied by bookmarks. This is also the behavior that results if you supply an invalid bookmark.
- `Client.Bookmarks.EPOCH` specifies that the subscription should begin from the beginning of the AMPS transaction log.
- `Client.Bookmarks.MOST_RECENT` specifies that the subscription should begin from the last-used message in the associated `BookmarkStore`. Alternatively, if this subscription has not been seen before, it instructs the subscription to begin with `EPOCH`. This is the most common value for this parameter, and is the value used in the preceding example. By using `MOST_RECENT`, the application automatically resumes from wherever the subscription left off, taking into account any messages that have already been processed and discarded.

When the `HAClient` re-subscribes after a disconnection and reconnection, it always uses `MOST_RECENT`, ensuring that the continued subscription always begins from the last message used before the disconnect, so that no messages are missed.

9.7. Conclusion

With only a few changes, most AMPS applications can take advantage of the `HAClient` and associated classes to become more highly-available and resilient. Using the `PublishStore`, publishers can ensure that every message published has actually been persisted by AMPS. Using `BookmarkStore`, subscribers can make sure that there are no gaps or duplicates in the messages received. `HAClient` makes both kinds of applications more resilient to network and server outages and temporary issues, and, by using the filebased `HAClient`, clients can recover their state after an unexpected termination or crash. Though `HAClient` provides useful defaults for the `Store`, `BookmarkStore`, `SubscriptionManager`, and `ServerChooser`, you can customize any or all of these to the specific needs of your application and architecture.

Chapter 10. AMPS Programming: Working with Commands

The AMPS clients provide named convenience methods for core AMPS functionality. These named methods work by creating messages and sending those messages to AMPS. All communication with AMPS occurs through messages.

You can use the `Command` object to customize the messages that AMPS sends. This is useful for more advanced scenarios where you need precise control over AMPS, in cases where you need to use an earlier version of the client to communicate with a more recent version of AMPS, or in cases where a named method is not available.

10.1. Understanding AMPS Messages

AMPS messages are represented in the client as `AMPS.Message` objects. The `Message` object is generic, and can represent any type of AMPS message, including both outgoing and incoming messages. This section includes a brief overview of elements common to AMPS command message. Full details of the commands to AMPS are provided in the *AMPS Command Reference Guide*.

All AMPS command messages contain the following elements:

- **Command.** The *command* tells AMPS how to interpret the message. Without a command, AMPS will reject the message. Examples of commands include `publish`, `subscribe`, and `sow`.
- **CommandId.** The *command id*, together with the name of the client, uniquely identifies a command to AMPS. The command ID can be used later on to refer to the command or the results of the command. For example, the command id for a `subscribe` message becomes the identifier for the subscription. The AMPS client provides a command id when the command requires one and no command id is set.

Most AMPS commands contain the following fields:

- **Topic.** The *topic* that the command applies to, or a regular expression that identifies a set of topics that the command applies to. For most commands, the topic is required. Commands such as `logon`, `start_timer`, and `stop_timer` do not apply to a specific topic, and do not need this field.
- **Ack Type.** The *ack type* tells AMPS how to acknowledge the message to the client. Each command has a default acknowledgement type that AMPS uses if no other type is provided.
- **Options.** The *options* are a comma-separated list of options that affect how AMPS processes and responds to the message.

Beyond these fields, different commands include fields that are relevant to that particular command. For example, SOW queries, subscriptions, and some forms of SOW deletes accept the **Filter** field, which specifies the filter to apply to the subscription or query. As another example, publish commands accept the **Expiration** field, which sets the SOW expiration for the message.

For full details on the options available for each command and the acknowledgement messages returned by AMPS, see the *AMPS Command Reference Guide*.

10.2. Creating and Populating the Command

To create a command, you simply construct a command object of the appropriate type:

```
Command command = new Command("sow");
```

Once created, you set the appropriate fields on the message. For example, the following code creates a publish message, setting the command, topic, data to publish, and an expiration for the message:

```
Command command = new Command("sow");  
command.setTopic("messages-sow");  
command.setFilter("/id > 20");
```

When sent to AMPS, AMPS performs a SOW query from the topic `messages-sow` using the filter of `/id > 20`. The results of sending this message to AMPS are no different than using the form of the `sow` method that sets these fields.

10.3. Using execute

Once you've created a message, use the `execute` method to send the message to AMPS. The `execute` method returns a `MessageStream` that provides response messages. The `executeAsync` method sends the message to AMPS, waits for a processed acknowledgement, then returns. Messages are processed on the client background thread.

For example, the following snippet sends the command created above:

```
client.execute(command);
```

This returns a `MessageStream` identical to the `MessageStream` returned by the equivalent `client.sow()` method.

You can also provide a message handler to receive acknowledgements, statistics, or the results of subscriptions and SOW queries. The AMPS client maintains a background thread that receives and processes incoming messages. The call to `executeAsync` returns on the main thread as soon as AMPS acknowledges the command as having been processed, and messages are received and processed on the background thread:

```
import com.crankuptheamps.client.Message;  
import com.crankuptheamps.client.MessageHandler;  
  
class SimpleMessageHandler implements MessageHandler {  
  
    // For sample purposes, just print the acknowledgement type and  
    // reason.  
  
    public void invoke(Message m)  
    {  
        System.out.println(m.getAckType() + " : " + m.getReason() );  
    }  
}
```

Then, to send the command and use the message handler, pass the command and the handler to `executeAsync()`.

```
client.executeAsync(command, new SimpleMessageHandler());
```

While this message handler simply prints the ack type and reason for sample purposes, message handlers in production applications are typically designed with a specific purpose. For example, your message handler may fill a work queue, or check for success and throw an exception if the command failed.

Notice that the `publish` command does not provide typically return results other than acknowledgement messages. To send a `publish` command, use the `executeAsync()` method with a null message handler:

```
client.executeAsync(publishCmd, null);
```

10.4. Command Cookbook

This section is a quick guide to commonly used AMPS commands. For the full range of options on AMPS commands, see the *AMPS Command Reference*.

Publishing

This section presents common recipes for publishing to a topic in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

The AMPS server does not return a stream of messages in response to a `publish` command.



AMPS `publish` commands do not return a stream of messages. A `publish` command must be used with asynchronous message processing, while passing an empty message handler.

Basic Publish

In its simplest form, a subscription needs only the topic to publish to and the data to publish. The AMPS client automatically constructs the necessary AMPS headers and formats the full `publish` command.

In many cases, a publisher only needs to use the basic `publish` command.

Table 10.1. Basic Publish

Header	Comment
Topic	<p>Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.</p> <p>Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.</p>

Header	Comment
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.

Publish With CorrelationId

AMPS provides publishers with a header field that can be used to contain arbitrary data, the `CorrelationId`.

Table 10.2. Publish With CorrelationId

Header	Comment
Topic	<p>Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.</p> <p>Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.</p>
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.
CorrelationId	<p>The <code>CorrelationId</code> to provide on the message. AMPS provides the <code>CorrelationId</code> to subscribers. The <code>CorrelationId</code> has no significance for AMPS.</p> <p>The <code>CorrelationId</code> may only contain characters that are valid in base-64 encoding.</p>

Publish With Explicit SOW Key

When publishing to a SOW topic that is configured to require an explicit SOW key, the publisher needs to set the `SowKey` header on the message.

Table 10.3. Publish with Explicit SOW Key

Header	Comment
Topic	<p>Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.</p> <p>Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.</p>
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.
SowKey	The SOW Key to use for this message. This header is only supported for publishes to a topic that requires an explicit SOW Key.

Command Cookbook: Subscribing

This section presents common recipes for subscribing to a topic in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic Subscription

In its simplest form, a subscription needs only the topic to subscribe to.

Table 10.4. Basic Subscription

Header	Comment
Topic	Sets the topic to subscribe to. All messages from the topic will be delivered on this subscription. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

Basic Subscription With Options

In its simplest form, a subscription needs only the topic to subscribe to. To add options to the subscription, set the `Options` header on the `Command`.

Table 10.5. Basic Subscription with Options

Header	Comment
Topic	Sets the topic to subscribe to. All messages from the topic will be delivered on this subscription. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Options	A comma-delimited set of options for this command. See the <i>AMPS Command Reference</i> for a description of supported options.

