WebSphere. software

WebSphere MQ 7.5

Managed File Transfer Performance Report for Linux

Configuration and Measurements for the following products:

WebSphere MQ 7.5

IBM Corporation
WebSphere MQ Performance Team
Sept 2012

Please take Note!
Before using this report, please be sure to read the paragraphs on “disclaimers”, “warranty and liability exclusion”, “errors and omissions” and the other general information paragraphs in the "Notices" section below.

First Edition, November 2012.
This edition applies to the Managed File Transfer component of WebSphere for Linux V7.5 (and to all subsequent releases and modifications until otherwise indicated in new editions).

Note to U.S. Government Users
Documentation related to restricted rights.
Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule contract with IBM Corp

Notices
DISCLAIMERS
The performance data contained in this report were measured in a controlled environment. Results obtained in other environments may vary significantly.

You should not assume that the information contained in this report has been submitted to any formal testing by IBM.

Any use of this information and implementation of any of the techniques are the responsibility of the licensed user. Much depends on the ability of the licensed user to evaluate the data and to project the results into their own operational environment.

WARRANTY AND LIABILITY EXCLUSION
The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore this statement may not apply to you.

In Germany and Austria, notwithstanding the above exclusions, IBM's warranty and liability are governed only by the respective terms applicable for Germany and Austria in the corresponding IBM program license agreement(s).

ERRORS AND OMISSIONS
The information set forth in this report could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; any such change will be incorporated in new editions of the information. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this information at any time and without notice.

INTENDED AUDIENCE
This report is intended for architects, systems programmers, analysts and programmers wanting to understand the performance characteristics of the Managed File Transfer component of WebSphere MQ V7.5. The information is not intended as the specification of any programming interface that is provided by WebSphere. It is assumed that the reader is familiar with the concepts and operation of the WebSphere MQ V7.5 Managed File Transfer component.

LOCAL AVAILABILITY
References in this report to IBM products or programs do not imply that IBM intends to make these available in all countries in which IBM operates. Consult your local IBM representative for information on the products and services currently available in your area.

ALTERNATIVE PRODUCTS AND SERVICES
Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

USE OF INFORMATION PROVIDED BY YOU
IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

TRADEMARKS AND SERVICE MARKS
The following terms used in this publication are trademarks of International Business Machines Corporation in the United States, other countries or both:
• IBM
• WebSphere
• DB2

Other company, product, and service names may be trademarks or service marks of others.

EXPORT REGULATIONS
You agree to comply with all applicable export and import laws and regulations.

How this document is arranged

Performance Headlines
Pages: 2-17

Chapter 2 details the performance headlines for the two scenarios (client and bindings). Each scenario is detailed fully with diagrams in this section. The headline tests show how the Chunk Size property for an agent, and show the effect of transferring files as a group of transfers verses transferring files as a single transfer.

We detail the time taken for each transfer to complete, and the associated CPU utilisation for the hardware in use.

Tuning Recommendations
Pages: 18-20

Chapter 3 discusses the appropriate tuning that should be applied to both the WebSphere MQ network, and Managed File Transfer agents.

Measurement Environment
Pages: 21

Chapter 4 gives an overview of the environment used to gather the performance results. This includes a detailed description of the hardware and software.

Contents

1 Overview 1
2 Performance Headlines 2
2.1 Agents Connecting in Bindings Mode 5
2.1.1 65635 ChunkSize 5
2.1.2 131072 ChunkSize 6
2.1.3 262144 ChunkSize 7
2.1.4 524228 ChunkSize 9
2.1.5 Test Summary 10
2.2 Agents Connecting in Client Mode 12
2.2.1 65635 ChunkSize 12
2.2.2 131072 ChunkSize 13
2.2.3 262144 ChunkSize 14
2.2.4 524228 ChunkSize 16
2.2.5 Test Summary 17
3 Tuning Recommendations 18
3.1 WMQ Setup 18
3.2 WMQ File Transfer Edition Setup 19
3.3 WMQ File Transfer Edition Transfer recommendations 20
4 Measurement Environment 21
4.1 Agents 21
4.2 WebSphere MQ 21
4.3 Operating System 21
4.4 Hardware 21

