

CAPLIN

Liberator 4.5

Benchmarks

January 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	2
1.4	Feedback	2
1.5	Acknowledgments	2
2	Overview	3
2.1	About the benchmark tests	3
	Message latency versus CPU usage	3
	Test setup	3
2.2	Headline results	4
2.3	Caplin's benchmark tools	5
3	Test scenarios	7
3.1	Standard Liberator tests	7
	Low updates	7
	Medium updates	8
	High updates	8
	Very high updates	9
3.2	JMX tests	9
4	Standard Liberator Tests – Results	10
4.1	Low updates	11
4.2	Medium updates	13
4.3	High updates	15
4.4	Very High Updates	17
4.5	Very high updates - batching	19
4.6	Message sizes	21
5	Liberator JMX Tests – Results	22
5.1	Low updates	23
5.2	High updates	26
6	How Caplin's benchmark tests were conducted	27
6.1	Test method	27
	Approach	27
	Test setup	27

6.2	Test configurations	29
6.3	Test software	29
	Caplin Liberator server	29
	Operating system.....	30
	Test DataSource application (Benchsrc)	30
	Test RTTP client application (Benchrttp)	30
6.4	Test hardware.....	30
	Main Server.....	30
	Main Harness.....	31
	Client Harnesses	31
6.5	The network	31
7	Frequently asked questions.....	32
7.1	What burst configuration should we use?	32
7.2	How many threads should we configure?	32
	DataSource threads.....	33
	Session threads	33
	CPU resource considerations	33
7.3	How do message sizes affect performance?	34
7.4	How much bandwidth will our Liberator use?	34
7.5	How many subscriptions can Liberator handle?.....	34
7.6	How much disk space will our Liberator need?	34
8	Glossary of terms and acronyms	35

1 Preface

1.1 What this document contains

This document details the results of a set of performance benchmark tests carried out on Caplin Liberator 4.5. It is hoped that the information provided in this report will assist customers in production capacity planning when deploying Liberator 4.5.

- ◆ Section 2.1 on page 3 gives a summary of the tests performed by Caplin Systems.
- ◆ Section 2.2 on page 4 summarizes the main results of the tests relating to typical performance profiles for Caplin Liberator.
- ◆ Section 3 on page 7 describes the test scenarios.
- ◆ Section 4 on page 10 contains the results, with performance graphs, of the standard tests.
- ◆ Section 5 on page 22 contains the results of tests that show how JMX™ monitoring clients connected to Liberator affect the latency of the messages Liberator sends to the standard clients.
- ◆ Section 6 on page 27 gives detailed information on how the benchmark tests were conducted.
- ◆ Section 7 on page 32 addresses frequently asked questions about how to configure and tune Liberator and its environment to achieve the required performance.

About Caplin document formats

This document is supplied in Portable document format (*.PDF* file), which you can read on-line using a suitable PDF reader such as Adobe Reader®. The document is formatted as a printable manual; you can print it from the PDF reader.

1.2 Who should read this document

This document is intended for anyone who is evaluating Caplin Liberator's performance characteristics, or who is planning to deploy Caplin Liberator. Typical readers would be:

- ◆ Technical Managers
- ◆ System Architects
- ◆ System Administrators

1.3 Related documents

◆ Caplin Xaqua: Overview

A business and technical overview of Caplin Xaqua.

◆ Caplin Liberator 4.5 Administration Guide

Describes the Caplin Liberator server and its place within Caplin Xaqua.

Explains how to install, configure, and manage the Liberator.

Includes configuration reference information, and a list of Liberator's log and debug messages.

1.4 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/>.

1.5 Acknowledgments

Adobe Reader is a registered trademark of Adobe Systems Incorporated in the United States and/or other countries.

AMD and *Opteron* are trademarks of Advanced Micro Devices, Inc.

Dell and *PowerEdge* are trademarks of Dell Inc in the United States and other countries.

Intel and *Intel Xeon* are registered trademarks of Intel Corporation in the U.S. and other countries.

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

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Enterprise Linux is a registered trademark of Red Hat, Inc. in the United States and other countries.

2 Overview

2.1 About the benchmark tests

The benchmark tests detailed in this document are designed to show how Caplin Liberator will perform on the Linux® platform, when deployed as real-time financial internet hub, streaming data updates to web-based financial trading applications.

The tests cover:

- ◆ Low update rates, typical of a low-end information portal
- ◆ Medium update rates, typical of a low-volatility trading system, such as credit trading
- ◆ High update rates, typical of a high-activity trading system such as FX
- ◆ Very high update rates, representing the most extreme online trading requirements

The main factor affecting the overall performance of Liberator is the power of the machine on which it runs. The tests were conducted on servers representing typical commercially available machines that can be used to host Web servers and server applications. A single Liberator instance was run on one machine, while test harnesses were run on other machines to provide data and client processes.

Message latency versus CPU usage

The key item measured in the tests was the end-to-end message latency against the number of logged in clients, and by implication the number of update messages being sent out to the totality of the connected client base.

Although some of the test results show CPU usage, in practice, end-to-end message latency is more significant as a measure of Liberator performance than CPU usage. Message latency has a direct impact on users and may increase long before CPU usage reaches its maximum. The aim of sizing a system incorporating Caplin Liberator should be to achieve a maximum desired message latency for a given maximum update rate.

Test setup

For detailed information on the test set up used at Caplin Systems, see section 6 "How Caplin's benchmark tests were conducted" on page 27.

Note: It is hoped that the information provided in this report will assist customers in production capacity planning. However, while the tests were designed to emulate real-world traffic and user scenarios, they were conducted using specific hardware running in an isolated environment, and therefore no guarantees can be made that identical results will be achieved in other environments.

2.2 Headline results

Details of the test scenarios, environments, full results and graphs can be seen in later sections, but here are some headline results for the tests run:

- ◆ 100,000 users each receive 1 message/second with a latency of only 6ms.
- ◆ 40,000 users each receive 10 messages/second with a latency of only 6ms.
- ◆ 9,000 users each receive 50 messages/second with a latency of only 5ms.
- ◆ 3,500 users each receive 100 messages/second with a latency of only 7ms, or
- ◆ 16,000 users each receive 100 messages/second with a latency of 60ms

2.3 Caplin's benchmark tools

Benchmarking a streaming server such as Liberator in a realistic manner is a challenge, because of the need to simulate the large numbers of users and high update rates that would be encountered in the real-world business environments where the server is typically deployed. Caplin Systems has produced a suite of tools, the Benchtools that make such benchmarks easier to set up and run.