Content Filtered Subscription

To provide a content filter on a subscription, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.6. Content Filtered Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Filter	Sets the content filter to be applied to the subscription. Only messages that match the content filter will be provided to the subscription.

Conflated Subscription to a SOW Topic

To request conflation on a subscription, set the `Options` property on the command to specify the conflation interval. When the topic has a SOW (including a view or a conflated topic), there is no need to provide a `conflation_key`.

Table 10.7. Conflated Subscription to a SOW topic

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Options	<p>A comma-separated list of options for the command. Set the conflation interval in the options.</p> <p>For example, to set a conflation interval of 250 milliseconds, provide the following option:</p> <pre>conflation=250ms</pre> <p>To set the conflation interval to 1 minute, provide an option of:</p> <pre>conflation=1m</pre>

Conflated Subscription to a Topic With No SOW

To request conflation on a subscription, set the `Options` property on the command to specify the conflation interval. When the topic does not have a SOW, you must provide a `conflation_key`.

Table 10.8. Conflated Subscription to a SOW topic

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Options	<p>A comma-separated list of options for the command. Set the conflation interval and the fields that determine a unique message in the options.</p> <p>For example, to set a conflation interval of 250 milliseconds for messages that have the same value for <code>/id</code>, provide the following option:</p> <pre>conflation=250ms,conflation_key=[/id]</pre> <p>To set the conflation interval to 1 minute for messages are the same as determined by a combination of <code>/customerId</code> and <code>/issueNumber</code> provide an option of:</p> <pre>conflation=1m, conflation_key=[/customerId,/ issueNumber]</pre>

Bookmark Subscription

To create a bookmark subscription, set the `Bookmark` property on the command. The value of this property can be either a specific bookmark, a timestamp, or one of the client-provided constants. The *AMPS User Guide* provides details on creating timestamps. Notice that the `MOST_RECENT` constant tells the AMPS client to find the appropriate message in the client bookmark store and begin the subscription at that point. In this case, the client sends that bookmark value to AMPS. The `Bookmark` option is only supported for topics that are recorded in an AMPS transaction log.

Table 10.9. Bookmark Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Bookmark	<p>Sets the point in the transaction log at which the subscription will begin. The bookmark provided can be a specific AMPS bookmark, a timestamp, or one of the client-provided constants.</p> <p>AMPS also accepts a comma-delimited list of bookmarks. In this case, AMPS begins the subscription from whichever of the bookmarks is earliest in the transaction log.</p>

Rate Controlled Bookmark Subscription

To create a bookmark subscription, set the `Bookmark` property on the command. The value of this property can be either a specific bookmark, a timestamp, or one of the client-provided constants. The *AMPS User Guide* provides details on creating timestamps. Notice that the `MOST_RECENT` constant tells the AMPS client to find the appropriate message in the client bookmark store and begin the subscription at that point. In this case, the client sends that bookmark value to AMPS. The `Bookmark` option is only supported for topics that are recorded in an AMPS transaction log.

Table 10.10. Bookmark Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Bookmark	<p>Sets the point in the transaction log at which the subscription will begin. The bookmark provided can be a specific AMPS bookmark, a timestamp, or one of the client-provided constants.</p> <p>AMPS also accepts a comma-delimited list of bookmarks. In this case, AMPS begins the subscription from whichever of the bookmarks is earliest in the transaction log.</p>
Options	A comma-separated list of options for the command. To control the rate at which AMPS delivers messages, the

Header	Comment
	options for the command must include a rate specifier. For example, to specify a limit of 750 messages per second, include <code>rate=750</code> in the options string.

Bookmark Subscription With Content Filter

To create a bookmark subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark, a timestamp, or one of the client-provided constants. The *AMPS User Guide* provides details on creating timestamps. Notice that the `MOST_RECENT` constant tells the AMPS client to find the appropriate message in the client bookmark store and begin the subscription at that point. In this case, the client sends that bookmark value to AMPS.

To add a filter to a bookmark subscription, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.11. Bookmark Subscription With Content Filter

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Bookmark	Sets the point in the transaction log at which the subscription will begin. The bookmark provided can be a specific AMPS bookmark, a timestamp, or one of the client-provided constants. AMPS also accepts a comma-delimited list of bookmarks. In this case, AMPS begins the subscription from whichever of the bookmarks is earliest in the transaction log.
Filter	Sets the content filter to be applied to the subscription. Only messages that match the content filter will be provided to the subscription.

Pausing a Bookmark Subscription

To pause a bookmark subscription, you must provide the subscription ID and the `pause` option on a subscribe command.

Table 10.12. Pause a Bookmark Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
SubId	A comma-delimited list of subscription IDs to pause.
Options	A comma-delimited list of options for the command. To pause a subscription, the options must include <code>pause</code> .

Resuming a Bookmark Subscription

To resume a bookmark subscription, you must provide the subscription ID and the `resume` option on a `subscribe` command.

Table 10.13. Resume a Bookmark Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
SubId	A comma-delimited list of subscription IDs to resume.
Options	A comma-delimited list of options for the command. To resume a subscription, the options must include <code>resume</code> .

Replacing the Filter on a Subscription

To replace the content filter on a subscription, provide the `SubId` of the subscription to be replaced, add the `replace` option, and set the `Filter` property on the command with the new filter. The *AMPS User Guide* provides details on the filter syntax.

Table 10.14. Replacing the Filter on a Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
SubId	The identifier for the subscription to update. The <code>SubId</code> is the <code>CommandId</code> for the original <code>subscribe</code> command.
Options	A comma-separated list of options. To replace the filter on a subscription, include <code>replace</code> in the list of options.
Filter	Sets the content filter to be applied to the subscription. Only messages that match the content filter will be provided to the subscription.

Subscribing to a Queue and Requesting a `max_backlog`

To subscribe to a queue and request a `max_backlog` greater than 1, use the `Options` field of the `subscribe` command to set the requested `max_backlog`.

Table 10.15. Requesting a `max_backlog`

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Options	A comma-separated list of options. To request a value for the <code>max_backlog</code> , pass the value in the options as follows:

Header	Comment
	<code>max_backlog=NN</code>
	For example, to request a max backlog of 7, your application would pass the following option:
	<code>max_backlog=7</code>

SOW Query

This section presents common recipes for querying a SOW topic in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic SOW Query

In its simplest form, a SOW query needs only the topic to query.

Table 10.16. Basic SOW Query

Header	Comment
Topic	Sets the topic to query. The SOW query returns all messages in the SOW. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

Basic SOW With Options

In its simplest form, a SOW needs only the topic to subscribe to. To add options to the subscription, set the `Options` header on the `Command`.

Table 10.17. Basic SOW Query with Options

Header	Comment
Topic	Sets the topic to query. The SOW query returns all messages in the SOW. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Options	A comma-delimited set of options for this command. See the <code>AMPS Command Reference</code> for a description of supported options.

SOW Query With Ordered Results

In its simplest form, a SOW needs only the topic to subscribe to. To return the results in a specific order, provide an ordering expression in the `OrderBy` header.

Table 10.18. Basic SOW Query with Ordered Results

Header	Comment
Topic	Sets the topic to query. The SOW query returns all messages in the SOW. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
OrderBy	<p>Orders the results returned as specified. Requires a comma-separated list of identifiers of the form:</p> <pre>/field [ASC DESC]</pre> <p>For example, to sort in descending order by <code>orderDate</code> so that the most recent orders are first, and ascending order by <code>customerName</code> for orders with the same date, you might use a specifier such as:</p> <pre>/orderDate DESC, /customerName ASC</pre> <p>If no sort order is specified for an identifier, AMPS defaults to ascending order.</p>

SOW Query With TopN Results

In its simplest form, a SOW needs only the topic to subscribe to. To return only a specific number of records, provide the number of records to return in the TopN header.

