

# Caplin Trader 1.4

---

## Integrating Caplin Trader With A Trading System

September 2009

# Contents

<b>1</b>	<b>Preface.....</b>	<b>1</b>
1.1	What this document contains.....	1
	About Caplin document formats .....	1
1.2	Who should read this document.....	1
1.3	Related documents.....	2
1.4	Typographical conventions.....	2
1.5	Feedback.....	3
1.6	Acknowledgments.....	3
<b>2</b>	<b>Overview.....</b>	<b>4</b>
<b>3</b>	<b>Trading concepts.....</b>	<b>6</b>
3.1	Trade Models.....	6
3.2	Trade Channels.....	6
3.3	Trades.....	6
3.4	Trade Events.....	7
3.5	Blotter Channels.....	7
<b>4</b>	<b>Example Trade Models.....</b>	<b>9</b>
4.1	Example Executable Streaming Price (ESP).....	9
4.2	Example Request for Stream (RFS).....	10
4.3	Example Order (ORD).....	11
4.4	Request for Quote (RFQ) with timeouts .....	12
<b>5</b>	<b>Configuring Trade Models.....</b>	<b>13</b>
5.1	RFQ example.....	14
<b>6</b>	<b>Using the Trading DataSource Java API.....</b>	<b>16</b>
6.1	Initialization.....	16
6.2	New Trade Channels.....	17
6.3	New Trades .....	17
6.4	Dealing with Events.....	18
6.5	Closing Trades.....	19
6.6	Closing Channels.....	19
6.7	Handling Blotter Channels.....	20
	Registering the BlotterTradeListener .....	20
	Implementing the BlotterTradeListener interface .....	22

<b>7</b>	<b>The Java Trading DataSource Example.....</b>	<b>24</b>
<b>8</b>	<b>Using the Trading DataSource C++ API.....</b>	<b>25</b>
8.1	Initialization.....	25
8.2	New Trade Channels.....	26
8.3	New Trades.....	26
8.4	Dealing with Events.....	27
8.5	Closing Trades.....	28
8.6	Closing Channels.....	28
8.7	Handling Blotter Channels.....	29
	Registering the BlotterListener .....	29
	Implementing the BlotterListener interface .....	30
<b>9</b>	<b>The C++ Trading DataSource Example.....</b>	<b>31</b>
<b>10</b>	<b>Configuring Caplin Liberator for trading.....</b>	<b>32</b>
10.1	Mapping trade messaging objects in Liberator.....	33
10.2	Routing trade messages to the Trading DataSource .....	35
10.3	Trading performance and integrity.....	36
	Optimizing client reconnection time .....	36
	Throttling .....	36
10.4	Security.....	38
	Session IDs .....	38
	Single sign-on and KeyMaster .....	38
<b>11</b>	<b>Appendix: Overview of throttling.....</b>	<b>39</b>
<b>12</b>	<b>Glossary of terms and acronyms.....</b>	<b>41</b>
	<b>Index.....</b>	<b>43</b>

# 1 Preface

## 1.1 What this document contains

This document describes how the Caplin Trading DataSource allows you to integrate Caplin Trader with your existing trading system.

Trading DataSources can be implemented in Java™ and C++. The document describes how to use both the Java and C++ APIs for this purpose.

### About Caplin document formats

This document is supplied in three formats:

- ◆ Portable document format (*.PDF* file), which you can read on-line using a suitable PDF reader such as Adobe Reader®. This version of the document is formatted as a printable manual; you can print it from the PDF reader.
- ◆ Web pages (*.HTML* files), which you can read on-line using a web browser. To read the web version of the document navigate to the *HTMLDoc\_m\_n* folder and open the file *index.html*.
- ◆ Microsoft HTML Help (*.CHM* file), which is an HTML format contained in a single file. To read a *.CHM* file just open it – no web browser is needed.

#### For the best reading experience

On the machine where your browser or PDF reader runs, install the following Microsoft Windows® fonts: Arial, Courier New, Times New Roman, Tahoma. You must have a suitable Microsoft license to use these fonts.

#### Restrictions on viewing .CHM files

You can only read *.CHM* files from Microsoft Windows.

Microsoft Windows security restrictions may prevent you from viewing the content of *.CHM* files that are located on network drives. To fix this either copy the file to a local hard drive on your PC (for example the Desktop), or ask your System Administrator to grant access to the file across the network. For more information see the Microsoft knowledge base article at <http://support.microsoft.com/kb/896054/>.

## 1.2 Who should read this document

This document is intended for Technical Managers, Enterprise Architects, and System Architects, who require an overview of the Caplin Trading DataSource and its Java and C++ APIs.

## 1.3 Related documents

### ◆ Caplin Trader Architecture

This document describes the architecture of Caplin Trader. It focuses on the use of the Caplin Platform in trading applications. It also identifies the areas in which the Platform can be integrated with your company's own and third-party systems.

### ◆ Caplin DataSource Overview

A technical overview of Caplin DataSource.

### ◆ Caplin Java Trading DataSource: API Documentation

This is the detailed Java API documentation for the Caplin Trading DataSource.

### ◆ Caplin C++ Trading DataSource: API Documentation

This is the detailed C++ API documentation for the Caplin Trading DataSource.

### ◆ Caplin Trader Trade Model Configuration XML Reference

This document defines the XML tags and attributes used to define Trade Models.

### ◆ Caplin Trader Client: API Specification

Documents the JavaScript libraries that allow developers to extend Caplin Trader Client by writing custom JavaScript code.

### ◆ Caplin KeyMaster Overview

KeyMaster integrates Caplin Liberator with an existing single sign-on system, so that end users do not have to explicitly log in to the Liberator server in addition to logging in to the enterprise's single sign-on server.

## 1.4 Typographical conventions

The following typographical conventions are used to identify particular elements within the text.

### Type

### Uses

**a-config-item**

Liberator configuration item

*a-config-parameter*

Parameter of a Liberator configuration item

*/AFolder/Afile.txt*

File names, folders and directories

```
Some code;
```

Program output and code examples

The `value=10` attribute is...

Code fragment in line with normal text

Some text in a dialog box

Dialog box output

Something typed in

User input – things you type at the computer keyboard

**XYZ Product Overview**

Document name

◆

Information bullet point

■

Action bullet point – an action you should perform

**Note:** Important Notes are enclosed within a box like this.  
Please pay particular attention to these points to ensure proper configuration and operation of the solution.

**Tip:** Useful information is enclosed within a box like this.  
Use these points to find out where to get more help on a topic.

## 1.5 Feedback

Customer feedback can only improve the quality of our product documentation, and we would welcome any comments, criticisms or suggestions you may have regarding this document.

Please email your feedback to [documentation@caplin.com](mailto:documentation@caplin.com).

## 1.6 Acknowledgments

*Adobe*, *Adobe® Reader*, and *Flex* are either registered trademarks or trademarks of Adobe Systems Incorporated is a registered trademark of Adobe Systems Incorporated in the United States and/or other countries.

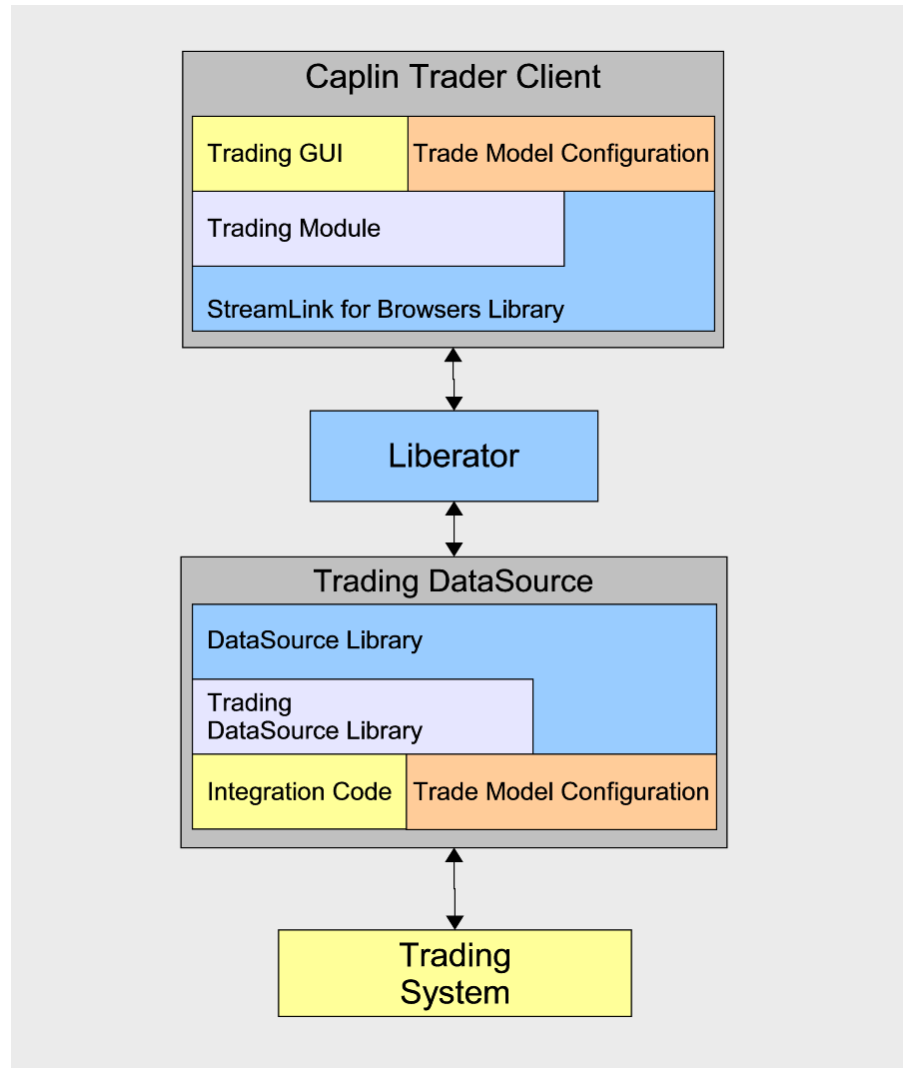
*Windows* is a registered trademark of Microsoft Corporation in the United States and other countries.

*Java* is a trademark of Sun Microsystems, Inc. in the U.S. or other countries.

## 2 Overview

Caplin Trader consists of a number of components (see the **Caplin Trader Architecture**). The main components used to integrate Caplin Trader with your trading system are the Trading DataSource and the Trading GUI.

The following diagram shows the basic architecture of the trading integration components and how they fit into Caplin Trader.



**Simplified Caplin Trader architecture showing only trading integration components**

## Trading system

The trading system represents your systems that support trade capture and execution.