The Benchtools suite was used to run the benchmark tests described in this document. These tools are also available for Caplin's customers to measure the performance of Caplin Liberator in their own environments.

The following diagram indicates how the Benchtool components integrate with Caplin Xaqua:

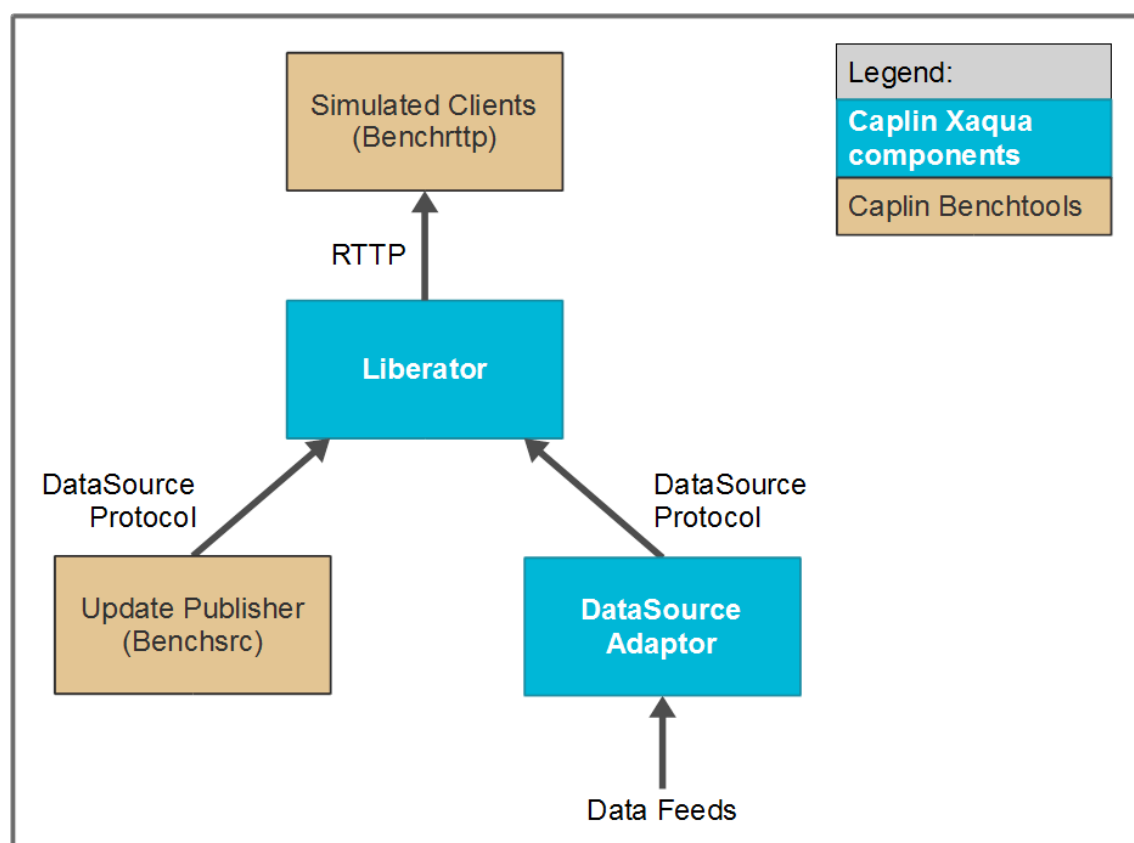


Figure 2–1: Benchtool components in Caplin Xaqua

The Benchtools package comprises two components:

- ◆ **Benchsrc**

Caplin's update publisher which generates updates of a configurable size and frequency.

- ◆ **Benchrttp**

Caplin's scalable client simulator which simulates concurrent RTTP sessions.

As with StreamLink clients, simulated clients use the RTTP protocol to establish a streaming connection to Liberator's HTTP port. Similar to other DataSources, Benchsrc sends updates to Liberator's DataSource port via the DataSource protocol.

The architecture shown in the diagram allows two test configurations:

1. Benchsrc generates updates.
2. Another publisher generates updates to a feed such as Reuters RMDS.

In both configurations Benchrttp clients consume the updates.

3 Test scenarios

The following sections describe the test scenarios used in these benchmarks. These scenarios are designed to simulate different types of activity and data rates, as commonly seen in real time financial applications, and they demonstrate the kind of performance that Liberator can achieve.

The exact numbers can vary significantly between different business scenarios, *and therefore Caplin always advises customers to run benchmarks that reflect their actual requirements and data update rates*. Caplin's benchmark tools make this easier to do; once the test environment is set up using these tools, it is easy to configure and test different scenarios. Benchmarks can either be run against a test back end, such as Benchsrc (which the tests described in this document use), or against real data supplied by the customer.

3.1 Standard Liberator tests

Each scenario consists of a source of data and a set of clients.

The source of data (benchsrc) is configured to supply a set of subjects, at a known update rate. Each client subscribes to a subset of these subjects, choosing at random. As the number of clients is increased, more of the source subjects are subscribed to, up to the point when all source subjects are subscribed to. In these tests, the point at which all source subjects are subscribed to is quite early in the test.

In all tests the message size is 54 bytes unless otherwise stated. This is a 5 field message, typical of financial applications.

Low updates

This scenario is a good base test, as it is typical of a low-end information portal.

Source subjects	1,000
Update rate per subject	0.5 updates/second
<i>Total source update rate</i>	<i>500 updates/second</i>
Subscriptions per client	2
<i>Update rate per client</i>	<i>1 update/second</i>

Medium updates

Compared to the Low updates scenario, this scenario has an increased update rate, increased number of subjects, and more subscriptions from each client. This level of subscriptions and data rates is typical of a low-volatility trading system, such as credit trading.