Tables

Table 1 Linux64 64KB chunk size for Single and Multiple instance transfers 5
Table 2 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers 5
Table 3 Linux64 128KB chunk size for Single and Multiple instance transfers 6
Table 4 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers 6
Table 5 Linux64 256KB chunk size for Single and Multiple instance transfers 7
Table 6 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers 8
Table 7 Linux64 512KB chunk size for Single and Multiple instance transfers 9
Table 8 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers 9
Table 9 Linux64 Best transfer speeds for Single and Multiple instance transfers..... 10
Table 10 Linux64-MD5 Best transfer speeds for Single & Multiple instance transfers .. 11
Table 11 Linux64-2 Channel - Best transfer speeds for Single & Multiple instances 11
Table 12 Linux64 64KB chunk size for Single and Multiple instance transfers 12
Table 13 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers 12
Table 14 Linux64 128KB chunk size for Single and Multiple instance transfers 13
Table 15 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers .. 13
Table 16 Linux64 256KB chunk size for Single and Multiple instance transfers 14
Table 17 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers .. 15
Table 18 Linux64 512KB chunk size for Single and Multiple instance transfers 16
Table 19 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers .. 16
Table 20 Linux64 Best transfer speeds for Single and Multiple instance transfers..... 17
Table 21 Linux64-MD5 Best transfer speeds for Single & Multiple instance transfers .. 17

Figures

Figure 1 Linux64 64KB chunk size for Single and Multiple instance transfers..... 5
Figure 2 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers..... 6
Figure 3 Linux64 128KB chunk size for Single and Multiple instance transfers..... 7
Figure 4 Linux64 MD5 128KB chunk size for Single and Multiple instance transfers 7
Figure 5 Linux64 256KB chunk size for Single and Multiple instance transfers..... 8
Figure 6 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers..... 8
Figure 7 Linux64 512KB chunk size for Single and Multiple instance transfers..... 9
Figure 8 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers... 10
Figure 9 Linux64 512KB chunk size, Single & Multiple instance transfers, 2 channels. 10
Figure 10 Linux64 64KB chunk size for Single and Multiple instance transfers..... 12
Figure 11 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers... 13
Figure 12 Linux64 128KB chunk size for Single and Multiple instance transfers..... 14
Figure 13 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers. 14
Figure 14 Linux64 256KB chunk size for Single and Multiple instance transfers..... 15
Figure 15 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers. 15
Figure 16 Linux64 512KB chunk size for Single and Multiple instance transfers..... 16
Figure 17 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers. 17

1 Overview

The Managed File Transfer (MFT) component of WebSphere MQ is a managed file transfer product that uses WebSphere MQ as its transport layer.

This performance report details WebSphere MQ MFT in a range of scenarios, giving the reader information on transfer times and CPU utilisation. The report is based on measurements taken from Intel hardware running the Red Hat Enterprise Linux operating system.

While producing this document, identical tests were run using MQ File Transfer Edition V7.4.0.1[1]. From those results it was concluded that the majority of the tests were close enough in all respects to the WebSphere MQ V7.5 MFT component to preclude their inclusion in this report. Tests were also run to compare the 32 bit and 64 bit versions of WMQ. The difference between the two sets of test results again was not significant and so has not been included in this report.

On each test an additional checksum calculation (MD5) test was performed. This did change the test throughputs and the results for 64 bit WebSphere MQ V7.5 have been included[2]. It is also worth noting that there was no code page conversion required between the machines and that the tests only measured text messages. Additional results were also gathered when using multiple WebSphere MQ channels in a cluster.

At the end of each block of results is a summary of the findings. It should be noted that results obtained and the inferences made depend on the test infrastructure hardware and any change could alter the results significantly. The reader is urged to use the findings in this report only as guidelines – this is particularly true for results were all of the values are very close.

[1] WebSphere MQ File Transfer Edition was previously a separate product, which now forms the integrated managed file transfer capability, available for optional installation, within WebSphere MQ V7.5.

[2] This document only reports Linux 64-bit results primarily because the differences between 32 and 64 bit operating system results were negligible and in the future, most customers will use a 64-bit OS. The MD5 checksum is optional but is a commonly used additional feature and has been included to demonstrate the overhead incurred. There is no relationship between the choice of using the 64-bit product and the inclusion of the MD5 results.