Table 10.19. SOW Query with TopN Results

Header	Comment
Topic	Sets the topic to query. The SOW query returns all messages in the SOW. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
TopN	<p>The maximum number of records to return. AMPS uses the OrderBy header to determine the order of the records.</p> <p>If no OrderBy header is provided, records are returned in an indeterminate order. In most cases, using an OrderBy header when you use the TopN header will guarantee that you get the records of interest.</p>
OrderBy	<p>Orders the results returned as specified. Requires a comma-separated list of identifiers of the form:</p> <pre>/field [ASC DESC]</pre> <p>For example, to sort in descending order by <code>orderDate</code> so that the most recent orders are first, and ascending order by <code>customerName</code> for orders with the same date, you might use a specifier such as:</p> <pre>/orderDate DESC, /customerName ASC</pre>

Header	Comment
	If no sort order is specified for an identifier, AMPS defaults to ascending order.

Content Filtered SOW Query

To provide a content filter on a SOW query, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.20. Content Filtered SOW Query Subscription

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be returned in response to the query.

Historical SOW Query

To create a historical SOW query, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps.

This command is only supported on SOW topics that have `History` enabled.

Table 10.21. Historical SOW Query

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.

Historical SOW Query With Content Filter

To create a historical SOW query, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps. To add a filter to the query, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

This command is only supported on SOW topics that have `History` enabled.

Table 10.22. Historical SOW Query With Content Filter

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

Header	Comment
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be provided to the query.

SOW Query for Specific Records

AMPS allows a consumer to query for specific records as identified by a set of `SowKeys`. For topics where AMPS assigns the `SowKey`, the `SowKey` for the record is the AMPS-assigned identifier. For topics configured to require a user-provided `SowKey`, the `SowKey` for the record is the original key provided when the record was published. The *AMPS User Guide* provides more details on SOW keys.

Table 10.23. SOW Query by SOW Key

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
SowKeys	A comma-delimited list of <code>SowKey</code> values. AMPS returns only the records specified in this list. For example, a valid format for a list of keys would be: 1853097931817257202,10402779940201650075,223638799

SOW Query with Pagination

AMPS allows a consumer to page through records in the SOW using the `top_n` and `skip_n` options. With this approach, the application uses the `top_n` option to limit the number of records returned to a single page worth of records. The application uses the `skip_n` option to set the number of records to skip ahead to get to the page to display, and sets the `OrderBy` header to specify the ordering for the records. For example, if 10 records fit on a page, and the pages are ordered by the `ClientName` field of the records, to display the fourth page, the application would set `top_n` to 10, `skip_n` to 30 (to skip the first three pages of records), and `OrderBy` to `/ClientName`.

Table 10.24. SOW Query by SOW Key

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
OrderBy	Orders the results returned as specified. Requires a comma-separated list of identifiers of the form: <code>/field [ASC DESC]</code> For example, to sort in descending order by <code>orderDate</code> so that the most recent orders are first, and ascending or-

Header	Comment
	<p>der by customerName for orders with the same date, you might use a specifier such as:</p> <pre>/orderDate DESC, /customerName ASC</pre> <p>If no sort order is specified for an identifier, AMPS defaults to ascending order.</p>
Options	<p>An options string that sets the top_n and skip_n values for this query. For example, to skip 100 records and return the next 10 records, use an options string such as:</p> <pre>top_n=10,skip_n=100</pre>

SOW and Subscribe

This section presents common recipes for atomic sow and subscribe in AMPS using the Command or Message interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic SOW and Subscribe

In its simplest form, a SOW and Subscribe needs only the topic to subscribe to.

Table 10.25. Basic SOW and Subscribe

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

SOW and Subscribe With Options

In its simplest form, a SOW and subscribe command needs only the topic to subscribe to. To add options to the subscription, set the Options header on the Command.

Table 10.26. Basic SOW and Subscribe with Options

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Options	<p>A comma-delimited set of options for this command. See the AMPS Command Reference for a full description of supported options.</p> <p>The most common options for this command are:</p> <pre>oof</pre> <p>Request out of order notifications</p>

Header	Comment
	timestamp Include timestamps on messages

Content Filtered SOW and Subscribe

To provide a content filter on a SOW and Subscribe, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.27. Content Filtered SOW and Subscribe

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be returned in response to the query.

Conflated SOW and Subscribe

To request conflation on the subscription for a SOW and subscribe, set the `Options` property on the command to specify the conflation interval. When the topic has a SOW (including a view or a conflated topic), there is no need to provide a `conflation_key`.

Table 10.28. Conflated SOW and Subscribe

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Options	<p>A comma-separated list of options for the command. Set the conflation interval in the options.</p> <p>For example, to set a conflation interval of 250 milliseconds, provide the following option:</p> <pre>conflation=250ms</pre> <p>To set the conflation interval to 1 minute, provide an option of:</p> <pre>conflation=1m</pre>

Historical SOW and Subscribe

To create a historical SOW query with a subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps. This command is only supported on SOW topics that are recorded in an AMPS transaction log. If the `Bookmark` provided is a value other than `NOW (0|1|)`, the SOW topic must have `History` enabled.

Table 10.29. Historical SOW and Subscribe

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.

Historical SOW and Subscribe With Content Filter

To create a historical SOW query with a subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps. This command is only supported on SOW topics that are recorded in an AMPS transaction log. If the `Bookmark` provided is a value other than `NOW (0|1|)`, the SOW topic must have `History` enabled.

Table 10.30. Historical SOW and Subscribe With Content Filter

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be provided to the query.

Delta Publishing

This section presents common recipes for publishing to a topic in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic Delta Publish

In its simplest form, a subscription needs only the topic to publish to and the data to publish. The AMPS client automatically constructs the necessary AMPS headers and formats the full `delta_publish` command.

In many cases, a publisher only needs to use the basic delta publish command.

Table 10.31. Basic Delta Publish

Header	Comment
Topic	Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.

Header	Comment
	Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.

Delta Publish With CorrelationId

AMPS provides publishers with a header field that can be used to contain arbitrary data, the `CorrelationId`. A delta publish message can be used to update the `CorrelationId` as well as the data within the message.

Table 10.32. Delta Publish With CorrelationId

Header	Comment
Topic	<p>Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.</p> <p>Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.</p>
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.
CorrelationId	<p>The <code>CorrelationId</code> to provide on the message. AMPS provides the <code>CorrelationId</code> to subscribers. The <code>CorrelationId</code> has no significance for AMPS.</p> <p>The <code>CorrelationId</code> may only contain characters that are valid in base-64 encoding.</p>

Delta Publish With Explicit SOW Key

When publishing to a SOW topic that is configured to require an explicit SOW key, the publisher needs to set the `SowKey` header on the message.

Table 10.33. Delta Publish with Explicit SOW Key

Header	Comment
Topic	<p>Sets the topic to publish to. The topic specified must be a literal topic name. Regular expression characters in the topic name are not interpreted.</p> <p>Some topics in AMPS, such as views and conflated topics, cannot be published to directly. Instead, a publisher must publish to the underlying topics.</p>

Header	Comment
Data	The data to publish to the topic. The AMPS client does not interpret, escape, or validate this data: the data is provided to the server verbatim.
SowKey	The SOW Key to use for this message. This header is only supported for publishes to a topic that requires an explicit SOW Key.

Delta Subscribing

This section presents common recipes for subscribing to a topic in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic Delta Subscription

In its simplest form, a delta subscription needs only the topic to subscribe to.

Table 10.34. Basic Delta Subscription

Header	Comment
Topic	Sets the topic to subscribe to. All messages from the topic will be delivered on this subscription. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

Basic Delta Subscription With Options

In its simplest form, a subscription needs only the topic to subscribe to. To add options to the subscription, set the `Options` header on the `Command`.

Table 10.35. Basic Delta Subscription

Header	Comment
Topic	Sets the topic to subscribe to. All messages from the topic will be delivered on this subscription. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Options	A comma-delimited set of options for this command. See the <i>AMPS Command Reference</i> for a description of supported options.

Content Filtered Delta Subscription

To provide a content filter on a subscription, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.36. Content Filtered Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Filter	Sets the content filter to be applied to the subscription. Only messages that match the content filter will be provided to the subscription.

Bookmark Delta Subscription

To create a bookmark subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark, a timestamp, or one of the client-provided constants. The *AMPS User Guide* provides details on creating timestamps. Notice that the `MOST_RECENT` constant tells the AMPS client to find the appropriate message in the client bookmark store and begin the subscription at that point. In this case, the client sends that bookmark value to AMPS.