Source subjects	4,000
Update rate per subject	0.5 updates/second
<i>Total source update rate</i>	<i>2,000 updates/second</i>
Subscriptions per client	20
<i>Update rate per client</i>	<i>10 updates/second</i>

High updates

Compared to the Low and Medium updates scenarios, this scenario has an even higher update rate, number of subjects, and number of subscriptions from each client. This is typical of a high-activity trading system such as FX.

Source subjects	10,000
Update rate per subject	0.5 updates/second
<i>Total source update rate</i>	<i>5,000 updates/second</i>
Subscriptions per client	100
<i>Update rate per client</i>	<i>50 updates/second</i>

Note: The Medium updates and High updates scenarios have changed since previously published benchmarks. In both cases, the number of source subjects has been increased, which means the update rate received by Liberator is higher. The update rate received by clients is still the same however.

Very high updates

This scenario simulates a very high-end single dealer platform, where each client has a large number of fast moving instruments on their screen.

This is the most onerous scenario tested. Compared to the other scenarios, it has a much larger data set, a higher update rate per object, and a higher number of subscriptions per client. It represents the most extreme online trading requirements.

Source subjects	20,000
Update rate per subject	1 update/second
<i>Total source update rate</i>	<i>20,000 updates/second</i>
Subscriptions per client	100
<i>Update rate per client</i>	<i>100 updates/second</i>

3.2 JMX tests

The JMX test scenarios are the same as the standard Liberator scenarios above, but with one or more JMX clients also connected to Liberator. They are designed to see how the introduction of JMX clients to monitor various aspects of Liberator can impact performance.

The JMX clients only affect the tests indirectly. Data to and from them is not measured in any way, so the measurements are all the same as for the standard tests.

4 Standard Liberator Tests – Results

Each of the three scenarios below shows two graphs. The first graph is a plot of the mean latency as the number of clients increases. Each point on the graph represents the mean of all the messages received in a 30 second period. The second graph shows more detail: the blue plot is the latency range (the maximum and minimum latency received in that period), the white line embedded in the latency range shows the mean latency, and the darker line shows the CPU usage of the Liberator throughout the test.

4.1 Low updates

The low updates scenario is really a test of how many users the system can support. The update rate of 1 message/second is low and not very indicative of financial applications. Each client in this scenario receives 54 bytes/second.

The following graph shows Liberator supporting 100,000 clients, each receiving 1 message/second, with mean latency reaching nearly 7 milliseconds. At 10,000 clients the mean latency is a mere 1 millisecond. At 100,000 clients Liberator is publishing 5.4 Mbytes/second.




Figure 4–1: Low Updates – Mean Latency

The next graph shows more detail on the latency of messages in this test run. The shaded area shows the range of latencies; the interesting part of this is the top, which represents the maximum latency of any message in that period. We can see that even at 100,000 clients, no single message took longer than 40ms. For the majority of the test, the maximum latency is about double the mean.

The other part of the graph is the CPU usage. At 100,000 clients, the CPU usage of the machine is only 10%. This shows that the machine and Liberator have a lot more to give in this scenario. However, the test was only set up for 100,000 clients, and more client harness machines would be needed to take it much further.




Figure 4-2: Low Updates – Details

4.2 Medium updates

The medium updates scenario increases the update rate for each client to 10 messages/second; this could represent a low end financial application. Each client receives 540 bytes/second.

The following graph shows that this test run reached 40,000 clients, with a mean latency of less than 6 milliseconds. That is a total of 400,000 messages/second being published by Liberator, at a bandwidth of 21.6 Mbytes/second.




Figure 4–3: Medium Updates – Mean Latency

The following details graph again shows maximum latencies only reaching peaks of about 40ms, and much lower than this for the majority of the test run. CPU usage reaches 40% of the 12-core server on which Liberator is running.




Figure 4–4: Medium Updates – Details for High Updates

4.3 High updates

The high updates scenario is more typical of a trading application. Each user receives 50 messages/second, using a bandwidth of 2.7 Kbytes/second.

The following graph shows that under this scenario Liberator manages to serve nearly 9,000 clients at a mean latency of about 5 milliseconds. At the end of the test Liberator is publishing 450,000 messages/second using a bandwidth of 24.3 Mbytes/second.




Figure 4–5: High Updates – Mean Latency

The next graph shows that the maximum latency is pretty low throughout the test. The CPU usage is 45% of the machine's capacity.




Figure 4–6: High Updates – Details

4.4 Very High Updates

The very high updates scenario represents a high end trading system. There is a much higher update rate at the source of the data, due to the larger number of subjects available to subscribe to. Each client receives 100 messages/second, using a bandwidth of 5.4 Kbytes/second.

The following graph shows a maximum number of 3,500 clients, with a mean latency of 7 milliseconds, and the majority of the test run achieving 3-5 milliseconds latency.