## Trading DataSource

The Trading DataSource is the interface between Caplin Trader and the trading system. Its job is to enable communication between clients and the trading system. It sits between Caplin Liberator and the trading system, handling messages sent between clients and the trading system, via Caplin Liberator. The Trading DataSource consists of the standard Caplin DataSource Library, the Trading DataSource Library, and the custom code required to integrate with your trading system.

The Trading DataSource provides a simple API that can be used to communicate with your trading system or can be integrated directly into it. This means the DataSource can be a stand alone process or part of an existing one. The Trading DataSource API is available in both Java and C++ and is built on top of the Caplin DataSource SDK. It gives full access to all the functionality of Caplin DataSource SDK; this allows you to send and receive custom messages in addition to trade messaging.

## Trading GUI

The Trading GUI is the part of the Caplin Trader Client that displays Trade Tickets and Trade Tiles. It can be customized to contain the correct information and understand the types of trades that can be performed. The Trading GUI consists of custom presentation code that interacts with the Trading DataSource through the Liberator using the client-side Trading Module and the StreamLink for Browsers Library.

The provided Trade Tickets and Trade Tiles can be used as a starting point for customization, or completely new trading displays can be created using the supplied APIs.

## Trade Model configuration

The Trade Model configuration is a set of XML files defining the Trade Models that are to be used by the Trading DataSource and Trading GUI. These definitions represent the trade life-cycle and provide an interface between the end user and the trading system. The same Trade Model configuration is used by both components to ensure they communicate and maintain a consistent state with one another.

Caplin Trader is not tied to any particular Trade Model; it can be configured to match your existing Trade Models along with any new Trade Models you wish to develop. Once configured with Trade Models, the Trading DataSource Library and Trading Module will control and verify the states and transitions allowed. This simplifies the integration needed as most of the logic is handled for you and is defined by your configuration.



## 3 Trading concepts

Caplin Trader uses a number of concepts to represent trading: models, channels, trades, and events.

### 3.1 Trade Models

A Trade Model represents a type of Trade, for example a Request for Quote (RFQ) or Executable Streaming Price (ESP). Trade Models consist of a number of states and transitions and are defined by configuration. The Trade Model controls the flow of a Trade by defining all the possible states the Trade can be in and the messages that cause transitions from one state to another.

### 3.2 Trade Channels

A Trade Channel represents a single user's communication between the Caplin Trader Client and the Trading DataSource. It is a private channel for bidirectional messaging and all messages relating to Trades for a user will be sent and received on the user's channel.

Caplin Trader Client opens a Trade Channel by subscribing to an object. The Caplin Liberator maps this subscription to a unique object name for that user (see [Mapping trade messaging objects in Liberator](#)<sup>[33]</sup>) and subscribes to the object from the Trading DataSource. When the Trading DataSource responds to this subscription, a private channel is effectively created for messages in both directions between the client and the Trading DataSource; this is the Trade Channel.

Many deployments would use a single Trade Channel. However, sometimes it is useful to have multiple Trading DataSources, each handling different asset classes. In this case the user could have a separate Trade Channel for each Trading DataSource; the client would be set up to subscribe to different object names for the different channels.

### 3.3 Trades

A Trade represents a single trade for a user. This could be an RFQ, an Execution on a streaming price, or any other type of trade. A Trade is typically initiated by the client and then the Trading DataSource processes events from the client and the trading system that transition the Trade between different states.

Multiple Trades can be in operation on the same Trade Channel either concurrently or one after the other. Each Trade has an associated RequestId set by the client and a TradeId set by the trading system. These ids are set by the first message sent by either side and are then included in subsequent messages to link the messages to the correct Trade.

A Trade is tied to a Trade Model; this relationship is determined by the first message sent by the client. Once the Trade Model for a Trade has been set, the state of the Trade transitions from the initial state to a final state according to the definition of the Trade Model.

## 3.4 Trade Events

A Trade Event typically represents an action by a client or an event from the trading system. An event originating from the trading system can represent an action by a dealer or an automated action. A Trade Event is raised either by receiving a message from a client, or directly which causes the DataSource to send a message to a client. Events are tied to a Trade and cause the Trade to move from one state to another as defined by the Trade Model.

A Trade Event contains a number of fields and values, which map directly onto a message sent or received by the Trading DataSource. Some message fields are mandatory and are part of the Trading API; for example, all Trade Events have a type which is represented by the MsgType field in the underlying message. Other fields are optional, some of which may be required by the Trade Model being used.

## 3.5 Blotter Channels

A Blotter Channel is a software channel that the Trading DataSource uses to send information to the Client for display in a blotter. In Caplin Trader Client the blotter is updated when a Trade changes state; for example; when a quote is requested, when a Trade is executed, and when a Trade is canceled.

Caplin Trader Client opens a Blotter Channel by subscribing to an object whose subject name begins with the string "/BLOTTER/" or contains the string "/FT/TRADEHISTORY/". The Liberator maps this subscription to a unique object name for the Caplin Trader Client end-user and subscribes to the object from the Trading DataSource. The Trading DataSource responds to this subscription by creating a unique Blotter Channel.

Subsequently, whenever the end-user interacts with the trading subsystem, the Trading DataSource creates a Blotter Event each time the client is notified of a Trade Event. As a result it may send a blotter message to the client across the Blotter Channel. The message typically contains information about the state of the Trade.

The following pictures show the blotter in Caplin Trader Client being updated at successive stages in the execution of an FX Trade using the RFS Trade Model. (For simplicity only the left hand side of the blotter is shown; in reality there are more fields on the right hand side of the blotter.)

1. When the end-user requests a quote the blotter entry is created and its status is set to Opened.

FX Blotter							
ID	Currency Pair	Dealt Currency	Status	B/S	Amount	Rate	S/D
1235491222769	AUDUSD	AUD	Opened	2-WAY	500,000		26 Feb 2009

2. As price quotes are streamed to the client the status changes to Price Update.

FX Blotter							
ID	Currency Pair	Dealt Currency	Status	B/S	Amount	Rate	S/D
1235491222769	AUDUSD	AUD	Price Update	2-WAY	500,000		26 Feb 2009

3. When the user clicks the Sell button on the Trade Ticket, the blotter entry status becomes Executing.

FX Blotter							
ID	Currency Pair	Dealt Currency	Status	B/S	Amount	Rate	S/D
1235491222769	AUDUSD	AUD	Executing	SELL	500,000	0.9180	26 Feb 2009

4. Finally the Trade is confirmed and the final sale details are sent to the blotter with status set to Done.

FX Blotter							
ID	Currency Pair	Dealt Currency	Status	B/S	Amount	Rate	S/D
1235491222769	AUDUSD	AUD	Done	SELL	500,000	0.9180	26 Feb 2009

### Successive updates to a blotter entry in Caplin Trader Client

The client should normally subscribe to a separate Blotter Channel for each asset class traded—for example, an FX blotter and an FI blotter—since the type of information that needs to be displayed on the blotter varies according to asset class. The client would be set up to subscribe to different object names for the different Blotter Channels, such as “/BLOTTER/FX” and “/BLOTTER/FI”.

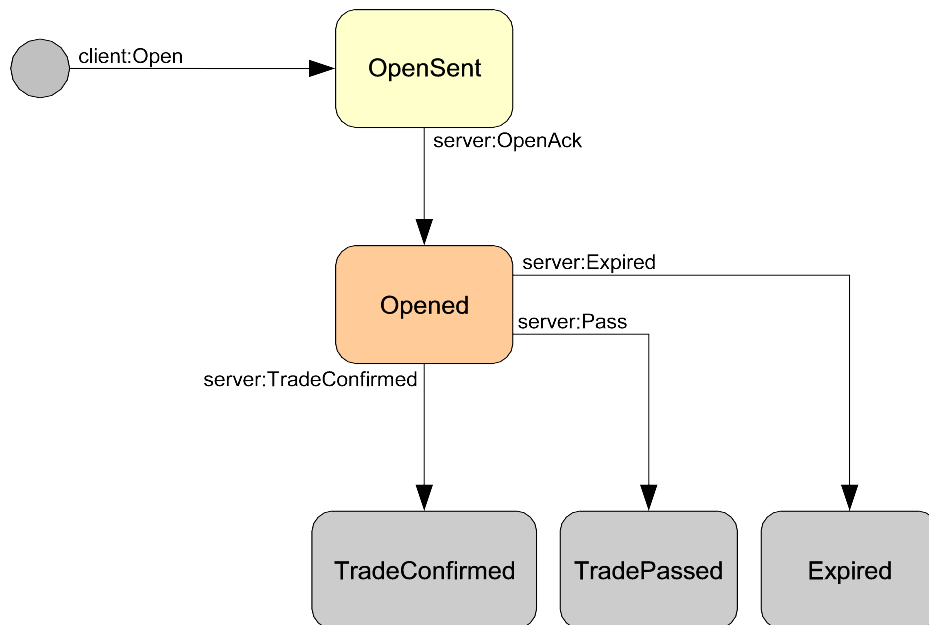
## 4 Example Trade Models

The following sections show examples of Trade Models that can be used with Caplin Trader. More complicated Trade Models can be used with little extra complexity of integration.

### 4.1 Example Executable Streaming Price (ESP)

This state diagram shows the example Executable Streaming Price (ESP) Trade Model.

This model is provided with Caplin Trader and is used by the demonstration Trading DataSource and the Caplin Trader Client Reference Implementation, for one-click trading via the Trade Tile.