Table 10.37. Bookmark Subscription

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.
Bookmark	<p>Sets the point in the transaction log at which the subscription will begin. The bookmark provided can be a specific AMPS bookmark, a timestamp, or one of the client-provided constants.</p> <p>AMPS also accepts a comma-delimited list of bookmarks. In this case, AMPS begins the subscription from whichever of the bookmarks is earliest in the transaction log.</p>

Bookmark Delta Subscription With Content Filter

To create a bookmark subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark, a timestamp, or one of the client-provided constants. The *AMPS User Guide* provides details on creating timestamps. Notice that the `MOST_RECENT` constant tells the AMPS client to find the appropriate message in the client bookmark store and begin the subscription at that point. In this case, the client sends that bookmark value to AMPS.

To add a filter to a bookmark subscription, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.38. Bookmark Delta Subscription With Content Filter

Header	Comment
Topic	Sets the topic to subscribe to. The topic provided can be either the exact name of the topic, or a regular expression that matches the names of the topics for the subscription.

Header	Comment
Bookmark	<p>Sets the point in the transaction log at which the subscription will begin. The bookmark provided can be a specific AMPS bookmark, a timestamp, or one of the client-provided constants.</p> <p>AMPS also accepts a comma-delimited list of bookmarks. In this case, AMPS begins the subscription from whichever of the bookmarks is earliest in the transaction log.</p>
Filter	Sets the content filter to be applied to the subscription. Only messages that match the content filter will be provided to the subscription.

SOW and Delta Subscribe

This section presents common recipes for atomic sow and delta subscribe in AMPS using the `Command` or `Message` interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Basic SOW and Delta Subscribe

In its simplest form, a SOW and Delta Subscribe needs only the topic to subscribe to.

Table 10.39. Basic SOW Query

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

SOW and Delta Subscribe With Options

In its simplest form, a SOW and subscribe command needs only the topic to subscribe to. To add options to the subscription, set the `Options` header on the `Command`.

Table 10.40. Basic SOW and Delta Subscribe with Options

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Options	<p>A comma-delimited set of options for this command. See the AMPS <code>Command Reference</code> for a full description of supported options.</p> <p>The most common options for this command are:</p> <p><code>oof</code> Request out of order notifications</p>

Header	Comment
	timestamp Include timestamps on messages

Content Filtered SOW and Delta Subscribe

To provide a content filter on a SOW and Delta Subscribe, set the `Filter` property on the command. The *AMPS User Guide* provides details on the filter syntax.

Table 10.41. Content Filtered SOW and Delta Subscribe

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be returned in response to the query.

Historical SOW and Subscribe

To create a historical SOW query with a subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps. This command is only supported on SOW topics that are recorded in an AMPS transaction log. If the `Bookmark` provided is a value other than `NOW (0 | 1 |)`, the SOW topic must have `History` enabled.

Table 10.42. Historical SOW and Subscribe

Header	Comment
Topic	Sets the topic to query and subscribe to. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.

Historical SOW and Delta Subscribe With Content Filter

To create a historical SOW query with a subscription, set the `Bookmark` property on the command. The property can be either a specific bookmark or a timestamp. The *AMPS User Guide* provides details on creating timestamps. This command is only supported on SOW topics that are recorded in an AMPS transaction log. If the `Bookmark` provided is a value other than `NOW (0 | 1 |)`, the SOW topic must have `History` enabled.

Table 10.43. Historical SOW and Delta Subscribe With Content Filter

Header	Comment
Topic	Sets the topic to query. The topic specified can be the literal topic name, or a regular expression that matches multiple topics.

Header	Comment
Bookmark	Sets the historical point in the SOW at which to query. The query returns the saved state of the records in the SOW as of the point in time specified in this header.
Filter	Sets the content filter to be applied to the query. Only messages that match the content filter will be provided to the query.

SOW Delete

This section presents common recipes for sending a `sow_delete` command using the Command or Message interfaces. This section provides information on how to configure the request to AMPS. You can adapt this information to your application and the specific interface you are using.

Delete All Records in a SOW

To delete all records in a SOW, provide a filter that evaluates to TRUE for every record in the SOW. By convention, 60East recommends `1=1` for the filter.

Table 10.44. Delete All Records in a SOW

Header	Comment
Topic	Sets the topic from which to remove records.
Filter	A filter specifying the messages to remove. By convention, use <code>1=1</code> to remove all records in the SOW.

Delete SOW Records Matching a Filter

To delete the records that match a particular filter, provide the filter in the `sow_delete` command.

Table 10.45. Delete All Records in a SOW

Header	Comment
Topic	Sets the topic from which to remove records.
Filter	A filter specifying the messages to remove.

Delete A Specific Message By Data

To delete a specific message, provide the data for the message to delete. With this form of SOW delete, AMPS deletes the message that would have been updated if the data were provided as a publish message. Notice that this form of `sow_delete` relies on the Key definition in the SOW configuration, and is not generally useful with explicitly-keyed SOW topics.

Table 10.46. Delete All Records in a SOW

Header	Comment
Topic	Sets the topic from which to remove records.

Header	Comment
Data	The message to remove.

Deleting Specific Messages Using Keys

To delete specific messages using SOW keys, provide the SOW keys for the message to delete.

Table 10.47. Delete All Records in a SOW

Header	Comment
Topic	Sets the topic from which to remove records.
SOWKeys	A comma-delimited list of SOWKeys that specify the messages to remove.

Acknowledging Messages from a Queue

To acknowledge messages from an AMPS queue, provide the bookmarks for the messages to acknowledge. Notice that this is the only form of the `sow_delete` command that can acknowledge messages from a queue, and that this form of `sow_delete` is not accepted for topics that are not queue topics.

Table 10.48. Acknowledging a queue message

Header	Comment
Topic	Sets the topic that contains the messages to acknowledge.
Bookmark	A comma-delimited list of Bookmarks that specify the messages to acknowledge.

Chapter 11. Utilities

The AMPS Java client includes a set of utilities and helper classes to make working with AMPS easier.

11.1. Composite Message Types

The client provides a pair of classes for creating and parsing composite message types.

- `CompositeMessageBuilder` allows you to assemble the parts of a composite message and then serialize them in a format suitable for AMPS.
- `CompositeMessageParser` extracts the individual parts of a composite message type

For more information regarding composite message types, refer to Chapter 4.3.

Building Composite Messages

To build a composite message, create an instance of `CompositeMessageBuilder`, and populate the parts. The `CompositeMessageBuilder` copies the parts provided, in order, to the underlying message. The builder simply writes to an internal buffer with the appropriate formatting, and does not allow you to update or change the individual parts of a message once they've been added to the builder.

The snippet below shows how to build a composite message that includes a JSON part, constructed as a string, and a binary part consisting of the bytes from an `ArrayList`.

```
StringBuilder sb = new StringBuilder();
sb.append("{\"data\":\"sample\"}");

List<Double> theData = new ArrayList<Double>();
// populate theData
...

// Create a byte array from the data: this is
// what the program will send.
ByteArrayOutputStream outBytes = new ByteArrayOutputStream();
ObjectOutputStream listWriter = new ObjectOutputStream(outBytes);
listWriter.writeObject(theData);

// Create the payload for the composite message.
CompositeMessageBuilder builder;

// Construct the composite
CompositeMessageBuilder builder = new CompositeMessageBuilder();
builder.append(sb.toString());
builder.append(outBytes.toByteArray(), 0, outBytes.size());

// send the message
```

```
String topic = "messages";

Field outMessage = new Field();
builder.setField(outMessage);

client.publish(topic.getBytes(), 0, topic.getBytes().length,
              outMessage.buffer, 0, outMessage.length);
```

Parsing Composite Messages

To parse a composite message, create an instance of `CompositeMessageParser`, then use the `parse()` method to parse the message provided by the AMPS client. The `CompositeMessageParser` gives you access to each part of the message as a sequence of bytes.

