

CAPLIN

StreamLink 5.0

StreamLink Configuration XML Reference

October 2011

CONFIDENTIAL

Contents

1	Preface.....	1
1.1	What this document contains.....	1
	About Caplin document formats	1
1.2	Who should read this document.....	1
1.3	Related documents.....	1
1.4	Typographical conventions	2
1.5	Feedback.....	2
1.6	Acknowledgments.....	2
2	Getting started.....	3
2.1	Technical assumptions and restrictions.....	3
2.2	Introduction to the configuration of StreamLink	3
2.3	Ordering and nesting of tags.....	5
3	XML Reference information.....	6
3.1	<Connections>.....	6
3.2	<FileLoggerConfiguration>.....	6
3.3	<LogFileName>.....	6
3.4	<LogFileNameSuffix>	7
3.5	<MaxLogFileSize>.....	7
3.6	<RttpConfiguration>.....	7
3.7	<RttpService>	8
3.8	<Server>	9
3.9	<ServerGroup>	10
3.10	<Type1Connection>.....	10
3.11	<Type1SSLConnection>	11
3.12	<Type2Connection>.....	11
3.13	<Type3Connection>.....	11

1 Preface

1.1 What this document contains

This reference document describes the XML-based configuration of StreamLink.

The information in this document applies to StreamLink version 5.0.

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 System Administrators and Software Developers who need to configure StreamLink.

1.3 Related documents

- ◆ **Caplin StreamLink Overview**

This document provides a technical overview of Caplin StreamLink.

1.4 Typographical conventions

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

Type	Uses
aMethod	Function or method name
<i>aParameter</i>	Parameter or variable name
<i>/AFolder/Afile.txt</i>	File names, folders and directories
<div>Some code;</div>	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
Glossary term	Items that appear in the “Glossary of terms and acronyms”
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.

Information about the applicability of a section is enclosed in a box like this.
For example: “This section only applies to version 1.3 of the product.”

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.

Visit our feedback web page at <https://support.caplin.com/documentfeedback/>.

1.6 Acknowledgments

Adobe® Reader 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.

2 Getting started

2.1 Technical assumptions and restrictions

XML

The XML markup defined in this document conforms to XML version 1.0 and the XML schema version defined at <http://www.w3.org/2001/XMLSchema>.

2.2 Introduction to the configuration of StreamLink

StreamLink can be configured to be highly resilient to server and network failures by defining an XML-based failover strategy and load balancing policy. The configuration defines the Liberator servers that StreamLink can connect to, the type of RTP connections that can be made, and the order in which StreamLink will attempt each connection.

Server Groups

Liberator servers can be logically arranged in named server groups, where each group can be defined to be of type 'balance' or 'order'. A server group can define a list of named Liberator servers in addition to other server groups.

Server Groups of type 'balance'

StreamLink will attempt to connect to each sever or server group in a 'balance' server group in a random order. Since other StreamLink clients will also attempt to connect in a random order, connections to Liberators will be statistically balanced across the defined servers and server groups.

Server Groups of type 'order'

StreamLink will attempt to connect to each server or server group in an 'order' server group in the order in which they are defined. This allows you to define an ordered list of backup servers or server groups that StreamLink will attempt to connect to when a connection fails or is lost. This configuration provides a reconnect and failover strategy.

Example: a simple load balancing configuration

```
<ServerGroup Type="balance">
  <Server Name="ServerA" .../>
  <Server Name="ServerB" .../>
</ServerGroup>
```

This example configuration defines two Liberator servers inside a 'balance' server group. When StreamLink first tries to connect to a Liberator server, it randomly chooses to connect to either "ServerA" or "ServerB". Since other StreamLink clients will also attempt to connect in a random order, the connections will be statistically balanced between "ServerA" and "ServerB".

Example: a simple failover configuration

```
<ServerGroup Type="order">
  <Server Name="PrimaryServer" .../>
  <Server Name="BackupServer" .../>
</ServerGroup>
```

The above example configuration defines two Liberator servers inside an 'order' server group. When StreamLink first tries to connect to a Liberator server, it will attempt to connect to "PrimaryServer". If this connection fails, or succeeds and then subsequently fails, then StreamLink will attempt to connect to "BackupServer1".