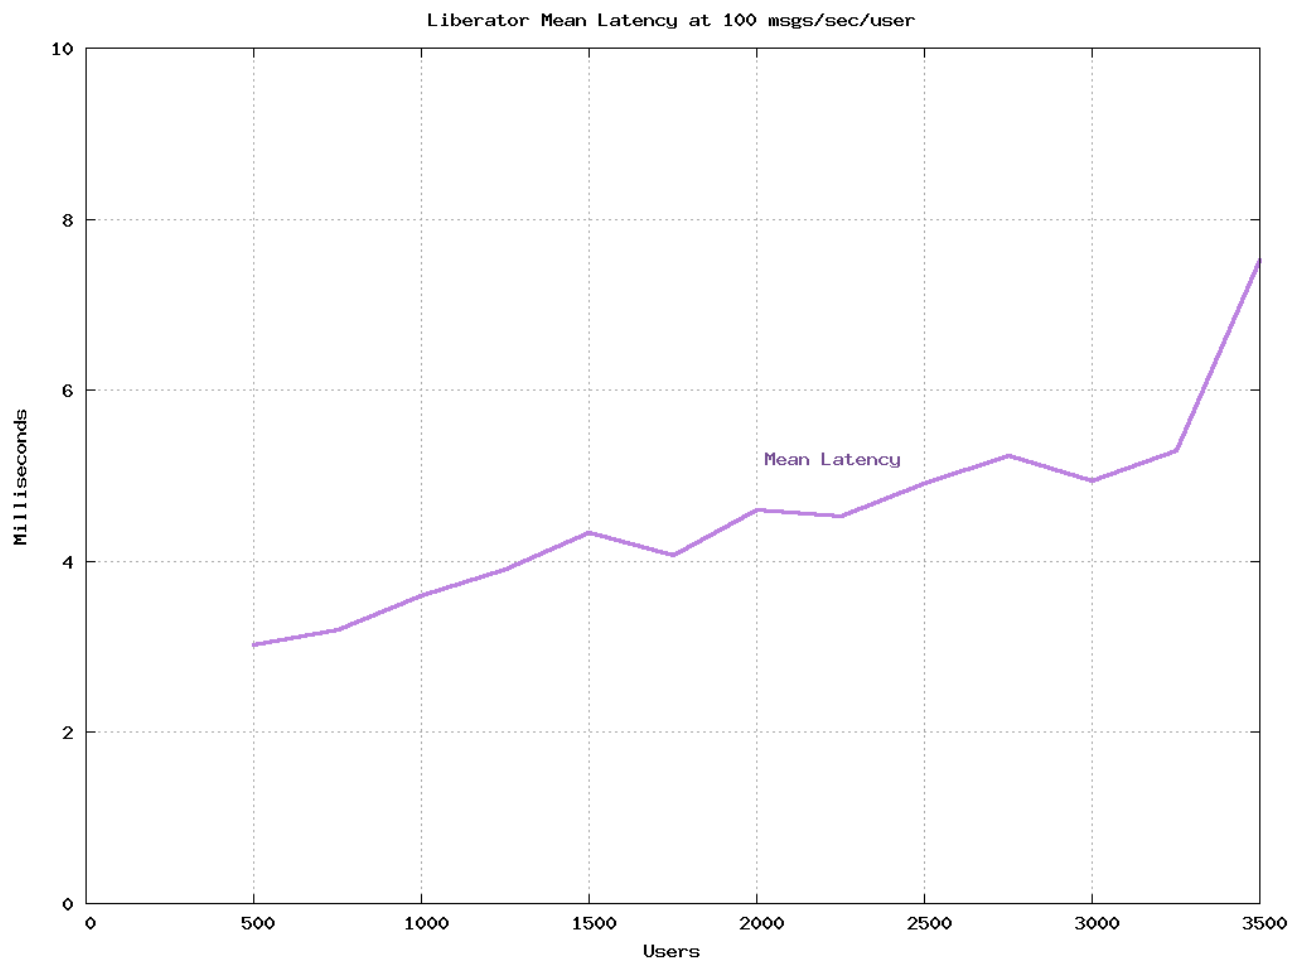


Figure 4–7: Very High Updates – Mean Latency

The next graph shows slightly higher maximum latencies, approaching 100 milliseconds at the end of the test run. CPU usage hits 60% at the 3,500 client mark.




Figure 4–8: Very High Updates – Details

4.5 Very high updates - batching

Liberator can be configured to batch messages together at the point they are sent to the client. This does not affect the data being sent, just how it is sent. Instead of sending a lot of small packets containing one message each, Liberator sends larger packets, less frequently, containing multiple messages.

The batching feature will only batch messages together when the update rate is over a configured amount, therefore it can be used to handle peak data rates. However, the scenarios tested here have a more uniform update rate so the batching is always active, when configured.

The following graph shows the mean latency of multiple test runs with different configurations for batching. The blue line shows the mean latency with no batching configured, this is the same run as the graphs above. The other three lines show how adding batching increases mean latency, as expected since we are delaying sending messages, but allows Liberator to support many more clients.

With 100 milliseconds batching, giving a mean latency of around 55 milliseconds, the test run reached over 16,000 clients receiving 100 messages/second each. In that configuration Liberator is publishing over 1.6 million messages per second which is 86.4 Mbytes/second, which could be approaching the limitations of the gigabit network.




Figure 4–9: Very High Updates – Batching

4.6 Message sizes

So far all the tests have been with messages of 54 bytes, this may seem like a small message, but it contains 5 fields, typical of a financial application. RTTP, the protocol used between Liberator and clients, is optimized to keep message sizes as small as possible.

The following graph shows a re-run of the high updates scenario, but with different sizes of message. The results for the original 54 byte message are plotted, along with runs using 108, 162 byte, and 216 byte messages.

The graph shows that as message size increases, latency is barely affected, but the ultimate number of clients Liberator can support is reduced.




Figure 4-4-10: Message Sizes

5 Liberator JMX Tests – Results

The JMX test results shown are intended as comparisons. The hardware used is not the same as for the standard Liberator tests and the test runs were not always complete. What the results do show is how monitoring different statistics, or connecting more JMX clients, affects the latency of the messages Liberator sends to the standard clients.

5.1 Low updates

The following graph shows three separate test runs of the low updates scenario. The graph compares the control run, which has no JMX clients connected, to two runs with a single JMX client connected. In one run, the JMX client monitored just DataSource peer statistics, in the other run the JMX client was monitoring all statistics.

You can see that in this scenario the JMX client does not introduce much latency, although a small, but clear, difference can be seen.




Figure 5–1: Liberator JMX Tests – Low Updates

The next graph shows the same three runs, but includes plots for the CPU usage of Liberator during the test. Again a clear difference can be seen, but if you look at the scale the tests only use about 1% CPU, so the difference is not significant.