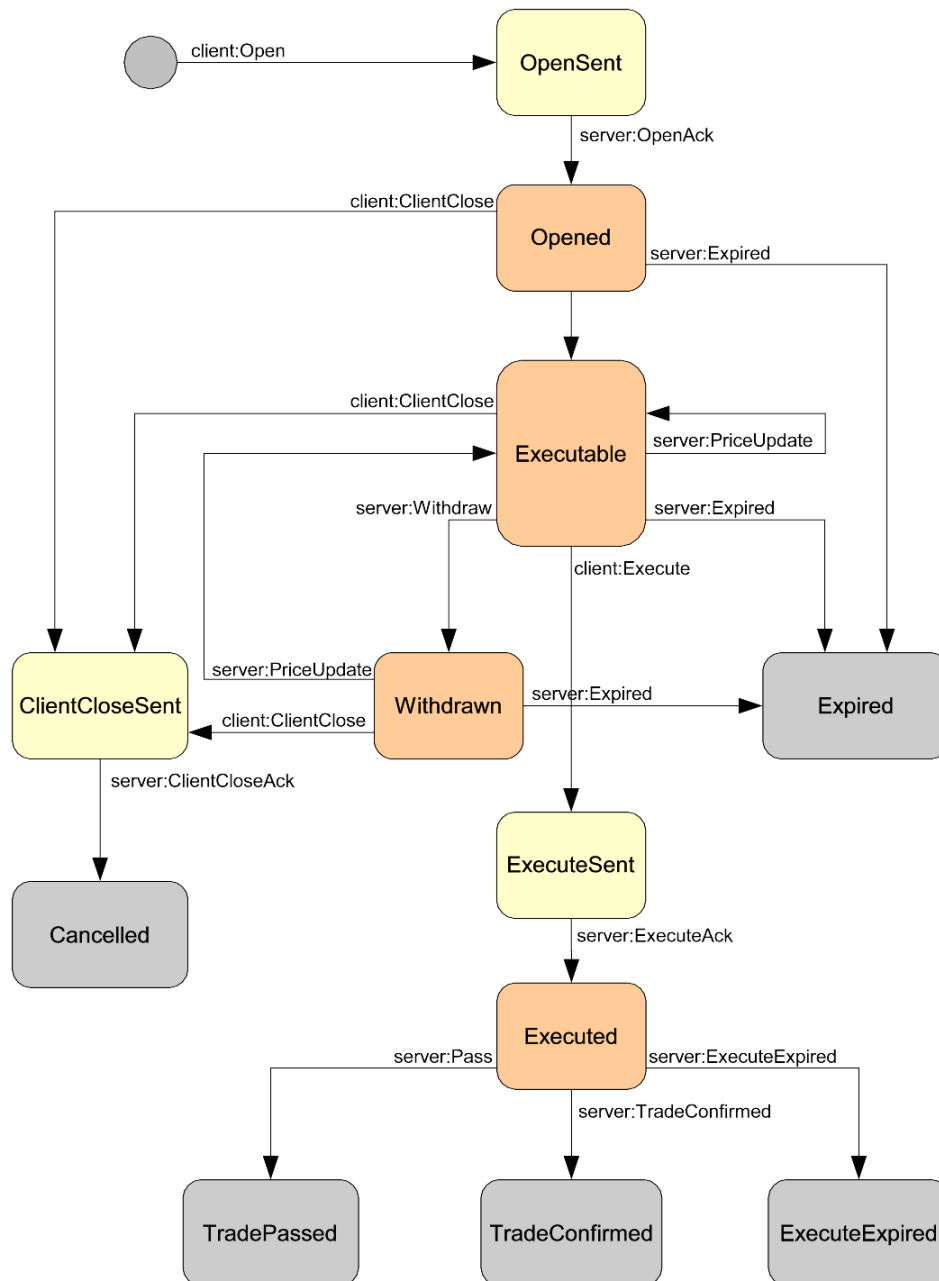


**State diagram for ESP Trade Model**

## 4.2 Example Request for Stream (RFS)

This state diagram shows the example Request For Stream (RFS)

Trade Model. This model is provided with Caplin Trader and is used by the demonstration Trading DataSource and the Caplin Trader Client Reference Implementation, for ticket-based Trades.




State diagram for Example RFS Trade Model

### 4.3 Example Order (ORD)

This state diagram shows the example Order (ORD) Trade Model.

This model differs from the ESP and RFS models in that it has two different transitions from the initial state. The standard transition is the client open event (client:Open), but the additional initial state transition (server:Restore) is used when the server restores a Trade from the trading system.

This model is provided with Caplin Trader and is used by the demonstration Trading DataSource. It is not currently used by the Caplin Trader Client, but will be included in a forthcoming version of the Caplin Trader Client Reference Implementation for ticket-based order Trades.



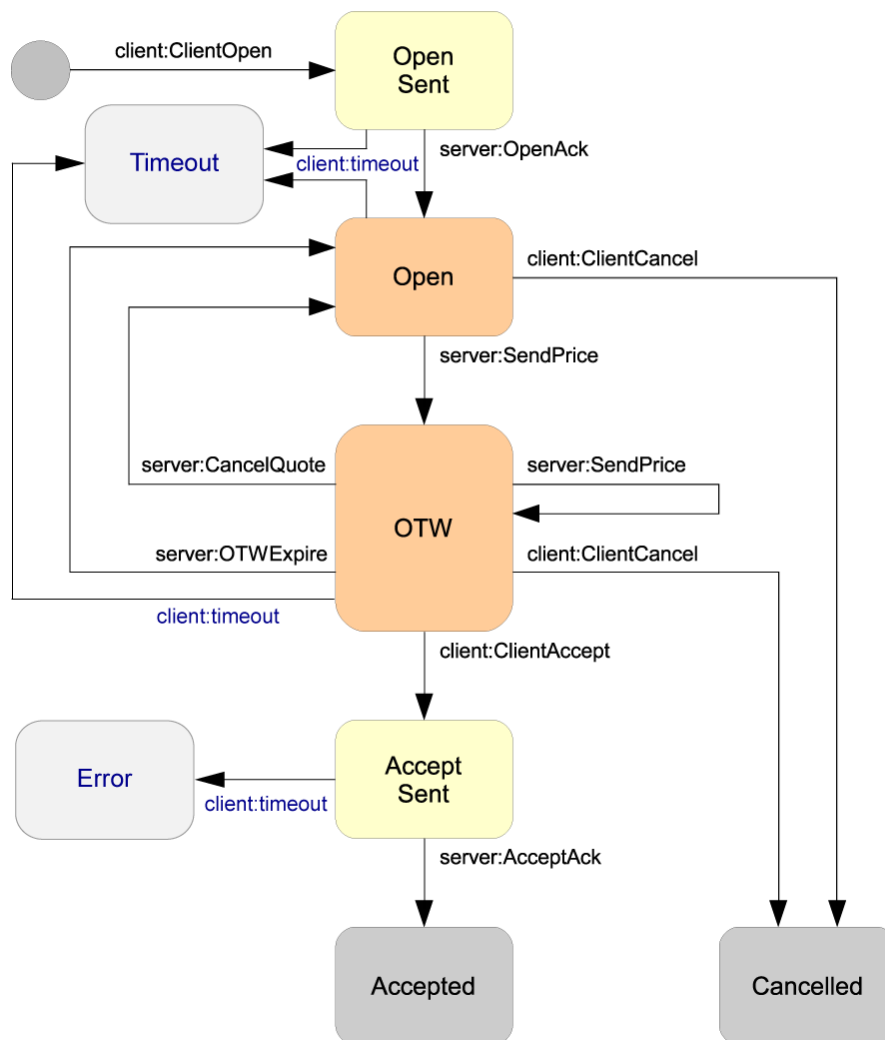
State diagram for Example ORD Trade Model

## 4.4 Request for Quote (RFQ) with timeouts

This state diagram shows a typical simple Request for Quote Trade Model.

This particular model is not used by the Caplin Trader Client Reference Implementation, but is typical of an RFQ workflow. It includes timeouts on events; these timeouts are implemented on the client side only, to ensure that the client does not hang if there is no response from the server. Additional states can easily be added before the Open state to account for credit checks and other validation steps as required.

The XML configuration that describes this trade model is shown in the [RFQ example](#)<sup>[14]</sup> section of [Configuring Trade Models](#)<sup>[13]</sup>.



State diagram for RFQ Trade Model

## 5 Configuring Trade Models

Caplin Trader is designed to work with any Trade Model; it can be configured to match your existing Trade Models or new models being developed.

Trade Models are configured using an XML definition file which defines

- ◆ The possible states a Trade can be in.
- ◆ The transitions a Trade can take from one state to another.
- ◆ The checks that are made before the transition can be made.

Any number of Trade Models can be configured and Trades can automatically pick out the relevant model to use.

The Trading DataSource kit includes example Trade Model XML definition files. These are Request for Stream (RFS), Executable Streaming Price (ESP), and Order (ORD). The XML definitions can be used as they are, or they can be adapted to your requirements.

For the detailed definition of the Trade Model XML, see the document **Caplin Trader Trade Model XML Reference**.



## 5.1 RFQ example

The following example is the XML configuration for the RFQ Trade Model shown in [Request for Quote \(RFQ\) with timeouts](#)<sup>[12]</sup>. Note that the `timeout` and `timeoutState` attributes on some of the `<state>` tags only apply to the Trade Model when it executes on a client. The timeouts ensure that the client does not hang if there is no response from the server.

```
<tradeModels>
  <tradeModel name="RFQ" initialState="Initial">
    <state name="Initial">
      <transition target="OpenSent"
        trigger="ClientOpen"
        source="client" />
    </state>

    <state name="OpenSent" timeout="10" timeoutState="Timeout">
      <transition target="Open"
        trigger="OpenAck"
        source="server" />
    </state>

    <state name="Open" timeout="60" timeoutState="Timeout">
      <transition target="OTW"
        trigger="PriceUpdate"
        source="server" />
      <transition target="Cancelled"
        trigger="ClientCancel"
        source="client" />
    </state>

    <state name="OTW" timeout="60" timeoutState="Timeout">
      <transition target="AcceptSent"
        trigger="ClientAccept"
        source="client" />
      <transition target="OTW"
        trigger="PriceUpdate"
        source="server" />
      <transition target="Open"
        trigger="OTWExpire"
        source="server" />
      <transition target="Cancelled"
        trigger="ClientCancel"
        source="client" />
    </state>

    <state name="AcceptSent" timeout="10" timeoutState="Error">
      <transition target="Accepted"
        trigger="AcceptAck"
        source="server" />
    </state>

    <state name="Accepted" />
    <state name="Cancelled" />
    <state name="Timeout" />
    <state name="Error" />
  </tradeModel>
</tradeModels>
```

See also the **Caplin Trader Trade Model Configuration XML Reference**.

## 6 Using the Trading DataSource Java API

This section provides a brief description of the main parts of the Trading DataSource Java API (see the **Caplin Java Trading DataSource: API Documentation** for full details).

To implement a Trading DataSource you set up listener objects to handle events occurring on Trades and create events to be sent back into the system.

### 6.1 Initialization

When your Trading DataSource application starts, it should create an instance of `TradingDataSource` and register itself as a `TradingApplicationListener`. The `TradingDataSource` starts up a `DataSource`, which then connects to the Liberator. The `DataSource` also sets up the necessary `DataSource` listeners internally to handle the trade messaging.

The following is a simple code extract showing the creation of a `TradingDataSource` within a custom application.