Example: a more advanced configuration

```
<ServerGroup Type="balance">
  <!-- 1st server group -->
  <ServerGroup Type="order">
    <Server Name="PrimaryA" .../>
    <Server Name="BackupA1" .../>
    <Server Name="BackupA2" .../>
  </ServerGroup>

  <!-- 2nd server group -->
  <ServerGroup Type="order">
    <Server Name="PrimaryB" .../>
    <Server Name="BackupB1" .../>
    <Server Name="BackupB2" .../>
  </ServerGroup>

  <!-- 3rd server group -->
  <ServerGroup Type="order">
    <Server Name="PrimaryC" .../>
    <Server Name="BackupC1" .../>
    <Server Name="BackupC2" .../>
  </ServerGroup>
</ServerGroup>
```

This example configuration defines three 'order' server groups inside an outer 'balance' server group. When StreamLink first tries to connect to a Liberator server, it randomly chooses between '1st server group', '2nd server group', and '3rd server group'.


In each 'order' group a primary server and two backup servers are defined. Within an 'order' group ('1st server group' for example), StreamLink will first attempt to connect to the primary server ('PrimaryA'). If this attempt fails it will then try to connect to the first backup server ('BackupA1'), followed by the second backup server ('BackupA2'). If StreamLink cannot connect to any server within this 'order' group, then it will choose one of the other (untried) 'order' server groups at random, and attempt to connect to the servers in that group.

During failover, StreamLink will attempt to connect to the next server in the current 'order' group, followed by the servers in other 'order' groups. StreamLink keeps a record of server connection attempts, and repeats the failover sequence a configurable number of times (see the `ReconnectAttempts` attribute of the `<RttpService>` configuration tag).

2.3 Ordering and nesting of tags

Each top level tag is shown below, together with the child tags that it can contain.

Tip: Advanced users may wish to consult the XML Schema (*ServiceConfiguration.xsd*) for definitive information on the ordering and nesting of tags. This file is supplied with the StreamLink product.

For a description of each tag and its attributes, see the [XML Reference information](#)  section.

<RttpConfiguration>

This is the outermost tag.

```
<RttpConfiguration>
  <RttpService></RttpService> (one only)
  <FileLogger></FileLogger> (zero or one)
</RttpConfiguration>
```

<RttpService>

```
<RttpService>
  <Connections></Connections> (one only)
  <ServerGroup></ServerGroup> (one only)
</RttpService>
```

<FileLoggerConfiguration>

```
<FileLoggerConfiguration>
  <LogFileName></LogFileName> (one only)
  <MaxLogFileSize></MaxLogFileSize> (zero or one)
  <LogFileNameSuffix></LogFileNameSuffix> (zero or one)
</FileLoggerConfiguration>
```

<Connections>

```
<Connections> (Must contain at least one of the following tags)
  <Type1Connection /> (one only)
  <Type1SSLConnection /> (one only)
  <Type2Connection /> (one only)
  <Type3Connection /> (one only)
</Connections>
```

<ServerGroup>

```
<ServerGroup> (Must contain at least one of the following tags)
  <ServerGroup></ServerGroup>
  <Server />
</ServerGroup>
```

3 XML Reference information

The following sections describe the StreamLink Configuration XML tags. They are arranged in alphabetical order of tag name.

For each tag the attributes you can use within it are listed and described in a table. The “Req?” column indicates whether the attribute is always required (“Y”) or is optional (“N”). Most attributes are optional. If you do not supply an optional attribute within an instance of the tag then the runtime behavior will be according to the default value of the attribute.

3.1 <Connections>

<Connections>

A list of RTTP connection types that StreamLink will try when it attempts to connect to a Liberator server. Each connection type is tried in the order in which they are defined. Valid connection types are:

- ◆ Type 1 (<Type1Connection>).
- ◆ Type 1 with SSL encryption (<Type1SSLConnection>).
- ◆ Type 2 (<Type1Connection>).
- ◆ Type 3 (<Type3Connection>).