For example, the following snippet parses and prints messages that contain a JSON part and a binary part that contains an array of doubles.

```
try (MessageStream stream = client.subscribe("messages"))
{
    for(Message message : stream)
    {
        int parts = parser.parse(message);

        String json = parser.getString(0);
        Field binary = new Field();
        parser.getField(1, binary);

        ByteArrayInputStream inBytes =
            new ByteArrayInputStream(binary.buffer, binary.position,
binary.length);
        ObjectInputStream listReader =
            new ObjectInputStream(inBytes);
        List<Double> theData = (List<Double>)listReader.readObject();

        System.out.println("Received message with " + parts + " parts.");
        System.out.println(json);
        for (Double d : theData)
        {
            System.out.print(d + " ");
        }
        System.out.print("\n");
    }
}
```

Notice that the receiving application is written with explicit knowledge of the structure and content of the composite message type.

11.2. NVFIX Messages

The client provides a pair of classes for creating and parsing NVFIX messages.

- `NVFIXBuilder` allows you to assemble a NVFIX message and then serialize it in a format suitable for AMPS.
- `NVFIXShredder` extracts the individual fields of a NVFIX message.

Building NVFIX Messages

To build a NVFIX message, create an instance of `NVFIXBuilder`, then add the fields of the message using `append()`. `NVFIXBuilder` copies the fields provided, in order, to the underlying message. The builder simply writes to an internal buffer with the appropriate formatting, and does not allow you to update or change the individual fields of a message once they've been added to the builder.

The snippet below shows how to build an NVFIX message and publish it to the AMPS client.

```
// Build the message payload

// Create a builder with 1024 bytes of initial capacity,
// using the default 0x01 delimiter.
NVFIXBuilder builder = new NVFIXBuilder(1024, (byte)1);

// Add fields
builder.append("test-field", "24");
builder.append("another", "Here's another field");
builder.append("data", "1234567890");

// Create a string for the topic
String topic = "test-topic";

client.connect(uri_);
System.out.println("connected..");
client.logon();

// publish the message, using the overload that takes a byte array and
// length for the topic and payload.

client.publish(topic.getBytes(), 0, topic.length(), builder.getBytes(), 0,
    builder.getSize());
```

Parsing NVFIX Messages

To parse a NVFIX message, create an instance of `NVFIXShredder`, then use the `toNVMap()` method to parse the message provided by the AMPS client. The `NVFIXShredder` gives you access to each field of the message in a map.

The snippet below shows how to parse and print an NVFIX message.

```
Client client = new Client("ConsoleSubscriber");

try
{
    // connect to the AMPS server and logon
    client.connect(uri_);
```

```
// subscribe to the test-topic topic.
// when a message arrives, print the message.

MessageStream ms = client.subscribe("test-topic");
try
{
    // Create a shredder -- since this just returns
    // the Map, we can reuse the same shredder.
    NVFIXShredder shredder = new NVFIXShredder((byte)1);

    for (Message m : ms)
    {
        // Skip messages with no data.
        if (m.getCommand() != Message.Command.SOW &&
            m.getCommand() != Message.Command.Publish) continue;

        System.out.println("Got a message");

        // Shred the message into a Map
        Map<CharSequence,CharSequence> fields = shredder.toNVMap(m.getData());

        // Iterate over the keys in the map and print the key and data
        for (CharSequence key : fields.keySet())
        {
            System.out.println("  " +key + "=" + fields.get(key));
        }
    }
}
finally // close the message stream to release the subscription
{ ms.close(); }
```

11.3. FIX Messages

The client provides a pair of classes for creating and parsing FIX messages.

- `FIXBuilder` allows you to assemble a FIX message and then serialize it in a format suitable for AMPS.
- `FIXShredder` extracts the individual fields of a FIX message.

Building FIX Messages

To build a FIX message, create an instance of `FIXBuilder`, then add the fields of the message using `append()`. `FIXBuilder` copies the fields provided, in order, to the underlying message. The builder simply writes to an internal buffer with the appropriate formatting, and does not allow you to update or change the individual fields of a message once they've been added to the builder.

The snippet below shows how to build a FIX message and publish it to the AMPS client.

```
// Build the message payload
```

```
// Create a builder with 1024 bytes of initial capacity,  
// using the default 0x01 delimiter.  
FIXBuilder builder = new FIXBuilder(1024, (byte)1);  
  
// Add fields  
builder.append("test-field", "24");  
builder.append("another", "Here's another field");  
builder.append("data", "1234567890");  
  
// Create a string for the topic  
String topic = "test-topic";  
  
client.connect(uri_);  
System.out.println("connected..");  
client.logon();  
  
// publish the message, using the overload that takes a byte array and  
// length for the topic and payload.  
  
client.publish(topic.getBytes(), 0, topic.length(), builder.getBytes(), 0,  
builder.getSize());
```

Parsing FIX Messages

To parse a FIX message, create an instance of `FIXShredder`, then use the `toMap()` method to parse the message provided by the AMPS client. The `FIXShredder` gives you access to each field of the message in a map.

The snippet below shows how to parse and print a FIX message.

```
Client client = new Client("ConsoleSubscriber");  
  
try  
{  
    // connect to the AMPS server and logon  
    client.connect(uri_);  
  
    // subscribe to the test-topic topic.  
    // when a message arrives, print the message.  
  
    MessageStream ms = client.subscribe("test-topic");  
    try  
    {  
        // Create a shredder -- since this just returns  
        // the Map, we can reuse the same shredder.  
        FIXShredder shredder = new FIXShredder((byte)1);  
  
        for (Message m : ms)  
        {  
            // Skip messages with no data.  
            if (m.getCommand() != Message.Command.SOW &&  
                m.getCommand() != Message.Command.Publish) continue;
```

```
System.out.println("Got a message");

// Shred the message into a Map
Map<Integer,CharSequence> fields = shredder.toMap(m.getData());

// Iterate over the keys in the map and print the key and data
for (Integer key : fields.keySet())
{
    System.out.println(" " +key + "=" + fields.get(key));
}
}
}
```

Chapter 12. Advanced Topics

12.1. C# Client Compatibility

AMPS clients are available for many languages. Many AMPS customers write clients using a variety of languages, often both Java and C#. While Java and C# are fundamentally different languages, they share enough syntax that it can be straightforward to port code between the two, and especially from Java to C#.

To aid in conversion from Java to C# (and from C# to Java), the C# client has a number of features that make it a little easier to bring code from Java to C#:

-
- C# Parameters that take lambda functions also take an interface type. The AMPS Java client defines interfaces such as `ClientMessageHandler`, which your application implements with a single `invoke()` method that is called when an event occurs. In C#, the AMPS client uses lambda functions and delegates to provide equivalent functionality. However, the same `*Handler` interfaces exist in C#, and instead of passing a lambda function, you may also implement these interfaces and pass in derived classes. While doing so would be inconvenient in C#, providing this symmetry allows your Java and C# to be ported interchangeably.
- Java-style method name conventions are used throughout AMPS. In .NET, method names often begin with a capitalized first letter (e.g. `Connect()` instead of `connect()`). However, the C# AMPS client retains the capitalization style of the Java client where possible, making porting straightforward.

12.2. Rebuilding the Client

In the rare occasion that you need to make customizations to the AMPS Java Client, the installation packages include the full source code for the AMPS client. Rebuilding the client is straightforward using the provided `build.xml`, and the Apache `ant` program. The `build.xml` file is located in the `api/client/java/` directory of the root of your AMPS installation. For example, if AMPS is installed in `/opt/AMPS/`, then the Java client's `build.xml` would be located in the `/opt/AMPS/api/client/java/` directory.

Assuming that a JDK version 1.6 or greater is installed and the Apache `ant` package has been installed, rebuilding the client is as simple as typing:

```
ant
```

from the command line in the same directory where the `build.xml` file is located in the AMPS install. Upon successful completion, the libraries will be located the `api/client/java/dist/lib` directory.

12.3. Transport Filtering

The AMPS Java client offers the ability to filter incoming and outgoing messages in the format they are sent and received on the network. This allows you to inspect or modify outgoing messages before they are sent to the network, and incoming messages as they arrive from the network. To create a transport filter, you implement the interface `TransportFilter`, construct an instance of the filter class, and install the filter with the `setTransportFilter` method on the transport.

The AMPS Java client does not validate any changes made by the transport filter. This interface is most useful for application debugging or transport development.