#### Creation of a TradingDataSource

```
public class MyTradingApp implements TradingApplicationListener
{
    static void main(String[] args)
    {
        new MyTradingApp();
    }

    MyTradingApp()
    {
        // Create a factory object for generating the trading state machines.
        StateMachineFactory myTradingStateMachineFactory = new StateMachineFactory();

        // Load the required Trade Models into the factory.
        myTradingStateMachineFactory.loadModels(new File("conf/MyESPStateModel.xml"));
        myTradingStateMachineFactory.loadModels(new File("conf/MyRFSStateModel.xml"));

        // Create the Trading DataSource.
        // This will create a DataSource object internally to manage the
        // communication with other DataSources, such as Caplin Liberator
        // and hence client applications.
        // TradingDataSource implements the standard DataSource callbacks,
        // which allow you to use the Trading API to communicate with
        // clients in the form of trade messages.

        tradingDataSource = new TradingDataSource
            (this, //Reference to this TradingApplicationListener
            "MyDataSourceConfig.xml",
            // The configuration file for this DataSource
            myTradingStateMachineFactory
            );

        // Once the TradingDataSource has been created
        // it has to be started explicitly.
        tradingDataSource.start();
    }

    // ...
}
```

The `TradingDataSource` processes the flow of a Trade by following the states defined by the Trade Model for the particular type of Trade. This processing is carried out by a state machine, implemented as a `StateMachine` object. Before creating the `TradingDataSource`, the Trading DataSource application creates a `StateMachineFactory`, and loads the factory with the required Trade Models. The models are defined in the XML configuration files discussed in [Configuring Trade Models](#)<sup>13</sup>. The `StateMachineFactory` is then passed to the `TradingDataSource` constructor, so that the DataSource can create a the state machine for each loaded Trade Model.

The `TradingDataSource` constructor is also supplied with an XML format configuration file for the DataSource. This file defines the connections that the DataSource makes with the other Caplin Platform components, such as Liberator.

## 6.2 New Trade Channels

The `TradingApplicationListener` is notified when new Trade Channels are created; this allows you to perform any necessary user specific initialization. You must also add a `ChannelListener` to the channel. The `ChannelListener` is an interface you must implement; it could be a new instance for each channel or a single global instance. Its job is to handle notifications on the channel about newly created Trades and Trades being closed.

The following code extract shows part of a sample implementation of a `TradingApplicationListener`. It shows a custom `ChannelListener` being added to the channel to handle Trades, and then a method being called on a hypothetical `tradingSystem` object to log in the user for the channel.

### Example implementation of `TradingApplicationListener.channelCreated()`

```
public void channelCreated(TradeChannel channel)
{
    channel.setChannelListener(new MyChannelListener(channel));

    // Handle new channel/user.
    // For example:
    tradingSystem.loginUser(channel.getUser());
}
```

## 6.3 New Trades

The `ChannelListener` is notified of new Trades when a client initiates them; the Trade is passed to it as a Trade object. When a Trade is created the `ChannelListener` can perform any initialization needed with the trading system and also add a `TradeListener` to the newly created Trade object. The `TradeListener` is an interface you must implement and could be a new instance for each Trade or a single global instance; its job is to handle all events for a Trade. You normally have different implementations of `TradeListener` that handle different trade types.

The following code extract shows part of a sample implementation of a `ChannelListener`. It shows a different custom `TradeListener` being added to the Trade object to handle its events, depending on the Trade Model used for the Trade.

#### Example implementation of ChannelListener.tradeCreated()

```
public void tradeCreated(Trade trade)
{
    // Check the Trade Model used
    // and create an appropriate listener to handle it.
    if (trade.getType().equals("ORD"))
    {
        trade.setTradeListener(new ORDTradeListener(trade));
    }
    else if (trade.getType().equals("RFQ"))
    {
        trade.setTradeListener(new RFQTradeListener(trade));
    }
}
```

How the new Trade is handled depends on the trading system to which the Trading DataSource is connected, and the nature of the API to that system.

The Trade object must be retained, so that it can be referred to when the trading system responds with an event (see [Dealing with events](#)<sup>[18]</sup>). Assume, for example, the trading system API supports a listener style interface with a listener object for each Trade. TradeListener.tradeCreated() can store the Trade object in a trading system listener object before calling the trading system. When the trading system subsequently raises an event on the Trade, it will call the listener, which can then refer to the Trade object as required (for example to create an event to pass on to the client).

Alternatively the trading system may not support a listener interface. For example, it may, just pass back an “event” with an ID relating to the Trade. In this case TradeListener.tradeCreated() would have to store the Trade object in a suitable data structure (say a hash table). This structure must be accessible by the code that handles events from the trading system. This code would typically use the ID returned in the trading system event as the key to extract the Trade object.

## 6.4 Dealing with Events

The TradeEvent object represents a Trade Event, which typically encapsulates a message between the client and the Trading DataSource. A TradeEvent has a type, which represents the type of the message, for example “Open”, “PriceUpdate” or “Execute”. It also has a number of fields to represent all the necessary information for that message, for example “BidPrice” or “Amount”.

The TradeListener is responsible for handling TradeEvents; it is notified when new events are received from the client. When a message is received from the client, it is processed by the Trading DataSource to verify that the event is valid based on the Trade Model before notifying the TradeListener of the event. The Trade and TradeEvent objects are passed to the TradeListener. The Trade has been updated with the data from the TradeEvent and can be used to create and send new TradeEvents. The TradeListener would then typically send a message on to the trading system or handle the event in some other way.

The following code extract shows an implementation of the TradeListener method to receive events.

#### Example implementation of TradeListener.receiveEvent()

```
public void receiveEvent(TradeEvent event)
{
    // Talk to trading system
}
```

Events raised by the trading system can be pushed into the Trading DataSource. This is done by creating a `TradeEvent` from the relevant `Trade` object, setting the necessary attributes, and asking the `Trade` object to send it. At this point the Trading DataSource will verify, through the Trade Model, that the event is allowed and contains all the necessary information, before sending the message off to the client.

The following code extract shows the typical custom code that would be written in the Trading DataSource to create an event and send it to a client.

**Custom code to create an event**

```
TradeEvent myEvent = trade.createEvent("PriceUpdate");
myEvent.addField("BidPrice", bidPrice);

// Add more fields
// ...

// Then send the event on to the client.
trade.sendEvent(myEvent);
```

## 6.5 Closing Trades

When a `Trade` reaches a final state it is closed. This could happen when the user or the trading system cancels the `Trade`, when the `Trade` is successfully executed, or when it is rejected. The final states are defined by the Trade Model and are the states that have no transitions to another state. The `ChannelListener` is notified when a `Trade` has reached this state, which allows the application to clean up any resources associated with that `Trade`.

The following code extract shows the implementation of the `ChannelListener` method for notifying closed Trades.

**Example implementation of `ChannelListener.tradeClosed()`**

```
public void tradeClosed(Trade trade)
{
    // Clean up
}
```

## 6.6 Closing Channels

When a user logs off the system the Trade Channel for that user is closed. This could also happen if the client application is designed to close Trade Channels when they are not in use. The `TradingApplicationListener` will be notified when a channel is closed, which allows any resources associated with that channel to be cleaned up. Once the Trade has been closed it can no longer be used to create or send events.

The following code extract shows the implementation of the `TradingApplicationListener` method for notifying closed channels.

**Example implementation of `TradingApplicationListener.channelClosed()`**

```
public void channelClosed(TradeChannel channel)
{
    // Clean up
}
```

## 6.7 Handling Blotter Channels

To handle Blotter Channels in the trading DataSource:

- ◆ Implement code in the `TradingApplicationListener` interface to register and deregister a `BlotterTradeListener`.
- ◆ Implement the `BlotterTradeListener` interface to construct and send blotter messages.  
This interface provides notification of Blotter Events through the life cycle of a Trade. Blotter events (class `BlotterEvent`) are created when the Trading DataSource has validated a state transition in the Trade Model and has sent a `TradeEvent` to the client.

In the code fragments shown in the next sections, the implementation of `BlotterTradeListener` is the example listener `AutoBlotterTradeListener` provided with the example Trading DataSource:

### Registering the BlotterTradeListener

When implementing the `TradingApplicationListener` interface (see [Initialization](#) <sup>16</sup>), add code to register and deregister a `BlotterTradeListener`.

#### Registering and deregistering BlotterTradeListener in TradingApplicationListener

```
public class MyTradingApp implements TradingApplicationListener
{
    ...
    private BlotterTradeListener autoBlotterTradeListener
        = AutoBlotterTradeListener.getInstance();
    ...

    // Called when the TradingDataSource has created a BlotterChannel
    // as a result of receiving a request for a blotter subject.

    public void blotterChannelCreated(BlotterChannel blotterChannel)
    {
        tradingDataSource.addBlotterTradeListener(
            blotterChannel,
            autoBlotterTradeListener
        );
    }

    // Called when a BlotterChannel is closed.
    // The channel is closed when the originally requested blotter
    // subject is discarded, or the connection to the Trading DataSource's
    // peer is lost, or the Trading DataSource is being shut down.

    public void blotterChannelClosed(BlotterChannel blotterChannel)
    {
        tradingDataSource.removeBlotterTradeListener(
            blotterChannel,
            autoBlotterTradeListener
        );
    }
}
```

Here is more detailed explanation of the previous code fragment:

- ◆ The `BlotterTradeListener` is the example listener `AutoBlotterTradeListener` provided with the example Trading DataSource. The static method `AutoBlotterTradeListener.getInstance()` returns an instance of the listener that can be used with this Blotter Channel:

```
private BlotterTradeListener autoBlotterTradeListener
    = AutoBlotterTradeListener.getInstance();
```

- ◆ Code the `blotterChannelCreated()` method to register the `BlotterTradeListener` with the Trading DataSource when the Blotter Channel is created, by calling `addBlotterTradeListener()`:

```
public void blotterChannelCreated(BlotterChannel blotterChannel)
{
    tradingDataSource.addBlotterTradeListener(
        blotterChannel,
        autoBlotterTradeListener
    );
}
```

- ◆ Code the `blotterChannelClosed()` method to deregister the `BlotterTradeListener` from the Trading DataSource when the Blotter Channel is closed:

```
public void blotterChannelClosed(BlotterChannel blotterChannel)
{
    tradingDataSource.removeBlotterTradeListener(
        blotterChannel,
        autoBlotterTradeListener
    );
}
```



## Implementing the BlotterTradeListener interface

The `BlotterTradeListener` has one method `receiveBlotterEvent()` that is called when the Trading DataSource has sent a Trade Event to the client. Note that the `BlotterEvent` is not an event in the execution of a Trade Model; it is merely a message containing both the `TradeEvent` for which a blotter entry is to be constructed and the `Blotter Channel` to send the entry on.

Here is a simple example of the `receiveBlotterEvent()` method in the `AutoBlotterTradeListener` implementation of `BlotterTradeListener`.