Attributes: This tag has no attributes.

3.2 <FileLoggerConfiguration>

<FileLoggerConfiguration>

When this tag is present, StreamLink will log the status of Liberator connections and other connection information to a log file (see <LogFileName>). A log file can help support staff diagnose connection problems that may arise between StreamLink and Liberator. If this tag is present, the child tag <LogFileName> must also be provided. If this tag is not present, Liberator will not create a log file.

Attributes: This tag has no attributes.

3.3 <LogFileName>

<LogFileName>

The name of the log file (see <FileLoggerConfiguration>).

Attributes: This tag has no attributes.

3.4 <LogFileNameSuffix>

<LogFileNameSuffix>

The suffix to apply to a log file when it is renamed for archiving (see <MaxLogFileSize>). If the suffix starts with a '.', then it is interpreted as a fixed suffix and is simply appended to the log file name to make the name of the old (archive) log file. For example, if the name of the log file is 'ABCLogFile.log' and the suffix is '.bak', the log file is renamed as 'ABCLogFile.log.bak'. If the suffix does not start with a '.', it is interpreted as a date/time pattern, which must be of the form 'MMddyyyyHHmmss'. For example, if the suffix is 'HHmm' and the log file was renamed at 08:20, its new name is 'ABCLogFile.log.0820'. If this tag is not supplied, the default suffix ".old" is applied; for example 'ABCLogFile.log' is renamed as 'ABCLogFile.log.old'.

Attributes: This tag has no attributes.

3.5 <MaxLogFileSize>

<MaxLogFileSize>

The maximum size of the log file in bytes (see <FileLoggerConfiguration>). When the log file reaches this size it is renamed for archiving (see <LogFileNameSuffix>) and a new log file is opened with the name as defined in <LogFileName>. If this tag is not supplied, the default log file size of 5,000,000 bytes is used.

Attributes: This tag has no attributes.

3.6 <RttpConfiguration>

<RttpConfiguration>

The outermost XML tag that configures the connection, failover, and load balancing policies that StreamLink will follow when it attempts to connect to a Liberator server.

Attributes: This tag has no attributes.

3.7 <RttpService>

<RttpService>

A named service defining the connection, failover, and load balancing policies that StreamLink follows when it attempts to connect to a Liberator server. The service identifies a named group of Liberator servers (<ServerGroup>) and the type of connections (<Connections>) that StreamLink can make to these servers.

Attributes:

Name	Type	Default	Req?	Description
HeartbeatInterval	integer	5000	N	The interval between heartbeat latency messages (in milliseconds). StreamLink can request a regular and time stamped heartbeat latency message from Liberator. StreamLink uses this timestamp, together with the clock offset between Liberator and StreamLink (see the 'SyncSize' attribute), to calculate one-way message latency. StreamLink stores the message latency and makes it available to the application through the StreamLink API. Set to 0 to disable heartbeat latency messages.
Name	string	(none)	Y	The name of the service.
NoopInterval	integer	10	N	The interval between NOOP commands (in seconds). StreamLink sends NOOP commands periodically to verify that the connection to Liberator is still open. A value of zero or less will disable NOOP commands. NOOP commands must be disabled when you are debugging code, because NOOP events can occur as you step through the code.
NoopTimeout	integer	5	N	The time (in seconds) to wait for a response to a NOOP command before StreamLink decides that the connection to Liberator has been lost.
ReconnectAttempts	integer	-1	N	StreamLink attempts to connect to each of the Liberator servers in the named server group in turn. If it does not succeed in establishing a connection, StreamLink repeats the connection sequence up to a maximum of 'ReconnectAttempts' times. A value of -1 means that there is no limit to the number of reconnection attempts.
ReconnectInterval	integer	5	N	The interval, in seconds, between connection attempts (see the 'ReconnectAttempts' attribute).
SyncInitialSpacing	integer	500	N	The time in milliseconds between successive sync messages, until the first 'SyncSize' messages have been sent to Liberator.