The client includes a sample filter, `TransportTraceFilter`, that simply writes incoming and outgoing buffers to an `OutputStream`.

12.4. Working with Messages & Byte Buffers

The AMPS Java client allows you to publish messages that contain data from byte buffers. When working with byte buffers in AMPS, it's best to follow these simple conventions.

AMPS provides overloaded `publish()` methods that allow you to publish messages from various formats. In this case, to publish a message using byte buffers, the message data must be provided as a `byte[]`. Along with the message data, the message topic, to which the message will be published to, must also be provided as a `byte[]`.

The example below shows how to serialize an object into a byte buffer, then publish the message to AMPS using the `publish()` method.

```
...
// create the topic string
String topic = "messages";

// create the object data
Employee emp = new Employee();

// create the field for the payload of the message
Field data = new Field();

try
{
    // create streams used to serialize the object and
    ByteArrayOutputStream bStream = new ByteArrayOutputStream();
    ObjectOutputStream oStream = new ObjectOutputStream(bStream);

    // serialize the object to the ObjectOutputStream
    oStream.writeObject(emp);

    // set the byte of the message to a field
    data.set(bStream.toByteArray());
}
catch (Exception e)
{
    e.printStackTrace();
}

//publish to the "messages" topic using byte buffers
client.publish(topic.getBytes(), 0, topic.getBytes().length, data.buffer, 0,
    data.buffer.length);
...
```


In addition to publishing messages, AMPS allows you to access the raw bytes in the data part of a message. The method `getDataRaw()` returns a `Field` that is composed of a byte buffer, position of the data in the buffer, and the length of the data. This data can then be deserialized and converted to an object for further use.

The example below shows how to access the raw bytes of a message, and then shows how to deserialize the bytes of that message to a Java object.

```
...  
  
// access the message raw data  
Field data = message.getDataRaw();  
  
// construct a ByteArrayInputStream object,  
// an ObjectInputStream object and a Field  
// to deserialize the data  
ByteArrayInputStream bais;  
ObjectInputStream ois;  
Field data;  
  
// construct the payload object  
Employee emp;  
  
ByteArrayInputStream  
try  
{  
    // deserialize the data using ObjectInputStream to an object  
    bais = new ByteArrayInputStream(data.buffer, data.position,  
    data.buffer.length);  
    ois = new ObjectInputStream(bais);  
    emp = (Employee) ois.readObject();  
}  
catch(Exception e)  
{  
  
...  
}
```

12.5. Providing SSL Certificates to the AMPS Java Client

The AMPS Java client includes support for Secure Sockets Layer (SSL) connections to AMPS. The client automatically attempts to make an SSL connection when the transport in the connection string is set to `tcps`, as described in Section 3.3.

There is no other change required to application code to use an SSL connection. However, to successfully make a SSL connection to AMPS, the Java runtime requires:

- A key to provide for the client connection, as stored in the keystore.
- A trust store to use to validate the server certificate used for the connection.

Most often, this means that the server certificate for the connection must be signed by a trusted certificate, and that trusted certificate must be in either the default trust store, or a trust store provided to the JVM at runtime.

This can also be accomplished by importing the server certificate into the default trust store or a trust store set by the application at run time.

The parameters that you need to provide depend on how you have configured your certificates. If the server certificate is in the default trust store, you only need to provide a key for the client connection. Otherwise, you need to provide both a client key and a trust store that can validate the server certificate.

There are three common methods of providing these certificates:

- Setting global parameters on the command line
- Setting global system properties
- Creating an SSL context object for a specific connection

These options are discussed in more detail below. In general, the first two methods are easiest for simple testing, and allow you to easily set the same options for all connections from the application. The third method is the most flexible, and lets an application use different options for different connections.

Java Virtual Machine implementations may differ somewhat in the implementation details for `javax.net.ssl`. For full details on providing certificates to the Java runtime, see the documentation for your JVM implementation. For example, documentation on the Oracle JVM implementation of `javax.net` is available at Java Secure Socket Extension (JSSE) Reference Guide [<https://docs.oracle.com/javase/8/docs/technotes/guides/security/jsse/JSSERefGuide.html>], and this section focuses on the Oracle JVM.

Setting Global Keystore and Truststore on the Command Line

For example, if you have a key installed into a keystore file and the server certificate is present in the default truststore or signed by a certificate in the default truststore, you might provide the client key to an application using the following commandline flags:

```
-Djavax.net.ssl.keyStore=<path to keystore file> \  
-Djavax.net.ssl.keyStorePassword=<password>
```

If the server certificate is in a different truststore, you might provide the following set of flags to specify both the keystore and the truststore to use:

```
-Djavax.net.ssl.trustStore=<path to truststore> \  
-Djavax.net.ssl.trustStorePassword=<truststore password> \  
-Djavax.net.ssl.keyStore=<path to keystore file> \  
-Djavax.net.ssl.keyStorePassword=<keystore password>
```

When you provide passwords on the command line, it is possible for other users on the system to see the password. This method is most often used during development and debugging.

Setting Global Keystore and Truststore Within the Application

Within an application, you can use the `System.setProperty` method to set flags before creating a `Client` or `HAClient`. For example, if you have a key installed into a keystore file and the server certificate is present in the default truststore, you might set these properties when your application starts.

```
System.setProperty("javax.net.ssl.keyStore","path to keystore file");
System.setProperty("javax.net.ssl.keyStorePassword","keystore password");
```

This is equivalent to the first set of command line arguments above. As with the command line arguments above, you can use this method to set the truststore arguments as necessary.

```
System.setProperty("javax.net.ssl.keyStore","path to keystore file");
System.setProperty("javax.net.ssl.keyStorePassword","keystore password");
System.setProperty("javax.net.ssl.trustStore","path to truststore file");
System.setProperty("javax.net.ssl.trustStorePassword","truststore password");
```

Creating an SSL Context

Setting options on the command line or with the `System.setProperty` method affects every connection made from the application. You can also create your own `SSLContext` for the application to use for the `tcps` transport. This method provides the most flexibility, and allows you to use different information for different connections.

The snippet below has the same end result as the first command line above for the AMPS connection (that is, it sets a keystore and uses the default trust store). However, because this `SSLContext` will only be used by the AMPS client, you can use a keystore and certificates that are only used for AMPS connections, rather than for the application as a whole. This may be required if the application needs to use SSL to connect to a different external system, and the keys or trust store used for AMPS are different than the keys or trust store for the other system.

```
// Snippet showing how to provide a keystore to the
// TCPSTransport for use of the AMPS Java client.

// Create a keystore. Assume that the keystore
// is in JKS format.

KeyStore ks = KeyStore.getInstance("JKS");

String password = ... ; // Obtain password for keystore

FileInputStream fis = null;
try
{
    fis = new java.io.FileInputStream("<path to keystore file>");
    ks.load(fis, password.toCharArray());
}
finally
{
    if (fis !=null ) fis.close();
}

// Get the key manager factory, using the default
// algorithm.
KeyManagerFactory kmf =
    KeyManagerFactory.getInstance(KeyManagerFactory.getDefaultAlgorithm());

// Initialize the factory with the keystore.
kmf.init(ks, password.toCharArray());
```

```
// Get the SSL context
SSLContext context = SSLContext.getInstance("TLS");

// Use the key manager just constructed, with defaults
// for the trust manager and randomness source.

context.init(kmf.getKeyManagers(), null, null);

// Set the SSLContext for the TCPS transport
// to the context just set up with the keystore.
TCPSTransport.setDefaultSSLContext(context);
```

Troubleshooting SSL Connectivity to AMPS

To troubleshoot SSL connectivity, run the application with the following flag:

```
-Djavax.net.debug=all
```

With this flag set, the `javax.net` library will produce detailed information about the connection, including extensive information designed to help identify and resolve certificate problems. In particular, this flag will produce information on the certificates that the application is using, whether the application can successfully load those certificates, and whether the application can validate the certificate provided by the AMPS server.

Chapter 13. Performance Tips and Best Practices

This chapter presents tips and techniques for writing high-performance applications with AMPS. This section presents principles and approaches that describe how to use the features of AMPS and the AMPS client libraries to achieve high performance and reliability.