### AutoBlotterTradeListener (example of BlotterTradeListener)

```
public class AutoBlotterTradeListener implements BlotterTradeListener
{
    // Called after a Trade Event has been validated and sent to the client.

    public void receiveBlotterEvent(BlotterEvent blotterEvent)
    {
        BlotterChannel blotterChannel = blotterEvent.getBlotterChannel();
        TradeEvent tradeEvent = blotterEvent.getTradeEvent();
        Trade trade = tradeEvent.getTrade();

        BlotterMessage blotterMessage = blotterChannel.createBlotterMessage();

        blotterMessage.addField("L1_AMOUNT", trade.getField("L1_AMOUNT"));
        blotterMessage.addField("ACCOUNT", trade.getField("ACCOUNT"));
        blotterMessage.addField("TRADING_TYPE", "demo");
        blotterMessage.addField("USER_NAME", trade.getChannel().getUser());
        ...

        blotterChannel.sendBlotterMessage(blotterMessage);
    }
}
```

Here is more detailed explanation of the previous code fragment:

- ◆ Obtain the `BlotterChannel` from the `BlotterEvent` that the Trading DataSource passed to the listener. From the `BlotterEvent` obtain the `TradeEvent` and the `Trade` to which the Trade Event relates.

```
BlotterChannel blotterChannel = blotterEvent.getBlotterChannel();
TradeEvent tradeEvent = blotterEvent.getTradeEvent();
Trade trade = tradeEvent.getTrade();
```

- ◆ Create a new blotter message on the `Blotter Channel`.

```
BlotterMessage blotterMessage = blotterChannel.createBlotterMessage();
```

- ◆ Populate the blotter message with the required fields and their values. Typically the field values are obtained from the Trade and from the user information associated with the Trade Channel.

```
blotterMessage.addField("L1_AMOUNT", trade.getField("L1_AMOUNT"));
blotterMessage.addField("ACCOUNT", trade.getField("ACCOUNT"));
blotterMessage.addField("TRADING_TYPE", "demo");
blotterMessage.addField("USER_NAME", trade.getChannel().getUser());
...
```

- ◆ Send the newly constructed blotter message to the client via the Blotter Channel.

```
blotterChannel.sendBlotterMessage(blotterMessage);
```

## 7 The Java Trading DataSource Example

The Java Trading DataSource is supplied with example code to help you get started. This is in the Java package `example`, which is located in the folder `examples/source`. The `example` package contains commented example code that shows how to use the Trading DataSource API, as described in [Using the Trading DataSource Java API](#)<sup>16</sup>. The package is set up to handle Request For Stream (RFS) and Executable Streaming Price (ESP) Trade Models, but could be easily adapted to handle any Trade Models you configure.

The example is also used as the Reference Implementation Trading DataSource for ESP and RFS Trades made through the Caplin Trader Client Reference Implementation, and shares the Trade Model definitions with the client.

## 8 Using the Trading DataSource C++ API

This section provides a brief description of the main parts of the Trading DataSource C++ API (see the **Caplin C++ Trading DataSource: API Documentation** for full details).

To implement a Trading DataSource you make use of callbacks on a class of your choice, to be notified of events occurring on Trades and to create events to be sent back into the system.

### 8.1 Initialization

When your Trading DataSource application starts, it should create an instance of `TradingDataSource` and register itself as a `TradingApplicationListener`.

The `TradingDataSource` starts up a `DataSource`, which then connects to the Liberator. The `DataSource` also sets up the necessary `DataSource` listeners internally to handle the trade messaging.

The following is a simple code extract showing the creation of a `TradingDataSource` within a custom application.

#### Creation of a TradingDataSource

```
using namespace Caplin::TradingDataSource;

class MyTradingApp : public TradingApplicationListener
{
public:
    MyTradingApp()
    {
        // Prepare a vector of filenames of the config files to use
        std::vector<std::string> configFiles;
        configFiles.push_back("conf/MyESPStateModel.xml");
        configFiles.push_back("conf/MyRFSStateModel.xml");

        // Create the Trading DataSource, passing in a pointer to a class that
        // implements TradingApplicationListener (so our 'this' pointer is fine)
        m_tradingDataSource.reset(new TradingDataSource(this,
                                                         "MyDataSourceConfig.conf",
                                                         configFiles));

        // Once the TradingDataSource has been created
        // it has to be started explicitly.
        m_tradingDataSource->start();
    }

    // ...
private:
    std::auto_ptr<TradingDataSource> m_tradingDataSource;
}
```

The `TradingDataSource` processes the flow of a Trade by following the states defined by the Trade Model for the particular type of Trade. This processing is carried out by a state machine, implemented internally as a `StateMachine` object. Before creating the `TradingDataSource`, the `Trading DataSource` application creates a `StateMachineFactory` internally and loads the factory with the required Trade Models. The models are defined in the XML configuration files discussed in [Configuring Trade Models](#) and their paths are passed in to the constructor of `TradingDataSource` either as a vector of strings or a single string.

The `TradingDataSource` constructor is also supplied with a configuration file for the `DataSource`. This file defines the connections that the `DataSource` makes with the other Caplin Platform components, such as Liberator.

## 8.2 New Trade Channels

The `TradingApplicationListener` is notified when new Trade Channels are created; this allows you to perform any necessary user specific initialization. You must also add a `ChannelListener` to the channel. The `ChannelListener` is an interface you must implement; it could be a new instance for each channel or a single global instance. Its job is to handle notifications on the channel about newly created Trades and Trades being closed.

The following code extract shows part of a sample implementation of a `TradingApplicationListener`. It shows a custom `ChannelListener` being added to the channel to handle Trades, and then a method being called on a hypothetical `tradingSystem` object to log in the user for the channel.

### Example implementation of `TradingApplicationListener::channelCreated()`

```
void MyTradingApplicationListener::channelCreated(TradeChannel& channel)
{
    channel.setChannelListener(m_pTradeChannelListener);

    // Handle new channel/user.
    // For example:
    tradingSystem.loginUser(channel.getUser());
}
```

## 8.3 New Trades

The `ChannelListener` is notified of new Trades when a client initiates them; the Trade is passed to it as a `Trade` object. When a Trade is created the `ChannelListener` can perform any initialization needed with the trading system and also add a `TradeListener` to the newly created `Trade` object. The `TradeListener` is an interface you must implement and could be a new instance for each Trade or a single global instance; its job is to handle all events for a Trade. You normally have different implementations of `TradeListener` that handle different trade types.

The following code extract shows part of a sample implementation of a `ChannelListener`. It shows a different custom `TradeListener` being added to the `Trade` object to handle its events, depending on the Trade Model used for the Trade.

### Example implementation of `ChannelListener::tradeCreated()`

```
void MyTradeChannelListener::tradeCreated(Trade& trade)
{
    if (trade.isType("ESP"))
    {
        trade.setTradeListener(m_pESPTradeListener);
    }
    else if (trade.isType("RFS"))
    {
        trade.setTradeListener(m_pRFSTradeListener);
    }
}
```

How the new Trade is handled depends on the trading system to which the Trading DataSource is connected, and the nature of the API to that system.

## 8.4 Dealing with Events

The `TradeEvent` object represents a Trade Event, which typically encapsulates a message between the client and the Trading DataSource. A `TradeEvent` has a type, which represents the type of the message, for example "Open", "PriceUpdate" or "Execute". It also has a number of fields to represent all the necessary information for that message, for example "BidPrice" or "Amount".

The `TradeListener` is responsible for handling `TradeEvents`; it is notified when new events are received from the client. When a message is received from the client it is processed by the Trading DataSource to verify that the event is valid based on the Trade Model before notifying the `TradeListener` of the event. The `Trade` and `TradeEvent` objects are passed to the `TradeListener`. The `Trade` has been updated with the data from the `TradeEvent` and can be used to create and send new `TradeEvents`. The `TradeListener` would then typically send a message on to the trading system or handle the event in some other way.

The following code extract shows an implementation of the `TradeListener` method to receive events.

### Example implementation of `TradeListener::receiveEvent()`

```
void RFQTradeListener::receiveEvent(Trade& trade,
                                    const TradeEvent& tradeEvent)
{
    // Talk to trading system
}
```

Events raised by the trading system can be pushed into the Trading DataSource. This is done by creating a `TradeEvent` from the relevant `Trade` object, setting the necessary attributes, and asking the `Trade` object to send it. At this point the Trading DataSource will verify, through the Trade Model, that the event is allowed and contains all the necessary information, before sending the message off to the client.

The following code extract shows the typical custom code that would be written in the Trading DataSource to create an event and send it to a client.

### Custom code to create an event

```
TradeEvent myEvent = trade.createEvent("PriceUpdate ");
myEvent.addField("BidPrice", bidPrice);

// Add more fields
// ...

// Then send the event on to the client.
trade.sendEvent(myEvent);
```

## 8.5 Closing Trades

When a Trade reaches a final state, it is closed. This could happen when the user or the trading system cancels the Trade, when the Trade is successfully executed, or when it is rejected. The final states are defined by the Trade Model and are the states that have no transitions to another state. The `ChannelListener` is notified when a Trade has reached this state, which allows the application to clean up any resources associated with that Trade. Once the Trade has been closed it can no longer be used to create or send events.

The following code extract shows the implementation of the `ChannelListener` method for notifying closed Trades.

### Example implementation of `ChannelListener.tradeClosed()`

```
void MyTradeChannelListener::tradeClosed(Trade& trade)
{
    // Clean up
    // ...
}
```

## 8.6 Closing Channels

When a user logs off the system the Trade Channel for that user is closed. This could also happen if the client application is designed to close Trade Channels when they are not in use. The `TradingApplicationListener` will be notified when a channel is closed, which allows any resources associated with that channel to be cleaned up.

The following code extract shows the implementation of the `TradingApplicationListener` method for notifying closed channels.

### Example implementation of `TradingApplicationListener.channelClosed()`

```
void MyTradingApplicationListener::channelClosed(TradeChannel& channel)
{
    // Clean up
    // ...
}
```

## 8.7 Handling Blotter Channels

To handle Blotter Channels in the trading DataSource:

- ◆ Implement code in the `TradingApplicationListener` class to register and deregister a `BlotterListener`.
- ◆ Implement the `BlotterListener` class to construct and send blotter messages.  
This interface provides notification of Blotter Events through the life cycle of a Trade. Blotter events are created when the Trading DataSource has validated a state transition in the Trade Model and has sent a `TradeEvent` to the client.

### Registering the BlotterListener

When implementing the `TradingApplicationListener` interface (see [Initialization](#)<sup>[25]</sup>), add code to register and deregister a `BlotterListener`.

- ◆ Code the `blotterChannelCreated()` method to register the `BlotterListener` with the Trading DataSource when the Blotter Channel is created, by calling `setBlotterListener()`.