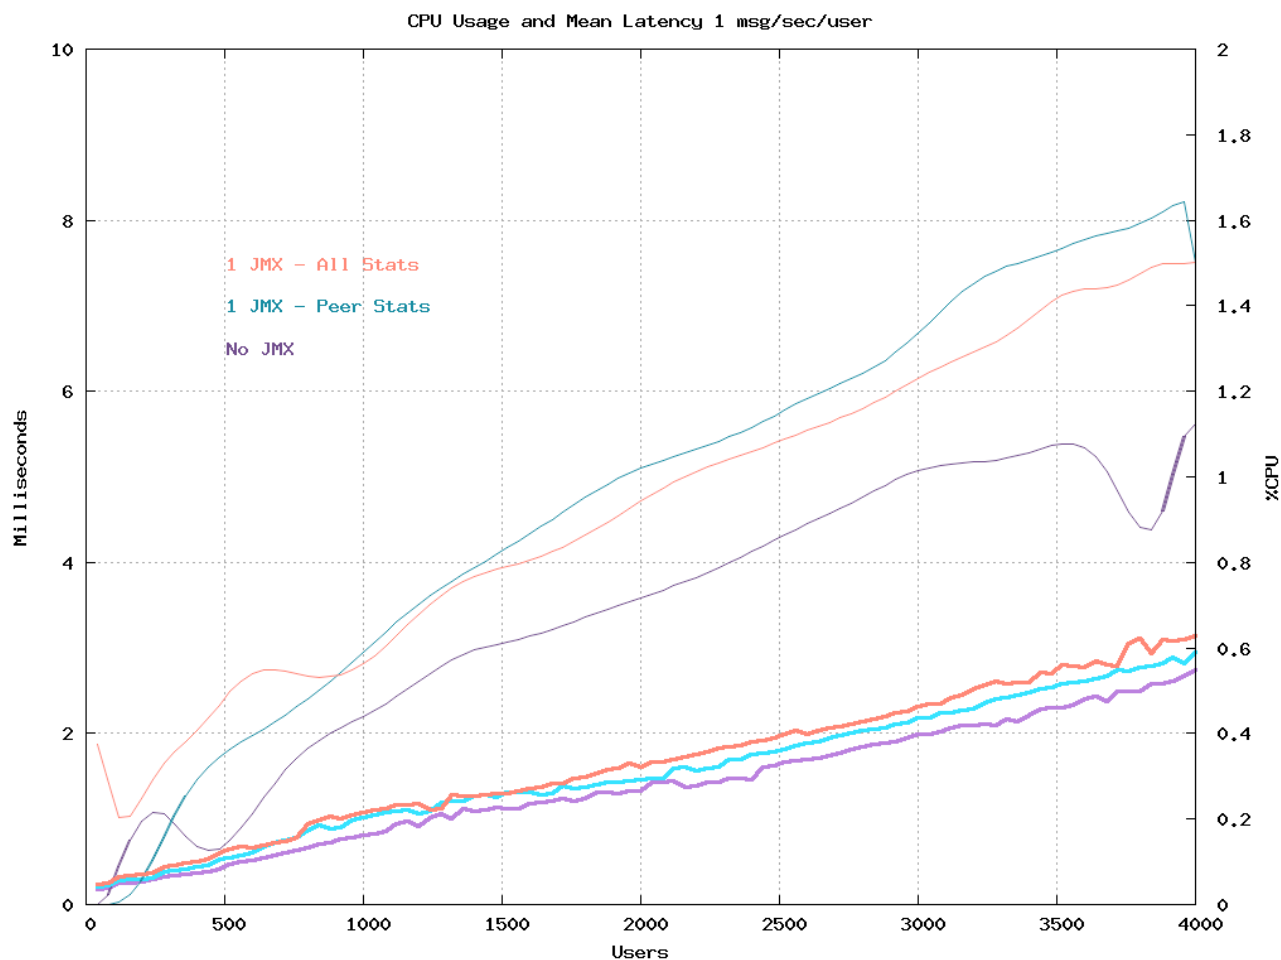


Figure 5–2: Liberators JMX Tests – Low Updates With CPU Usage

The next graph shows the effect of connecting 10 JMX clients instead of just 1. There is now a clearer difference in CPU usage, although it is still very low. It also shows that the mean latency is not greatly affected at these low update levels.