Specific techniques (for example, the details on how to write a message handler) are described in other parts of the AMPS documentation and referenced here. Other techniques require information specific to the application (for example, determining the minimum set of information required in a message), and are best done as part of your application design.

All of the recommendations in this section are general guidelines. There are few, if any, universal rules for performance: at times, a design decision that is absolutely necessary to meet the requirements for an application might reduce performance somewhat. For example, your application might involve sending large binary data that cannot be incrementally updated. That application will use more bandwidth per message than an application that sends 100-byte messages with fields that can be incrementally updated. However, since the application depends on being able to deliver the binary payloads, this difference in bandwidth consumption is a part of the requirements for the application, not a design decision that can be optimized.

13.1. Measure Performance and Set Goals

The most important tools for creating high performance applications that use AMPS are clear goals and accurate measurement. Without accurate measurement, it's impossible to know whether a particular change has improved performance or not. Without clear goals, it's difficult to know whether a given result is sufficient, or whether you need to continue improving performance.

60East recommends that your measurements include baseline metrics for the part of your message processing that does not involve AMPS. As an example, imagine your task is to reduce the amount of time that elapses between when an order is sent and when the processed response is received from 100ms in total to 85ms in total. To achieve this reduction, you might first measure the processing that your application performs on the order. If that processing consumes 65ms, the most effective optimization may be to improve the order processing. On the other hand, if processing an order consumes 15ms, then optimizing message delivery or network utilization may be the most effective way to meet your goals.

When measuring performance, simulate your production environment as closely as possible. For example, AMPS is highly parallelized, so sending a pattern of subscriptions and publishes from a single test client that would normally come from 20 clients will produce a very different performance profile. Likewise, AMPS can typically perform at rates that fill the available bandwidth. Performance measured on a 1GbE connection may be very different than performance measured over a 10GbE connection. Consider the characteristics of your data, and the number of messages you expect to store and process. A 1GB data set consisting of 1 million records will perform differently than a 1GB data set consisting of 10 million records, or a 1GB data set consisting of 100 records.

When collecting information about performance, 60East recommends enabling persistence for the Statistics Database (`stats.db`), so you can easily collect historical data on both AMPS and the operating system. For example, a dip in performance correlated with high CPU and memory usage at the same time each day may be correlated with other activity on the system (such as cron jobs or close of business processing). In a situation like that, where the performance reduction is based on factors external to the AMPS application, the overall system metrics captured in `stats.db` can help you re-create the external state and understand the state of the system as a whole. AMPS collects

the statistics in memory by default, and persisting that data into a database does not typically have a measurable effect on performance itself, but makes measuring and tuning performance much easier.

For performance testing, 60East recommends using dedicated hardware for AMPS to eliminate the effects of other processes. If dedicated hardware is not available and other processes are consuming resources, 60East recommends disabling AMPS NUMA tuning to ensure that AMPS threads do not unnecessarily compete with other processes during performance tuning.

13.2. Simplify Message Format and Contents

AMPS supports a wide range of message types, and is capable of filtering and processing large and complex messages. For many applications, the simplicity of being able to use messages that contain the full information is the most important consideration. For other applications, however, achieving the minimum possible latency and the maximum possible network utilization is important enough to warrant choosing a simplified message format.

To simplify message contents, carefully consider the information that downstream processors require. If a downstream process will not use information in the message, there is no need to send the information. For example, consider an application that provides orders from a UI. In such an application, the object that represents the order often contains information relevant to the local state of the application that is not relevant to a downstream system. Rather than simply serializing the full object, your application may perform better if you serialize only the fields that a downstream system will take action on.

To simplify message format, choose the simplest format that can convey the information that your application needs. The general principle is that the simpler the message format is, the more quickly AMPS and client libraries can parse messages of that type. Likewise, the more complicated the structure of each message is, the more work is required to parse the message. For the highest levels of performance, 60East recommends keeping the message structure simple and preferring message formats such as NVFIX, BFlat, or JSON as compared with more complicated formats such as XML or BSON.

13.3. Use Content Filtering Where Possible

AMPS content filtering helps your application perform better by ensuring that your application only receives the messages that it needs. Wherever possible, we recommend using content filtering to precisely specify which messages your application needs. In particular, if at any point your application is receiving a message, parsing the message, and then determining whether to act on the message or not, 60East recommends using content filters to ensure that your application only receives messages that it needs to act on.

13.4. Use Asynchronous Message Processing

The synchronous message processing interface is straightforward, and presents a convenient interface for getting started with AMPS.

However, the `MessageStream` used by the synchronous interface makes a full copy of each message and provides it from the background reader thread to the thread that consumes the message. This memory overhead and synchronization between the reader thread and consumer thread happens regardless of whether the application needs all of the header fields in the message or even processes the message. The `MessageStream` also does not take into account the speed at which your program is consuming messages, and will read messages into memory as fast as the network and processor allow. If your application cannot consume messages at wire speed, this can lead to increasing memory consumption as the application falls further behind the `MessageStream`.

Most applications see improved performance by using a `MessageHandler`. With this approach, the `MessageHandler` does minimal work. If more extensive processing is needed, the `MessageHandler` dispatches the work to another thread: but it does this only when the work is necessary, and it only saves the part of the message needed to accomplish the work.

13.5. Use Hash Indexes Where Possible

When querying a SOW, hash indexes on SOW topics are supported for exact matching on string data as described in the *AMPS User Guide*. A hash index can perform many times faster than a parallel query. If the query pattern for your application can take advantage of hash indexes, 60East recommends creating those hash indexes on your SOW topics.

13.6. Use a Failed Write Handler and Exception Listener

In many cases, particularly during the early stages of development, performance problems can point to defects in the application. Even after the application is tuned, monitoring for failure is important to keep applications running smoothly.

60East recommends always installing a failed write handler if your application is publishing messages. This will help you to quickly identify cases where AMPS is rejecting publishes due to entitlement failures, message type mismatches, or other similar problems.

60East recommends always installing an exception listener if your application is using asynchronous message processing. This will help you to identify and correct any problems with your message handler.

13.7. Reduce Bandwidth Requirements

In many applications that use AMPS, network bandwidth is the single most important factor in overall performance. Your application can use bandwidth most efficiently by reducing message size. For example, rather than serializing an entire object, you might serialize only the fields that the remote process needs to act on, as mentioned above. Likewise, rather than sending one message that contains a collected set of information that processors will need to extract, consider sending a message in the units that processors will work with. This can reduce bandwidth to processors substantially. For example, rather than sending a single message with all of the activity for a single customer over a given period of time (such as a trading day), consider breaking out the record into the individual transactions for the customer.

Tune Batch Size for SOW Queries

As described in Section 6.3, tuning the batch size for SOW queries can improve overall performance by improving network utilization. In addition, because the AMPS header is only parsed once per batch, a larger batch size can dramatically improve processing performance for smaller messages.

The AMPS clients default to a batch size of 10. This provides generally good performance for most transactional messages (such as order records or inventory records). For large messages, particularly messages greater than a megabyte in size, a batch size of 1 may reduce memory pressure in the client and improve performance.

With smaller messages (for example, message sizes of a few hundred bytes), 60East recommends measuring performance with larger batch sizes such as 50 or 100 . For large messages, reducing the batch size may improve overall performance by requiring less memory consumption on the AMPS server.

Conflate Fast-Changing Information

If your data source publishes information faster than your clients need to consume it, consider using a conflated topic. For example, in a system that presents a user interface and displays fast-moving data, it is common for the data to change at a rate faster than the user interface can format and render the data. In this case, a conflated topic can both reduce bandwidth and simplify processing in the user interface.

Minimize Bandwidth for Updates

If your application uses a SOW and processes frequent updates, consider using delta publish and delta subscribe to reduce the size of the messages transmitted. These features are designed to minimize bandwidth while still providing full-fidelity data streams.

Conflate Queue Acknowledgements

The AMPS clients include the ability to conflate acknowledgements back to AMPS as queue messages are processed. Using these features, with an appropriate `max_backlog`, can reduce the amount of network traffic required for acknowledgements.

Use a Transaction Log When Monitoring Publish Failures

When a topic is not covered by a transaction log, AMPS returns acknowledgment messages for every publish that requests one. This ensures that each message is acknowledged, even when AMPS has no persistent record of the messages in the topic. However, acknowledging each message requires more network traffic for each publish message.