#### Registering the BlotterListener in TradingApplicationListener

```
void DemoTradingSource::blotterChannelCreated(BlotterChannel& blotterChannel)
{
    m_tradingDataSource->setBlotterListener(blotterChannel,
                                           this           //The BlotterListener
                                           );
}
```

In this example the `DemoTradingSource` implements both `TradingApplicationListener` and `BlotterListener`, so the `BlotterListener` argument of `setBlotterListener` is passed as `this`.

- ◆ Code the `blotterChannelClosed()` method to deregister the `BlotterListener` from the Trading DataSource when the Blotter Channel is closed:

#### Deregistering the BlotterListener in TradingApplicationListener

```
void DemoTradingSource::blotterChannelClosed(BlotterChannel& blotterChannel)
{
    m_tradingDataSource->removeBlotterListener(blotterChannel);
}
```



## Implementing the BlotterListener interface

The `BlotterListener` has one method `receiveBlotterEvent()` that is called when the Trading DataSource has validated a state transition in the Trade Model and has sent a `TradeEvent` to the client.

Here is a simple example of the `receiveBlotterEvent()` method. In this example the `DemoTradingSource` implements `BlotterListener`, and so contains the code of `receiveBlotterEvent()`.

### Example of `BlotterListener::receiveBlotterEvent()`

```
void DemoTradingSource::receiveBlotterEvent (
    TradingDataSource::BlotterChannel& blotterChannel,
    const TradingDataSource::Trade& trade,
    const TradingDataSource::TradeEvent& tradeEvent
)
{
    BlotterMessage blotterMessage = blotterChannel.createMessage();

    blotterMessage.addFields(trade.getFields());

    std::string status = tradeEvent.getField("Status");
    blotterMessage.addField("Status", status);
    blotterMessage.addField("TradeDate", "20080808");
    blotterMessage.addField("TimeStamp", "123456");
    blotterMessage.addField("UserName", trade.getUser());
    ...

    blotterChannel.sendMessage(blotterMessage);
}
```

Here is more detailed explanation of the previous code fragment:

- ◆ Create a new blotter message on the Blotter Channel.

```
BlotterMessage blotterMessage = blotterChannel.createMessage();
```

- ◆ Populate the blotter message with the required fields and their values. Typically the field values are obtained from the Trade.

```
blotterMessage.addFields(trade.getFields());

std::string status = tradeEvent.getField("Status");
blotterMessage.addField("Status", status);
blotterMessage.addField("TradeDate", "20080808");
blotterMessage.addField("TimeStamp", "123456");
blotterMessage.addField("UserName", trade.getUser());
...
```

- ◆ Send the newly constructed blotter message to the client via the Blotter Channel.

```
blotterChannel.sendMessage(blotterMessage);
```

## 9 The C++ Trading DataSource Example

The C++ Trading DataSource is supplied with an example application to help you get started.

The Visual Studio 2005 project file for the example is in the folder *examples\DemoTradingSource*.

The source files *DemoTradingSource.cpp* and *DemoTradingSource.h* contain the example code that shows how to use the Trading DataSource API as described in this document. The example is set up to handle Request For Stream (RFS) and Executable Streaming Price (ESP) Trade Models, but could be easily adapted to handle any Trade Models you configure.

## 10 Configuring Caplin Liberator for trading

Caplin Liberator is an integral part of Caplin Trader (see the diagram in the [Overview](#)<sup>[4]</sup>) and therefore must be correctly configured to support trading activity.

The following aspects of trading activity are determined through Liberator configuration:

- ◆ Associating a unique user with a trade message and ensuring that one user cannot trade on behalf of another; see [Mapping trade messaging objects in Liberator](#)<sup>[33]</sup>.
- ◆ [Routing of trade messages](#)<sup>[35]</sup> to the correct Trading DataSource.
- ◆ [Trading performance and integrity](#)<sup>[36]</sup>.
- ◆ [Security](#)<sup>[38]</sup>.

The following table lists the Liberator configuration items relevant to trading:

Configuration item	Parameter	See section
add-object	<i>name</i>	<a href="#">Mapping trade messaging objects in Liberator</a> <sup>[33]</sup>
add-object	<i>throttle-times</i>	<a href="#">Throttling</a> <sup>[36]</sup>
add-object	<i>discard-timeout</i>	<a href="#">Throttling</a> <sup>[36]</sup>
add-peer	—	<a href="#">Routing trade messages to the Trading DataSource</a> <sup>[35]</sup>
add-data-service	—	<a href="#">Routing trade messages to the Trading DataSource</a> <sup>[35]</sup>
object-map	—	<a href="#">Mapping trade messaging objects in Liberator</a> <sup>[33]</sup>
output-queue-size	—	<a href="#">Optimizing client reconnection time</a> <sup>[36]</sup>
session-id-len	—	<a href="#">Session IDs</a> <sup>[38]</sup>

**Note:** The configuration items discussed here specifically relate to using Liberator in a trading environment, and in particular to support Caplin Trader Client. There are many other aspects of Liberator functionality and performance that also need to be configured in an implementation of the Caplin Platform. For more information about configuring Liberator, see the **Liberator Administration Guide**.

**Tip:** The examples of Liberator configuration shown in the following sections are derived from the configuration file supplied with the Caplin Trader Reference Implementation. This file is *rtttd.conf*, located in `$CT_INSTALL_DIR/apps/caplin/Liberator/etc/`, where `$CT_INSTALL_DIR` is the directory in which Caplin Trader has been installed.

## 10.1 Mapping trade messaging objects in Liberator

Trade messages are passed between client and Liberator as updates to subscriptions. The subject of the update identifies it as a trade message. The Liberator configuration must define a base subject name for these subscriptions, as a "built-in" Liberator object that is created when the Liberator starts up. For example, in the Caplin Trader Reference Implementation, the Liberator configuration file defines an object called `/FT/TRADE`, so that all trade messages sent between client and Liberator have a subject name that starts with `/FT/TRADE`.

```
add-object
    name      /FT/TRADE
    ...
end-object
```

**Tip:** The **add-object** configuration item can also take additional optional parameters, as shown by the ... in the previous example. For more information, see the Configuration Reference section of the **Liberator Administration Guide**, and the section on [Throttling](#)<sup>36</sup>.

A client opens a Trade Channel by subscribing to one of the built-in Liberator trade message objects, for example, `/FT/TRADE`. Because Liberator and the Trading DataSource need to manage many users who are simultaneously trading, the Liberator maps the generic trade message objects onto user specific object names. This defines the unique channel over which each user trades. The mapping is defined using the **object-map** configuration item.

### Example object mapping

```
object-map /FT/TRADE/%1 /FT/TRADE/%1/%U
```

`%U` is the unique user name (Liberator login name) associated with the Liberator session.

`%1` represents any variable length string appearing in the subject name of the subscription.

When a user called "UserA" connects to Liberator and trades, the trade messages sent between the client and Liberator have a subject name of the form:

`/FT/TRADE/FX/<identification-of-this-trade>`

Before passing an incoming trade message to the Trading DataSource, Liberator maps the subject name in the message according to the **object-map** configuration, so that the subject of the passed on message is:

`/FT/TRADE/FX/< identification-of-this-trade>/UserA-1`

where **UserA-1** is the unique user name assigned to this Liberator login of "UserA".

This transformation allows the Trading DataSource to distinguish between trade messages from "UserA" and trade messages relating to other users. The Trade Channel for "UserA" is uniquely defined by the combination of the strings `/FT/TRADE/FX/` and `UserA-1` in the message subject.

Similarly, when Liberator receives a trade message with subject `"/FT/TRADE/FX/< identification-of-this-trade>/UserA-1"` from the Trading DataSource, it can readily determine the client connection over which it should forward the message to the client, and the forwarded message is given the subject `"/FT/TRADE/FX/< identification-of-this-trade>/"`.

## Preventing identity theft

The %1 parameter in the **object-map** helps to prevent a user from trading using another user's identity. For example, if "UserB", should attempt to fake a trade message so that it appears to come from "UserA", Liberator will reject the message. The faked trade message from the "UserB" client would have the subject

```
/FT/TRADE/FX/<identification-of-this-trade>/UserA
```

However the object mapping in the Liberator transforms this subject name to

```
/FT/TRADE/FX/<identification-of-this-trade>/UserA/UserB-1
```

which Liberator recognizes as a subject for which no subscription exists, so the message is rejected.

## 10.2 Routing trade messages to the Trading DataSource

The Liberator configuration defines the DataSource peers to which Liberator is connected and the data services that Liberator provides to subscribing clients. In Caplin Trader this configuration must include the Trading DataSources and their associated data services.

**Tip:** For more information about data services, see the Caplin DataSource Overview.

For example, in the Caplin Trader Reference Implementation, the Liberator configuration defines connections to two Trading DataSources, one for FX trading and one for FI trading:

### Example of Trading DataSource configuration in Liberator

```
# fxtradesource
add-peer
  remote-id      17
  remote-type    active
  remote-name    fxtradesrc
  label          fxtradesrc
end-peer

#fitradesource
add-peer
  remote-id      18
  remote-type    active
  remote-name    fitradesrc
  label          fitradesrc
end-peer

...

add-data-service
  service-name    fx-trade-data
  include-pattern ^/FT/TRADE/FX

  add-source-group
    required
    add-priority
    label          fxtradesrc
    end-priority
  end-source-group
end-data-service

add-data-service
  service-name    fi-trade-data
  include-pattern ^/FT/TRADE/FI

  add-source-group
    required
    add-priority
    label          fitradesrc
    end-priority
  end-source-group
end-data-service
```

The **add-peer** configuration items define connections to two Trading DataSources, one for FX Trading (`fxtradesrc`) and one for FI Trading (`fitradesrc`).

There is a data service for each of these Trading DataSources, defined by the **add-data-service** configuration item. The FX service (`service-name fx-trade-data`) has an *include-pattern* parameter that ensures all trade message subscriptions whose subject begins with `/FT/TRADE/FX` are directed to the FX Trading DataSource. Similarly the FI Service definition (`service-name fi-trade-data`) directs to the FI Trading DataSource any trade message subscription whose subject begins with `/FT/TRADE/FI`.

Note that both the *include-patterns* start with the base subject name for trade messages as defined in an **add-object** configuration item; see [Mapping trade messaging objects in Liberator](#)<sup>[33]</sup>.

## 10.3 Trading performance and integrity

Liberator configuration items are used to:

- ◆ Optimize client reconnection time.
- ◆ Reduce the performance impact of high update rates through throttling.

### Optimizing client reconnection time

If a client session becomes disconnected, Liberator will store update messages for the client until the client reconnects. This optimizes the reconnect time – on reconnection the stored updates are sent to the client as though the connection had not been lost.

The **output-queue-size** configuration item defines the maximum number of such messages that Liberator will store for each client. If the client reconnects after this limit has been reached, Liberator effectively discards the messages in the queue and instead sends the client the full image data for each subscribed object, which can take a significant time.