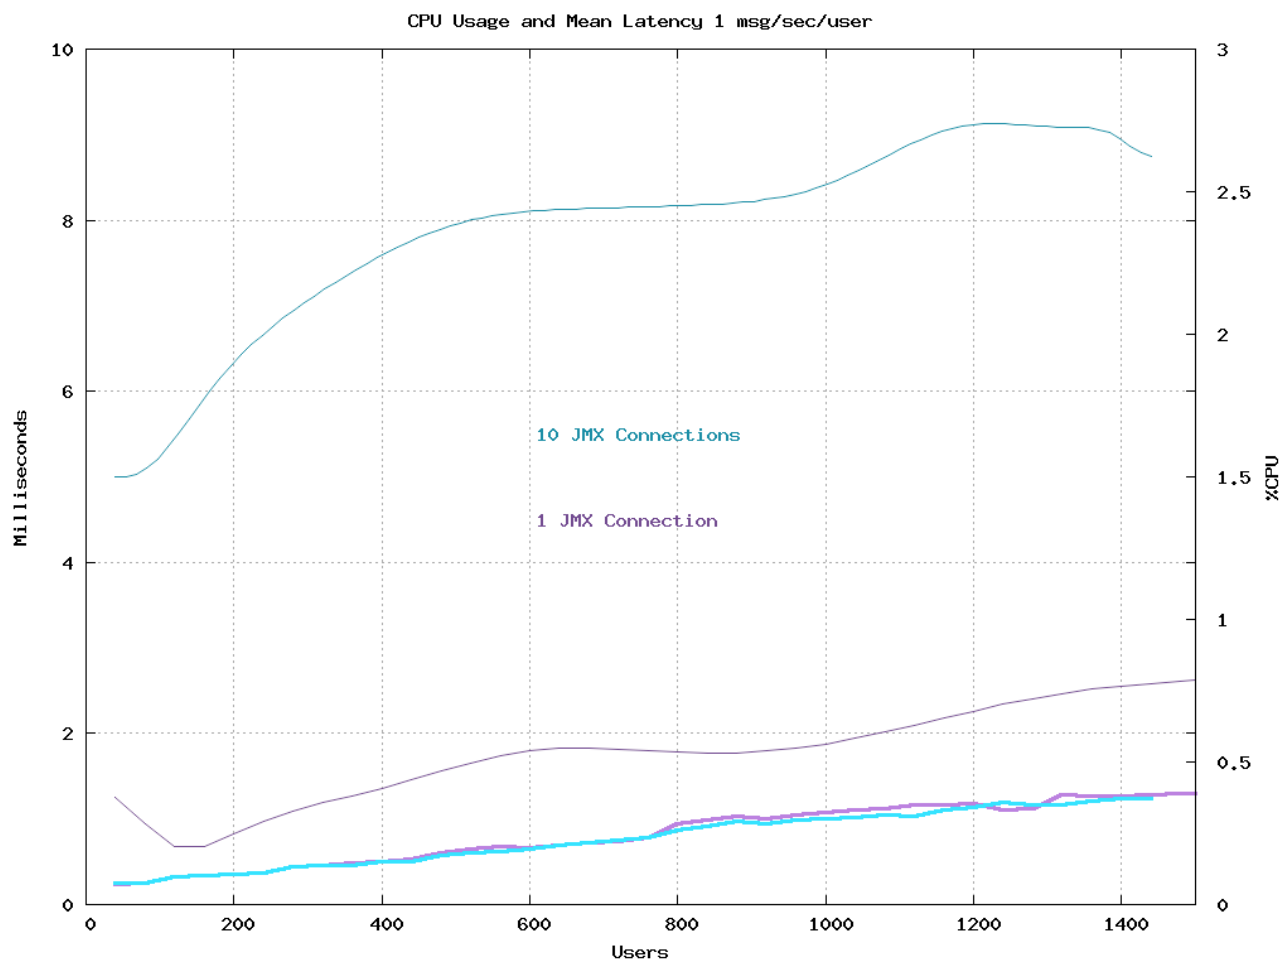


Figure 5–3: Liberator JMX Tests – Low Updates With 10 JMX Clients

5.2 High updates

The high updates scenario stresses the system a lot more. The following graph shows a control run with no JMX monitoring, a run with 1 JMX client, and a run with 10 JMX clients. Again, latency is not affected much, even at these higher update rates, but you can see from the CPU usage that Liberator has to work a lot harder to achieve this latency.

Looking at the 2000 client mark, the control run uses 21% CPU, the 1 JMX client run uses about 34%, and the 10 JMX client uses about 38%.




Figure 5–4: Liberator JMX Tests – High Updates

6 How Caplin's benchmark tests were conducted

The following sections describe the test method used, give information about the test configurations, and detail the test software, test hardware, and the network used.

6.1 Test method

Approach

Although the benchmark consisted of several different tests, they all followed a similar method.

Each test consisted of one or more DataSources publishing messages into a Liberator which pushed the messages out to a set of subscribing clients through RTTP connections. Each subscribed object was updated at a regular rate by the supplying DataSource. Additional clients were logged on to the Liberator throughout the test run to determine the effect of increasing the load on the Liberator.

Test setup

The multiple RTTP client connections were simulated using a specially written application called Benchrttp. The DataSource application supplying the Liberator (Benchsrc) was also specially written. Both Benchrttp and Benchsrc are controllable using a command protocol, thus allowing message rates and number of clients to be remotely managed using scripts.

The following diagram shows the hardware configuration used for the tests, and how the test software was distributed across the hardware.

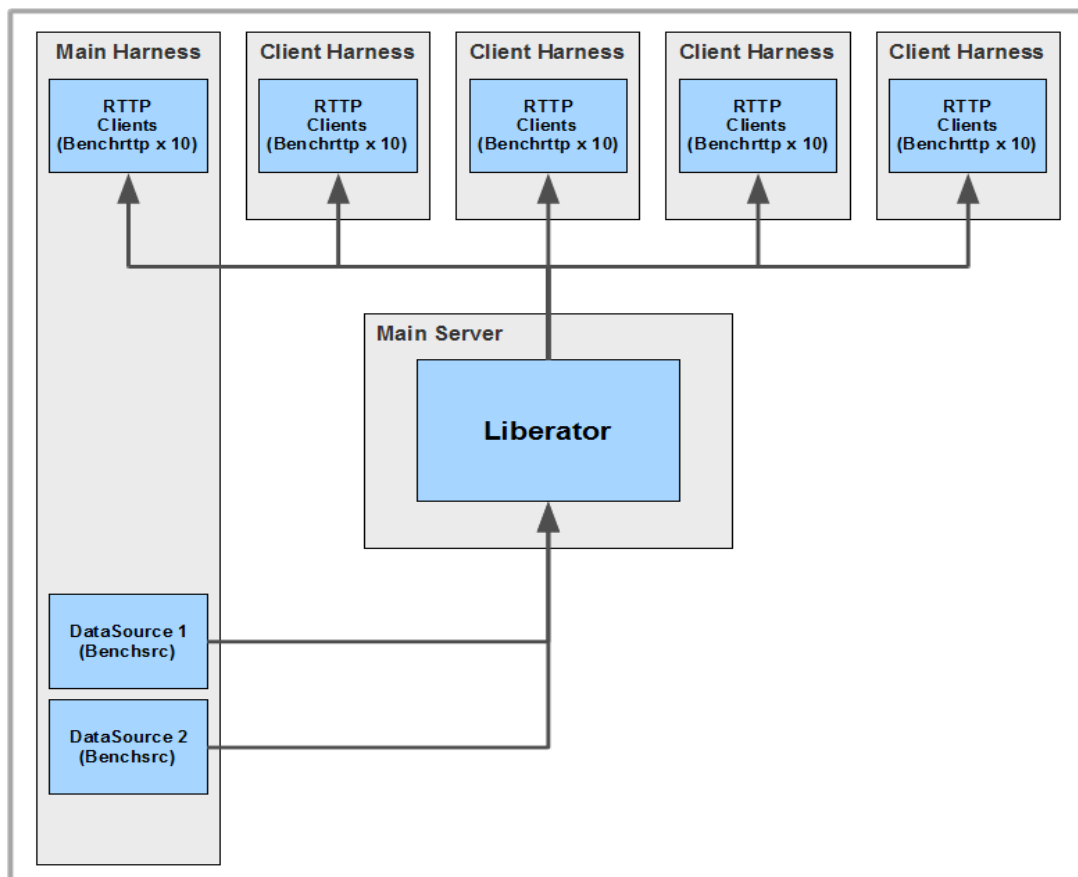


Figure 6–1: Benchmark test setup

The Liberator server was hosted on a machine of its own (Main Server). A separate machine (Main Harness) was used to host up to four DataSources (Benchsrc) feeding the Liberator. Main Harness and four further machines (Client Harness's), hosted up to ten instances each of Benchrttp.