Name	Type	Default	Req?	Description
SyncSize	integer	12	N	The number of sync messages that StreamLink uses to calculate the clock offset between Liberator and StreamLink. StreamLink periodically sends time stamped sync messages to Liberator. Liberator adds its own timestamp to each sync message it receives, and immediately returns the sync message to StreamLink. StreamLink uses the sync message with the quickest round trip time from the last 'SyncSize' messages to calculate the clock offset between Liberator and StreamLink. Set the value of 'SyncInitialSpacing' to be less than the value of 'SyncSpacing', to reduce the time it takes to send and receive the first 'SyncSize' messages. StreamLink stores the clock offset and makes it available to the application through the StreamLink API.
SyncSpacing	integer	5000	N	The time in milliseconds between successive sync messages, after the first 'SyncSize' messages have been sent to Liberator. Set to 0 to disable sync messages. If sync messages are disabled, StreamLink will not calculate the clock offset between Liberator and StreamLink (see 'SyncSize' and 'HeartbeatInterval').

3.8 <Server>

<Server>

Defines a single Liberator server that StreamLink can connect to.

Attributes:

Name	Type	Default	Req?	Description
Address	string	(none)	Y	The address of the server. This can either be a hostname (for example myliberator.example.com) or an IP address (for example 192.168.226.1).
Name	string	(none)	Y	The name that identifies the server in this XML configuration.

3.9 <ServerGroup>

<ServerGroup>

The Liberator servers that StreamLink can connect to as part of the failover and load balancing policy. A server group can define a list of Liberator servers (<Server>) in addition to other server groups (<ServerGroup>). StreamLink attempts to connect to these servers and server groups in accordance with the value of the Type attribute.

Attributes:

Name	Type	Default	Req?	Description
Type	string	(none)	Y	Defines the sequence in which embedded Servers/ServerGroups are processed. Valid values are: balance - Randomly choose between the remaining (untried) servers and/or server groups, as specified in this <ServerGroup> tag. This setting can be used to balance the loading across the available servers. order - Connect to each server and/or server group specified in this <ServerGroup> tag in the order specified. This setting can be used to define the failover order of connection attempts.

3.10 <Type1Connection>

<Type1Connection>

Defines an unencrypted type 1 (direct TCP/IP socket) RTTP connection to the server. Type 1 connections are not available in StreamLink for Silverlight.

Attributes:

Name	Type	Default	Req?	Description
Port	string	(none)	Y	The port number that Liberator will listen on for StreamLink connections.

3.11 <Type1SSLConnection>

<Type1SSLConnection>

Defines a type 1 (direct TCP/IP socket) RTTP connection to the server that is encrypted with SSL. Type 1 SSL connections are only available in StreamLink.NET and are only supported by more recent versions of Liberator. For more information contact Caplin Support.

Attributes:

Name	Type	Default	Req?	Description
Port	string	(none)	Y	The port number that Liberator will listen on for StreamLink connections.

3.12 <Type2Connection>

<Type2Connection>

Defines a type 2 (streaming HTTP/HTTPS) RTTP connection to the server.

Attributes:

Name	Type	Default	Req?	Description
Port	string	(none)	Y	The port number that Liberator will listen on for StreamLink connections.
Protocol	string	(none)	Y	The connection protocol. Valid values are HTTP and HTTPS (case insensitive).

3.13 <Type3Connection>

<Type3Connection>

Defines a type 3 (polling HTTP/HTTPS) RTTP connection to the server.

Attributes:

Name	Type	Default	Req?	Description
Port	string	(none)	Y	The port number that Liberator will listen on for StreamLink connections.
Protocol	string	(none)	Y	The connection protocol. Valid values are HTTP and HTTPS (case insensitive).

Contact Us

Caplin Systems Ltd
Cutlers Court
115 Houndsditch
London EC3A 7BR
Telephone: +44 20 7826 9600
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.