In a trading environment, the client is likely to be subscribed to a large number of instruments that are updating frequently, so the default **output-queue-size** value of 64 may be too small to optimize reconnection, even for transient connection losses. It is therefore suggested that **output-queue-size** be increased to **512** messages.

### Throttling

Liberator's throttling feature is a mechanism for reducing the performance impact of high update rates – see the appendix [Overview of throttling](#)<sup>[39]</sup>. For optimal performance it is highly desirable to throttle the instrument data updates that are streamed to clients, but:

**Note:** Messages on Trade Channels must **not** be throttled.

Trade messages must not be throttled because two trade messages sent within a single throttle period would be merged into one message. The resulting behavior would be undefined – for example, the Trade could fail or could execute incorrectly.

To meet these two requirements, throttling instrument data updates but not throttling trade messages, define separate Liberator objects for the instrument data subscriptions and for the messaging subscriptions. You can then set separate throttle levels for the messages relating to these two types of subscription.

**Example of throttle settings in the Caplin Trader Reference Implementation:**

```
add-object
  name      /FX
  type      20
  only-changed-fields
  throttle-times 0.25 2
end-object

add-object
  name      /FI
  type      20
  only-changed-fields
  throttle-times 0.25 2
end-object

...

add-object
  name      /FT/TRADE
  type      20
  throttle-times 0
  discard-timeout 0
end-object

...
```

This configuration shows that streamed data about FX and FI instruments is sent to clients as updates to the subscriptions in the Liberator directories /FX and /FI respectively. For example, updates to the FX currency pair EURUSD would be sent as messages with the subject name /FX/EURUSD.

Both FX and FI updates are subject to throttling (*throttle-times 0.25 2*), where the default throttle time is 0.25 seconds (the first entry in the list). The configuration parameter *only-changed-fields* ensures that just the changed fields in an update are forwarded to the client, thus optimizing the performance of the real time instrument display.

The **add-object** configuration item for trade messaging (*name /FT/TRADE*) has different parameter settings:

- ◆ *throttle-times 0* ensures that the objects used for trade messaging (*/FT/TRADE/...*) do not have throttling enabled.

**Note:** The throttle time *must* be explicitly set to zero here, so as to override any globally defined throttle time (configuration item **object-throttle-times**).

- ◆ The *only-changed-fields* parameter is false (by default), so all trade messages are sent in their entirety.

This is a performance optimization. The additional processing needed on Liberator to send the client all fields in each trade message is less than the processing needed to reconstruct messages in the client if the Liberator only sent the updated fields.

- ◆ *discard-timeout 0* ensures that Liberator removes a trade message subscription from its cache as soon as the Trade Channel is closed (the user has logged off Liberator and has therefore unsubscribed from /FT/TRADE/).



## 10.4 Security

Follow these recommendations to maximize the security of Trades.

Also see [Preventing identity theft](#)<sup>[34]</sup> in [Mapping trade messaging objects in Liberator](#)<sup>[33]</sup>.

### Session IDs

When a client connects to Liberator, the Liberator generates a unique session ID. This identifies the session in subsequent (RTTP) message exchanges between the client and Liberator across this connection. The client includes the session ID in every message sent to Liberator. The session ID is generated using a cryptographically secure pseudo-random number generator, the major feature of such a generator being that is difficult for a third party to predict its output by observing its previous outputs.

The Liberator configuration item **session-id-len** defines the length in characters of the unique identifier for a session. Its default value is 12, for backwards compatibility with older versions of Liberator and client StreamLink libraries, where the length of the session identifier cannot be changed.

**Note:** It is recommended that in Caplin Trader installations **session-id-len** be increased to **22** characters. This makes it extremely unlikely that a third party could successfully guess a session ID so as to impersonate a legitimate end-user.

### Single sign-on and KeyMaster

If you have a single sign-on system in place, Caplin recommends using the Caplin KeyMaster product in conjunction with this system, so that Caplin Trader users can access Liberator in a secure manner through the single sign-on. For more information, see the **Caplin KeyMaster Overview**.

## 11 Appendix: Overview of throttling

In a fast moving market, data updates can be generated more frequently than an end user can actually notice. For example, a user does not actually need to see every update if an item updates ten times in a second. Additionally, system resources and performance can be adversely impacted by such high update rates being fed through to client applications.

Liberator's throttling feature is a mechanism for reducing the performance impact of high update rates. When throttling is enabled, Liberator accumulates all the updates for a data item during an interval called the throttle time. At the end of this interval Liberator sends just the latest updated values of the item to the subscribed clients. The rate at which updates are sent to clients is therefore reduced or "throttled".

The following table shows how throttling works.

Consider a data item containing the bid and ask prices for stock in company ABC. Clients are subscribed to the item /ABC consisting two fields "bid" and "ask" (the bid price and the ask price). The table shows the succession of field updates that Liberator receives from the price feed via a DataSource adapter, the values that it holds in its cache, and the throttled updates that are actually sent to clients subscribed to /ABC.

Time	Updates received by Liberator	/ABC in Liberator cache		Updates sent to clients
		bid=	ask=	
T0 Start of throttle period 1		51.160	52.032	
T1	bid=51.162, ask=52.037	51.162	52.037	
T2	bid=51.155	51.155	52.037	
T3	bid=51.158	51.158	52.037	
T4 End of throttle period 1		51.158	52.037	bid=51.158, ask=52.037
T4 Start of throttle period 2		51.158	52.037	
T5	ask=52.041	51.158	51.041	
T6	ask=52.040	51.158	51.040	
T7 End of throttle period 2		51.158	51.040	ask=51.040

During the first throttle period the bid price changes three times (at times T1, T2, and T3) and the ask price changes just once (at time T1). At the end of this throttle period (time T4) Liberator sends just the most recently updated bid and ask prices to the clients, so the clients do not see the bid price updates at times T1 and T2.

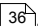
During the second throttle period the ask price changes twice (at times T5 and T6), but the bid price does not change at all, so at the end of this period (time T7) Liberator sends only the most recently updated ask price to the clients.

If an item is not updated at all during a throttle period, the very next time the item is updated in any subsequent throttle period Liberator immediately sends the updated values to its subscribed clients. This ensures that throttling does not introduce unnecessary delay in propagating updates to clients.

Throttling can be configured per data item; this allows all the items in a particular directory or even an individual item in a directory to be throttled by specific amounts.

Liberator can supply the same item to multiple users at different throttle levels. This allows users who need to view a large number of items, but who have low specification computers or slow network connections to the server, to receive data at a speed that suits their environment.

Client applications can change the level of throttling for specified items, groups of items, or all items globally.

**Note:** Messages on Trade Channels must **not** be throttled; see [Throttling](#) .

**Tip:** In some older Caplin documents throttling is sometimes called conflation.

## 12 Glossary of terms and acronyms

This section contains a glossary of terms, abbreviations, and acronyms used in this document.

Term	Definition
<b>API</b>	<u>A</u> pplication <u>P</u> rogramming <u>I</u> nterface
<b>Caplin Platform</b>	The Caplin Platform is a suite of software products for on-line financial trading and Web delivery of real-time market data.
<b>Caplin Trader</b>	Caplin Trader is a complete platform and toolkit for building multi-product trading portals. It is built on the <b>Caplin Platform</b> .
<b>Caplin Trader Client</b>	Caplin Trader Client is a platform neutral Web application that provides a rich trading workstation in a browser. It is based on an Ajax framework called <b>webcentric</b> , into which you can place any Web content created in HTML, Ajax, Adobe Flex™, or any other similar technology.
<b>Blotter</b>	A record of the details of Trades made by an end-user of Caplin Trader. In Caplin Trader Client the blotter is displayed in a dedicated panel and is dynamically updated as the user progresses through a Trade.
<b>Blotter Channel</b>	See <a href="#">Blotter Channels</a> <sup>[7]</sup> .
<b>Data service</b>	An abstraction layer within a <b>DataSource</b> that allows the DataSource to request objects, based on their subject names, without needing to know the specific DataSources from which the objects originate. A data service is defined through configuration. For more information see the <b>Caplin DataSource Overview</b> .
<b>DataSource</b>	DataSources are software adapters within the <b>Caplin Platform</b> that connect the Platform to external sources of real time data and external <b>trading systems</b> . In other Caplin documents DataSources are also called DataSource adapters.
<b>ESP</b>	<u>E</u> xecutable <u>S</u> teaming <u>P</u> rice <b>T</b> rade <b>M</b> odel.
<b>Liberator</b>	Caplin Liberator is a bidirectional streaming push server designed to deliver market data and trade messages over any network that supports Web traffic.
<b>RFQ</b>	<u>R</u> equest for <u>Q</u> uote <b>T</b> rade <b>M</b> odel.
<b>SDK</b>	<u>S</u> oftware <u>D</u> evelopment <u>K</u> it
<b>Throttling</b>	A <b>Liberator</b> feature for reducing the performance impact of high update rates for data streamed to clients. See <a href="#">Appendix: Overview of throttling</a> <sup>[39]</sup> .
<b>Trade</b>	In this document the term “Trade” represents a single trade for a user. This could be an <b>RFQ</b> , an execution on a streaming price ( <b>ESP</b> ), or any other type of trade. See <a href="#">Trades</a> <sup>[6]</sup> .
<b>Trade Channel</b>	A single user’s communication between the client application and the <b>Trading DataSource</b> . See <a href="#">Trade Channels</a> <sup>[6]</sup> .
<b>Trade Event</b>	An action by a client or an event from the <b>trading system</b> . See <a href="#">Trade Events</a> <sup>[7]</sup> .

Term	Definition
<b>Trading DataSource</b>	The <b>DataSource</b> used to integrate <b>Caplin Trader</b> with a <b>trading system</b> .
<b>Trading DataSource API</b>	The <b>API</b> to the <b>Trading DataSource</b> that allows the DataSource to be integrated with a <b>trading system</b> .
<b>Trade Model</b>	A Trade Model represents a type of Trade, for example a Request for Quote ( <b>RFQ</b> ) or Executable Streaming Price ( <b>ESP</b> ). Trade Models consist of a number of states and transitions. See <a href="#">Trade Models</a> .
<b>Trading system</b>	The term used in this document to refer to systems that support trade capture.