The number of Benchrttp instances required for each test was determined by the maximum number of simulated clients needed to run the test; enough Benchrttp resource was required to ensure that Liberator limits could be reached before any limits imposed by Benchrttp.

For more detailed information about the hardware and software used to run the tests, see the following sections.

6.2 Test configurations

Most of the tests determined the average latency of update messages delivered to clients, against the number of subscribing (logged on) clients. This gave a measure of how the Liberator performed as the overall client update rate increased, for the following reason.

Each client subscribed to a number of objects (the number varied according to the test). Because each subscription required Liberator to update those objects on the client in line with the updates from the DataSources, increasing the number of logged on clients had the effect of proportionally increasing the Liberator's client update load.

For example, if each client subscribed to 100 objects and there were 2,000 clients logged on to the Liberator, then the Liberator would have to supply $100 \times 2,000 = 200,000$ object updates per second to all the clients.

The latency of the client update messages was measured at the client end. In all cases measurements were taken from the application log files and the results were plotted.

The most important Liberator parameters affecting performance are:

- ◆ The number of DataSource threads.
- ◆ The number of session threads (**threads-num** configuration parameter).
- ◆ The value of the **burst-max** configuration parameter.

Many of the tests were performed with varying numbers of DataSource threads and/or session threads, in order to determine the impact of the thread settings on Liberator performance. The test results show the optimum configuration for each scenario.

6.3 Test software

Caplin Liberator server

The tests were run against a Caplin Liberator 4.5 server. The server configuration only had a few changes from the default settings, as follows:

- ◆ Log Cycling. This was modified to prevent the very large amounts of data being processed from using up too much disk space.
- ◆ System Max Files. This was increased to 131072 to allow high numbers of clients to connect.
- ◆ Object Throttling. The default object throttling of 1 second was turned off for benchmarking. This allowed all clients to receive all the messages they were subscribed to.
- ◆ Threads. The numbers of DataSource threads and session threads were adjusted according to the needs of the individual tests. Many of the tests were repeatedly run with different numbers of threads in order to determine the impact of the thread settings on Liberator performance.
- ◆ Burst Settings. The default setting for **burst-max** is 0.5. This was altered to 0 for most of the tests. The "Very high updates - batching" test (page 19) used several values for this parameter.

Operating system

The operating system used on all the server hardware was Linux – CentOS 5.5 or Red Hat® Enterprise Linux® version 4. This was a standard configuration with only one significant change; the number of open file descriptors was increased to allow Liberator to support high numbers of client connections. This is detailed in the **Caplin Liberator Administration Guide**.

Test DataSource application (Benchsrc)

The Benchsrc test tool was configured to produce updates to sets of objects at a known message rate. Benchsrc is a single threaded application, so for those tests where the Liberator was configured to use multiple DataSource threads there was an instance of Benchsrc for each DataSource thread.

Test RTTP client application (Benchrttp)

The Benchrttp test tool was configured to request sets of objects that were published by the DataSource. It measured the number of messages being received and also the latency of the messages. The simulated clients all used RTTP Type 2 connections – this type of connection is HTTP tunnelled.

6.4 Test hardware

The benchmark tests used up to 6 machines as shown in the diagram in Figure 6–1 on page 28. Caplin Liberator ran on a dedicated machine (Main Server). The test RTTP client processes (Benchrttp) ran on up to 5 separate machines to spread the load (Main Harness and Client Harnesses). The test DataSources (Benchsrc) ran on one of the Benchrttp machines (Main Harness).

Main Server

The main server runs Liberator, which is the component being tested.

Components	Liberator
Vendor	Dell
Model	PowerEdge R415
Processors	2 x Six-Core AMD Opteron™ Processor 4180
Memory	16GB
Operating System	CentOS 5.5 (Kernel 2.6.18-194.el5 64bit)

Main Harness

This machine is used to provide the data for the tests. It also runs some of the client harnesses so that latency can be measured using the same clock from source to destination.

Components	Benchsrc, Benchrtt(s)
Vendor	HP
Model	ProLiant DL585
Processors	4 x Dual-Core AMD Opteron™ Processor 8220 SE
Memory	8GB
Operating System	Red Hat Enterprise Linux ES release 4 (Kernel 2.6.9-55.ELsmp 64bit)

Client Harnesses

These machines run the rest of the client harnesses. Latency is still measured on these machines, to assert that thresholds are not broken, but graphs are all based on times recorded on the main harness machine.

Components	Benchrtt(s)
Vendor	Dell
Model	PowerEdge R210
Processors	Quad-Core Intel® Xeon® CPU X3460
Memory	4GB
Operating System	CentOS 5.5 (Kernel 2.6.18-194.el5 64bit)

6.5 The network

The test machines were set up on a Gigabit network which was used solely for transmitting the test data. Any control messages or terminal access used a separate network.

7 Frequently asked questions

The following sections discuss various issues concerning how to configure and tune Liberator and its environment to achieve the required performance.

7.1 What burst configuration should we use?

Liberator can be configured to batch messages together to improve overall performance. Liberator's default configuration uses a 0.5 second batch time (for low latency messaging this should be reduced). The configuration option for this is called **burst-max**. With this configuration, Liberator may batch together outgoing messages to a client, delaying them by up to 0.5 seconds. If the message rate is slower than the burst setting the messages will be sent immediately. Batching messages in this way means that with a fairly constant message rate an average latency of half the burst-max setting will be introduced. However, the benefit of this is that the Liberator and network can cope with higher message rates. There is clearly a trade off between latency and message rate, which is why the tests carried out in these Benchmarks show results for different settings of burst-max (see section 4.5 "Very high updates - batching" on page 19).

It is clear that a low **burst-max** setting (for example 0.1 sec) can improve overall average latency and achieve much higher messages rates than with no batching. In some cases, increasing burst-max to the default of 0.5 sec will allow Liberator to achieve even higher message rates, but this is at the expense of increased latency.