When a topic is covered by a transaction log, AMPS conflates persisted acknowledgments. Conflation is possible in this case because AMPS has a full record of the messages and does not have to store additional state to conflate the acknowledgements. With conflated acknowledgements, AMPS will send a success acknowledgement periodically that covers all messages up to that point. If a message fails, AMPS immediately sends the conflated success acknowledgement for all previous messages and the failure acknowledgement for the failed message.

Combine Conflation and Deltas

In many cases, using an approach that combines delta publishes to a SOW with delta subscriptions to a conflated topic can dramatically reduce bandwidth to the application with no loss of information.

13.8. Limit Unnecessary Copies

One of the most effective ways to increase performance is to limit the amount of data copied within your application.

For example, if your message handler submits work to a set of processors that only use the `Data` and `Bookmark` from a `Message`, create a data structure that holds only those fields and copy that information into instances of that data structure rather than copying the entire `Message`. While this approach requires a few extra lines of code, the performance benefits can be substantial.

When publishing messages to AMPS, avoid unnecessary copies of the data. For example, if you have the data in a byte array, use the `publish` methods that use a byte array rather than converting the data to a string unnecessarily. Likewise, if you have the data in the form of a string, avoid converting it to a byte array where possible.

13.9. Manage Publish Stores

When using a publish store, the Client holds messages until they are acknowledged as persisted by AMPS, as determined by the replication configuration for the AMPS instance.

In the event that an instance with `sync` replication goes offline, the publish store for the Client will grow, since the messages are not being fully persisted. To avoid this problem, 60East recommends that an instance that uses `sync` replication always configure Actions to automatically downgrade the replication link if the remote instance goes offline for a period of time, and upgrade the link when the remote instance comes back online.

See the "High Availability and Replication" chapter in the *User Guide* for more information on replication, `sync` and `async` acknowledgement modes, and the Actions used to manage replication.

13.10. Work with 60East as Necessary

60East offers performance advice adapted for your specific usage through your support agreement. Once you've set your performance goals, worked through the general best practices and applied the practices that make sense for your application, 60East can help with detailed performance tuning, including recommendations that are specific to your use case and performance needs.

Appendix A. Exceptions

The following table details each of the exception types thrown by AMPS.

Table A.1. Exceptions supported in Client and HAClient

Exception	When	Notes
<code>AlreadyConnectedException</code>	Connecting	Thrown when <code>connect()</code> is called on a <code>Client</code> that is already connected.
<code>AMPSException</code>	Anytime	Base class for all AMPS exceptions.
<code>AuthenticationException</code>	Anytime	Indicates that an authentication failure occurred on the server.
<code>BadFilterException</code>	Subscribing	This typically indicates a syntax error in a filter expression.
<code>BadRegexTopicException</code>	Subscribing	Indicates that a malformed regular expression was found in the topic name.
<code>CommandException</code>	Anytime	Base class for all exceptions relating to commands sent to AMPS.
<code>ConnectionException</code>	Anytime	Base class for all exceptions relating to the state of the AMPS connection.
<code>ConnectionRefusedException</code>	Connecting	The connection was actively refused by the server. Validate that the server is running, that network connectivity is available, and the settings on the client match those on the server.
<code>DisconnectedException</code>	Anytime	No connection is available when AMPS needed to send data to the server <i>or</i> the user's disconnect handler threw an exception.
<code>InvalidTopicException</code>	SOW query	A SOW query was attempted on a topic not configured for SOW on the server.
<code>InvalidTransportOptionsException</code>	Connecting	An invalid option or option value was specified in the URI.
<code>InvalidURIException</code>	Connecting	The URI string provided to <code>connect()</code> was formatted improperly.
<code>MessageTypeException</code>	Connecting	The class for a given transport's message type was not found in AMPS.
<code>MessageTypeNotFoundException</code>	Connecting	The message type specified in the URI was not found in AMPS.
<code>NameInUseException</code>	Connecting	The client name (specified when instantiating <code>Client</code>) is already in use on the server.
<code>RetryOperationException</code>	Anytime	An error occurred that caused processing of the last command to be aborted. Try issuing the command again.
<code>StreamException</code>	Anytime	Indicates that data corruption has occurred on the connection between the client and server. This usually indicates an internal error inside of AMPS—please contact AMPS support.

Exceptions

Exception	When	Notes
SubscriptionAlreadyExistsException	Subscribing	A subscription has been requested using the same CommandId as another subscription. Create a unique CommandId for every subscription.
TimeoutException	Anytime	A timeout occurred waiting for a response to a command.
TransportTypeException	Connecting	Thrown when a transport type, unknown to AMPS is selected in the URI.
UnknownException	Anytime	Thrown when an internal error occurs. Contact AMPS support immediately.

Index

A

amps_client-javadoc.jar, 2
amps_client-sources.jar, 2
amps_client.jar, 2, 3, 4

B

Base class for exceptions, 16
bookmarkSubscribe() method, 43

C

certificate management, 77
Client Interface, 37
Client.setDisconnectHandler() method, 18
connect() method, 19
connectAndLogon() method, 38
connection parameters, 7
 tcp_keepalive, 8
 tcp_linger, 7
 tcp_nodelay, 7
 tcp_rcvbuf, 7
 tcp_sndbuf, 7
ConnectToNextUri() method, 20
createFileBacked(), 36
createFileBacked() method, 42
create_memory_backed(), 36

D

discard() method, 44
Disconnect handler, 19
disconnect() method, 41
DisconnectedException, 19
Disconnection, 18, 38

E

EPOCH parameter, 45
Exception base class, 16
ExceptionListener, 19
executeSOWQuery(), 24

F

failover, 20
File-based stores, 37

G

getTopic() method, 14

H

HAClient, 35

HAClient Client, 37
Heartbeat, 40
Hello world message, 5

I

Interfaces
 Message, 13, 75
 MessageHandler, 13, 14
invoke() method, 12

J

JAR files
 amps_client-javadoc.jar, 2
 amps_client-sources.jar, 2
 amps_client.jar, 2, 3, 4

M

Memory-based stores, 37
Message interface, 13,
MessageHandler
 implementation, 11
MessageHandler interface, 13, 14
Method
 bookmarkSubscribe(), 43
 connect(), 19
 connectAndLogon(), 38, 38
 ConnectToNextUri(), 20
 createFileBacked(), 42
 createMemoryBacked(), 36
 create_file_backed(), 36
 discard(), 44
 disconnect(), 38, 41
 executeSOWQuery(), 24
 getTopic(), 14
 invoke(), 12
 printStackTrace(), 16
 publish(), 40, 41
 setBookmarkStore(), 37
 setDisconnectHandler(), 18
 setHeartbeat(), 40
 setPublishStore(), 37
 sow(), 24
 sowAndSubscribe(), 25
 sowDelete(), 26
 subscribe(), 9, 12, 14, 15, 45
 toString(), 16
 unpersistedCount(), 40, 41
MOST_RECENT parameter, 45

N

NOW parameter, 45
null returns, 17

O

OOF, 26

P

Parameters

 EPOCH, 45

 MOST_RECENT, 45

 NOW, 45

printStackTrace() method, 16

publish around errors, 16

publish failures, 22

publish() method, 40, 41

R

Recompiling client, 75

ReconnectDelayStrategy, 38

reconnection delay, 38

reconnection timeout, 38

S

ServerChooser Interface, 39

setBookmarkStore(), 37

setDisconnectHandler() method, 18

setHeartbeat() method, 40

setPublishStore(), 37

sow(), 24

sowAndSubscribe() method, 25

sowDelete(), 26

SO_KEEPALIVE, 8

SO_LINGER, 7

SSL, 77

stderr, 16

Store interface, 40

SubId, 44

subscribe, 16

subscribe(), 9

subscribe() method, 12, 14, 15, 45

T

tcp_keepalive, 8

tcp_linger, 7

tcp_nodelay, 7, 7

tcp_rcvbuf, 7

tcp_sndbuf, 7, 7

TestPublisher.class, 4

toString() method, 16

U

unpersistedCount() method, 40, 41