# Index

## - A -

- Abbreviations, definitions 41
- Acronyms, definitions 41
- add-data-service
  - Liberator configuration example 35
  - Liberator configuration item 32
- add-object
  - discard-timeout parameter 36
  - Liberator configuration example 33
  - Liberator configuration item 32
  - throttle-times parameter 36
- add-peer
  - Liberator configuration example 35
  - Liberator configuration item 32
- add-source-group
  - Liberator configuration example 35
- API
  - C++ Trading DataSource, using 25
  - Java Trading DataSource, using 16
  - Trading DataSource, overview 4
  - Trading DataSource, reference documentation 2
- API (C++)
  - ChannelListener 28
  - ChannelListener interface 26
  - TradeEvent object 27
  - TradeListener interface 26
  - TradeListener object 27
  - TradingApplicationListener 25, 26
  - TradingDataSource 25
- API (Java)
  - ChannelListener 19
  - ChannelListener interface 17
  - TradeEvent object 18
  - TradeListener interface 17
  - TradeListener object 18
  - TradingApplicationListener 16, 17
  - TradingDataSource 16
- Architecture
  - of Caplin Trader 2
- asset class
  - handled by Trading DataSource 6

## - B -

- Blotter channel
  - definition of and example 7
  - handling in C++ 29
  - handling in Java 20
- BlotterTradeListener (C++)
  - registering with TradingApplicationListener 29
- BlotterTradeListener (Java)
  - registering with TradingApplicationListener 20
- BlotterTradeListener interface
  - implementing in C++ 30
  - implementing in Java 22

## - C -

- C++
  - API for Trading DataSource, examples 25
- Cache
  - in Liberator 39
- Caplin Trader
  - architecture of 2
- Caplin Trading DataSource
  - API reference documentation 2
- Channel
  - blotter channel 7
  - trade channel 6
- ChannelListener (C++)
  - notification of new trades 26
- ChannelListener (Java)
  - notification of new trades 17
- ChannelListener (C++)
  - adding to a trade channel 26
  - notifying channel closure 28
  - notifying trade closure 28
- ChannelListener (Java)
  - adding to a trade channel 17
  - notifying channel closure 19
  - notifying trade closure 19
- Configuration
  - of Liberator for trading 32
  - of RFQ trade model using XML 14
  - of trade model 4, 13
- Conflation 39

## - D -

Data service  
    in Liberator, for trading 35  
DataSource listener (C++)  
    handling trade messages 25  
DataSource listener (Java)  
    handling trade messages 16  
discard-timeout  
    parameter of add-object 36

## - E -

ESP Trade Model  
    in example package of Java Trading  
    DataSource kit 24  
    in Trading DataSource kit 13  
    state diagram 9  
Event  
    definition for trading 7  
    from trading system (C++) 26  
    from trading system (Java) 17  
    handling by listener object 16, 25  
    represented by TradeEvent object (C++)  
    27  
    represented by TradeEvent object (Java)  
    18  
Examples  
    closing a trade (C++) 28  
    closing a trade (Java) 19  
    closing a trade channel (C++) 28  
    closing a trade channel (Java) 19  
    creating a new trade channel (C++) 26  
    creating a new trade channel (Java) 17  
    creation of trade event (C++) 27  
    creation of trade event (Java) 18  
    implementation of ChannelListener (C++)  
    26  
    implementation of ChannelListener (Java)  
    17  
    implementation of TradeListener (C++)  
    27  
    implementation of TradeListener (Java)  
    18  
    trade model configuration (RFQ) 14  
    Trading DataSource example (C++) 31  
    Trading DataSource example (Java) 24  
Examples (C++)

    initialization of Trading DataSource 25  
Examples (Java)  
    initialization of Trading DataSource 16  
Executable Streaming Price (ESP)  
    example of trade model 6  
    state diagram 9

## - F -

Factory (C++)  
    for trading state machines 25  
Factory (Java)  
    for trading state machines 16  
Field  
    in trade event 7

## - G -

Glossary 41  
Guard 6

## - I -

Identity theft  
    preventing 33  
Integrity  
    of trading 36

## - J -

Java  
    API for Trading DataSource, examples  
    16  
    trademark 3

## - K -

KeyMaster 38

## - L -

Liberator  
    access through KeyMaster 38  
    cache 39  
    mapping trade messaging objects 33  
    optimizing client reconnection time 36

- Liberator
  - routing trade messages 35
  - session ID 38
  - throttling 36
- Liberator configuration for trading
  - summary of configuration items 32
- Listener interface
  - in trading system (C++) 26
  - in trading system (Java) 17
- Listener object
  - ChannelListener (C++) 26, 28
  - ChannelListener (Java) 17, 19
  - TradeListener (C++) 26, 27
  - TradeListener (Java) 17, 18
- Listener object (C++)
  - TradingApplicationListener 25
- Listener object (Java)
  - TradingApplicationListener 16

## - M -

- Message
  - encapsulated by TradeEvent (C++) 27
  - encapsulated by TradeEvent (Java) 18
- Messages
  - causing state transitions 6
  - custom 4
  - for trading 4
  - on blotter channel 7
  - on trade channels 6
  - raising trade event 7
  - transmitting RequestId and TradeId 6
- MsgType field
  - in trade event 7

## - N -

- Note: 39
  - General Liberator configuration 32
  - Setting for session-id-len 38
  - Throttling trade messages 36
  - Throttle-time setting for trade messages 36

## - O -

- object-map

- Liberator configuration example 33
- Liberator configuration item 32
- use in preventing identity theft 33
- ORD trade model
  - in Trading DataSource kit 13
  - state diagram 11
- Order (ORD)
  - state diagram 11
- output-queue-size
  - Liberator configuration – definition 36
  - Liberator configuration item 32

## - P -

- Performance
  - of trading 36

## - R -

- Readership 1
- Reconnection time
  - optimizing for client 36
- Reference Implementation Trading DataSource 24
- Request for Quote (RFQ)
  - example of trade model 6
  - state diagram 12
- Request for Stream (RFS)
  - state diagram 10
- RequestId 6
- RFQ Trade Model
  - example XML configuration 14
  - state diagram 12
- RFS Trade Model
  - in example package of Java Trading DataSource kit 24
  - in Trading DataSource kit 13
  - state diagram 10

## - S -

- Security
  - of Trades 38
- Session ID
  - in Liberator 38
  - size of 38
- session-id-len



- session-id-len
  - Liberator configuration item 32
  - recommended size 38
- Single sign-on system 38
- State
  - concept, in trade model 6
  - defining in XML configuration 14
- State machine 16, 25
- T -**
- Terms, glossary of 41
- throttle-times
  - parameter of add-object 36
- Throttling
  - overview 39
  - settings for trading 36
- timeout
  - in Trade Model on client 12
  - in Trade Model XML configuration 14
- Tip:
  - Location of reference Liberator configuration file 32
  - Use of the term “conflation” 39
- Trade
  - closing (C++) 28
  - closing (Java) 19
  - definition of 6
  - handling new Trade using ChannelListener (C++) 26
  - handling new trade using ChannelListener (Java) 17
  - security of 38
- Trade Channel
  - definition of 6
  - Liberator configuration 33
  - notifying closure (C++) 28
  - notifying closure (Java) 19
  - notifying creation (C++) 26
  - notifying creation (Java) 17
- Trade Event
  - definition of 7
  - handling by listener object 16, 25
- Trade messaging
  - configuration for routing to Trading DataSource 35
  - role of TradeEvent object (C++) 27
  - role of TradeEvent object (Java) 18
  - role of TradeListener object (C++) 27
  - role of TradeListener object (Java) 18
  - setting discard-timeout in Liberator 36
  - setting throttle-times in Liberator 36
- Trade messaging (C++)
  - handling via DataSource listeners 25
- Trade messaging (Java)
  - handling via DataSource listeners 16
- Trade messaging object
  - in Liberator 33
- Trade model
  - configuration overview 4
  - configuring 13
  - definition of 6
  - ESP state diagram 9
  - ORD state diagram 11
  - relationship to trade 6
  - relationship to Trading DataSource 4
  - RFQ state diagram 12
  - RFS state diagram 10
  - supported in Trading DataSource kit 13
- Trade object 17, 18, 26, 27
- Trade System
  - events raised by 18, 27
  - messages from TradeListener (C++) 27
  - messages from TradeListener (Java) 18
- TradeEvent object (C++) 27
- TradeEvent object (Java) 18
- TradeId 6
- TradeListener (Java)
  - custom 17
- TradeListener (C++)
  - custom 26
  - handling trade events 27
- TradeListener (Java)
  - handling trade event 18
- Trading DataSource
  - API overview 4
  - example of Liberator configuration for 35
  - overview and architecture 4
  - relationship to trade models 4
  - standard trade models in kit 13
  - using the C++ API 25
  - using the Java API 16
- Trading GUI
  - overview 4
- Trading integrity 36
- Trading performance 36

- Trading state machine (C++) 25
- Trading state machine (Java) 16
- Trading System
  - cancelling a trade (C++) 28
  - cancelling a trade (Java) 19
  - events from 7
  - events from (C++) 26
  - events from (Java) 17
  - initialization within (C++) 26
  - initialization within (Java) 17
  - listener interface (C++) 26
  - listener interface (Java) 17
  - overview 4
  - processing events 6
- TradingApplicationListener (C++)
  - notifying when trade channel created 26
- TradingApplicationListener (C++)
  - example 25
  - registering BlotterTradeListener 29
- TradingApplicationListener (Java)
  - example 16
  - notifying when trade channel created 17
  - registering BlotterTradeListener 20
- TradingDataSource (C++)
  - creating in custom application 25
- TradingDataSource (Java)
  - creating in custom application 16
- Transition
  - concept, in trade model 6
  - defining in XML configuration 14
  - of trade state 6

**- V -**

- Value
  - in trade event 7

**- X -**

- XML configuration
  - example for RFQ trade model 14
  - in TradingDataSource kit 13
  - overview 4

## Contact Us

Caplin Systems Ltd  
Triton Court  
14 Finsbury Square  
London EC2A 1BR  
Telephone: +44 20 7826 9600  
Fax: +44 20 7826 9610  
[www.caplin.com](http://www.caplin.com)

The information contained in this publication is subject to UK, US and international copyright laws and treaties and all rights are reserved. No part of this publication may be reproduced or transmitted in any form or by any means without the written authorization of an Officer of Caplin Systems Limited.

Various Caplin technologies described in this document are the subject of patent applications. All trademarks, company names, logos and service marks/names ("Marks") displayed in this publication are the property of Caplin or other third parties and may be registered trademarks. You are not permitted to use any Mark without the prior written consent of Caplin or the owner of that Mark.

This publication is provided "as is" without warranty of any kind, either express or implied, including, but not limited to, warranties of merchantability, fitness for a particular purpose, or non-infringement.

This publication could include technical inaccuracies or typographical errors and is subject to change without notice. Changes are periodically added to the information herein; these changes will be incorporated in new editions of this publication.

Caplin Systems Limited may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time.

This publication may contain links to third-party web sites; Caplin Systems Limited is not responsible for the content of such sites.