7.2 How many threads should we configure?

There isn't a simple answer to this question.

Liberator's configuration option **threads-num** sets the number of session threads used to service client connections. Liberator also implements multiple DataSource threads – there is one thread per DataSource connection.

How many threads you require depends on the maximum anticipated update rate from Liberator's DataSource peers, the maximum expected number of subscribing clients, the subscription profile of the clients (see "DataSource threads" below), and the hardware characteristics of the machine running the Liberator (CPU speed and number of CPUs).

Caplin Systems can provide customers with expert guidance on configuring the optimum number of Liberator threads to meet their particular requirements.

DataSource threads

DataSource threads enable Liberator to better handle high update rates from its DataSources. Even if there is only one DataSource instance feeding Liberator, you can still configure more than one connection to the DataSource so that the Liberator's performance can benefit from using multiple DataSource threads.

There is also a relationship between the number of *clients* using the Liberator and the number of DataSource threads required, but this depends on the subscription profile of the clients, as follows.

At one extreme, if each client subscribes to a different set of objects, then the update demand on the DataSources will increase roughly in proportion to the number of subscribing clients. So in this case, as the maximum anticipated number of clients increases, more DataSource threads will be required in order to achieve the same message latency.

At the other extreme, if each client subscribes to the same set of objects, the update demand on the DataSources will remain constant as the number of subscribing clients increases. So in this case the number of DataSource threads required will *not* depend on the number of subscribing clients.

You should configure enough DataSource connections (and hence DataSource threads) to allow Liberator to handle the maximum anticipated update rate from the DataSource(s) with acceptable message latency for the maximum number of clients expected to use the Liberator.

Don't use more DataSource threads than you need, because the additional threads will not produce a significant additional improvement in performance. There will be a limit above which adding more DataSource connections has little extra benefit, because the limiting factor becomes the number of session threads.

Session threads

Provided there are enough DataSource threads to handle updates from the DataSources, increasing the number of session threads will allow more clients to subscribe with no unacceptable increase in message latency. However there will be an upper limit on the number of session threads, beyond which performance will decrease – see "CPU resource considerations" below.

CPU resource considerations

Once CPU resource limits have been reached on the machine running the Liberator, adding more threads will not necessarily improve performance. In these tests it was usually the case that 10 session threads (on a 12 core server) gave the best results, this left CPU cores available for the DataSource thread or threads.

7.3 How do message sizes affect performance?

When high numbers of clients and messages are used, the size of the message plays a significant part in the overall performance. Larger update messages decrease Liberator's maximum effective message rate. Liberator also performs slightly better when handling update messages consisting of a small number of large fields rather than messages containing a larger number of smaller fields.

Section 4.6 "Message sizes" on page 21 shows the impact of message size on the latency, CPU, and number of clients Liberator can support.

7.4 How much bandwidth will our Liberator use?

In the context of Liberator performance, bandwidth is the update rate delivered to all subscribed clients in bytes/sec. In most of these tests each message is about 54 bytes long; this may seem small, but it is because RTTP is a very efficient protocol. The message contains five fields, a typical update in financial applications.

Each test gives some details of the bandwidth used per client and overall.

7.5 How many subscriptions can Liberator handle?

The number of subscriptions a client has does not significantly affect performance directly, rather the number of messages is far more significant. This can be controlled using throttling.

7.6 How much disk space will our Liberator need?

Liberator uses approximately 60 megabytes of disk space when it is installed. Running Liberator requires extra disk space for log files.

The amount of disk space needed for the log files depends entirely on messages rates and client activity. Liberator supports configuration options to control the cycling of log files, so it is possible to limit how much disk space is used and how much information is saved in log files. Log files can grow to high single-digit gigabytes per day in some setups.

8 Glossary of terms and acronyms

This section contains a glossary of terms and acronyms relating to the Liberator benchmark.

Term	Definition
Benchsrc	A Caplin benchmark test tool that provides DataSource messages as input to a Caplin Liberator server. It is primarily intended to be used in conjunction with Benchrttp and the control scripts from the Caplin Benchmarking kit.
Benchrttp	A Caplin benchmark test tool that connects to a Liberator server and simulates a configurable number of clients and contributors of streamed RTTP data. It is primarily intended to be used in conjunction with Benchsrc and the control scripts from the Caplin Benchmarking kit.
burst-max	<p>The efficiency of Caplin Liberator can be increased by writing user output in defined "bursts", particularly in a system with a large number of clients where bursting batches together small messages before outputting them to a client.</p> <p>burst-max is a Liberator configuration parameter that controls bursting. It is the maximum time in seconds of client update buffering before Liberator will send updates to a client.</p> <p>For more information see the Caplin Liberator Administration Guide.</p>
Caplin Liberator	Caplin Liberator is a real-time financial internet hub that delivers trade messages and market data to and from subscribers over any network.
Caplin Xaqua	<p>A framework for building single-dealer platforms that enables banks to deliver multi-product trading direct to client desktops.</p> <p>Caplin Xaqua was formerly called "the Caplin Platform".</p>
DataSource	DataSource is the internal communications infrastructure used by Caplin Xaqua's server components such as Caplin Liberator , Caplin Transformer, and DataSource adapters .
DataSource adapter	A DataSource application that integrates with an external (non-Caplin) system, exchanging data and/or messages with that system.
DataSource application	A Caplin Xaqua application that uses the Caplin DataSource APIs to communicate with other Caplin Xaqua applications via the DataSource protocol.
DataSource peer	Alternative name for a DataSource application .
JMX	<p><u>J</u>ava <u>M</u>anagement <u>E</u>xtensions</p> <p>A Java technology for application and network management.</p>

Term	Definition
RTTP	<u>R</u> ea <u>T</u> ime <u>T</u> ext <u>P</u> rotocol. Caplin's protocol for streaming real-time financial data from Caplin Liberator servers to client applications, and for transmitting trade messages between clients and Liberator in both directions.
StreamLink	The StreamLink libraries connect client applications to Caplin Liberator via the RTTP protocol. They provide an object oriented API that gives access to RTTP functionality.
Throttling	A technique used by Caplin Liberator to improve performance by restricting the rate at which object updates are sent to a client.

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.