2 Performance Headlines

The measurements for the performance headlines are based on the time taken to transfer a set of files and the associated CPU cost. A single performance measurement will use 2GB worth of files, with the size of the files varying as follows:

- 1MB
- 10MB
- 100MB

For example, when using a 1MB file then the test will transfer 2000 files in a single performance run. Each test varies the file size, but keeping the same overall MB transferred constant thus demonstrating the cost of the open and close file operations on transfer time and CPU usage.

The performance headlines demonstrate the effect of altering the agent's Chunk Size property (see http://pic.dhe.ibm.com/infocenter/wmqv7/v7r5/index.jsp "WebSphere MQ Managed File Transfer→Reference→Configuring→The agent.properties file" for more details on setting this property). The Chunk Size defines the size of the MQ message that the agent will use to transfer the files. The following Chunk Sizes (defined in bytes) have been used:

- 65536
- 131072
- 262144 (this is the agent's default value)
- 524288

To demonstrate the multithreaded capability of the agent, a multiple transfer test were run and compared to the equivalent single transfer test. The multiple transfer test divides the number of files transferred in the single transfer test into ten threads. The threads were then run simultaneously.
All files were transferred using text mode as opposed to binary mode. Each file transferred was the same size for a given performance run but contained random data. Transfers were submitted using the documented XML format.

The results are laid out in the chapters 2.1 and 2.2. Each test case has its own results table and associated graph. The first set of tables and figures show the reader the results for each chunk size (agentChunkSize) property has on the transfer time for a particular file size. At the end of the chapter is a summary that highlights the best combinations of chunk size and file size for single and multiple threaded tests.

Agents Connecting in Bindings Mode
In this scenario each agent is connected to a local queue manager in *bindings* mode. The two local queue managers and a third coordinating queue manager are clustered (cluster name is 'FTE'). The two local queue managers are connected via Sender/Receiver channel pairs. A third queue manager is located on another machine, and is used as the coordination Queue Manager. The following diagram details the exact scenario:

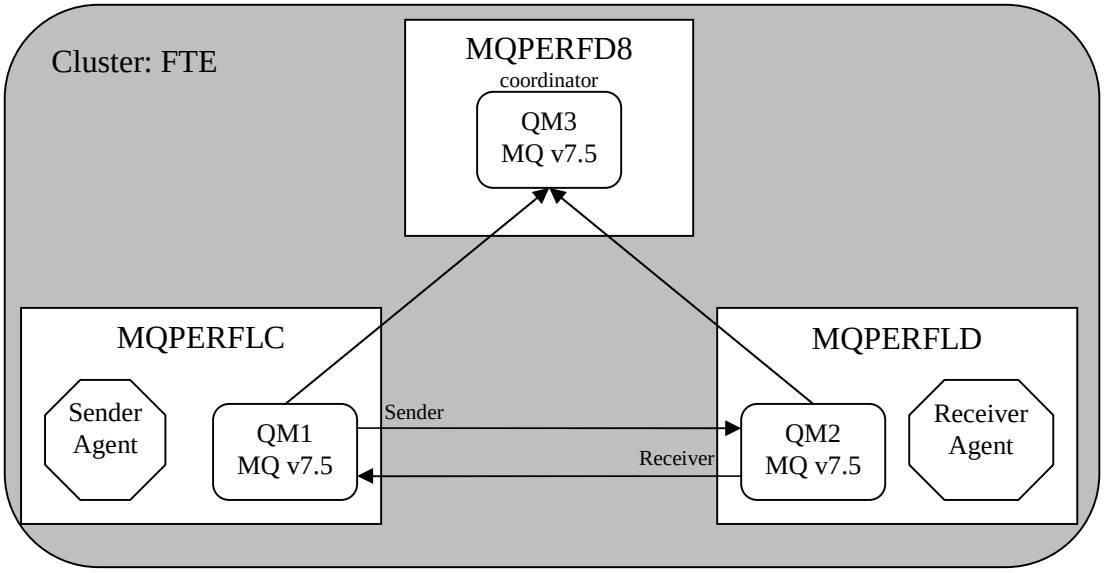

Agents Connecting in Client Mode

In this scenario each agent is connected to the same single remote queue manager in client mode. A second queue manager is placed on forth machine to act as the coordination queue manager. This coordination queue manager is not highly utilised as it is not directly involved in the transfers and so will have little or no effect on the Sender CPU values that are collected. The coordinator queue manager and agent queue manager are clustered (cluster name is 'FTE'). The following diagram details the exact scenario:

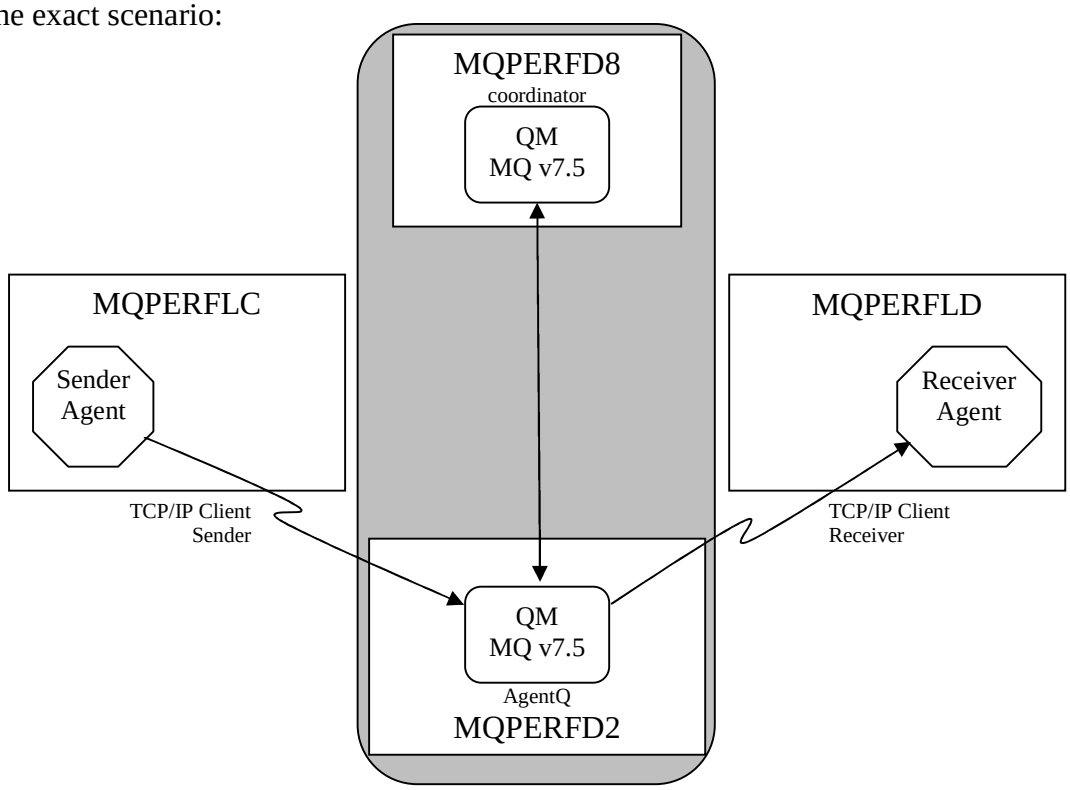

In the following sections, the transfer speeds and CPU costs are grouped by chunk size and show the comparative costs for single and multithreaded transfers.

Multiple WMQ channels in a cluster for Agents connecting in Bindings Mode

Recent additions to base WMQ now allow for the use of multiple channels to be used in a cluster instead of the default cluster transmit channel. To turn on multiple channels the following changes need to be applied (a more detailed description can be found in the MQ MFT documentation):-

- Set the *agentMultipleChannelsEnabled* property to true (destination agent only).
- In the channel definitions of the receiving queue manager, define a cluster-receiver channel for each channel used, e.g.
  `DEFINE CHANNEL(TO.DESTQMNAME_1) TRPTYPE(TCP) CHLTYPE(CLUSRCVR) CLUSTER(FTECLUSTER) CONNAME('<QMName>(<QMListenerPort>)')`
- In the queue definitions of the sending queue manager, define a queue manager alias corresponding to each channel used in the receiving queue manager, e.g.
  `DEFINE QREMOTE(SYSTEM.FTE.DESTQMNAME_1) RQMNAME(DESTQMNAME) CLUSTER(FTECLUSTER)`
  SYSTEM.FTE.DESTQMNAME_n is the mandatory naming convention for queue manager aliases because the sending agent searches for queue manager aliases of this format. There is no need to define a default XMITQ for the queue manager.

2.1 Agents Connecting in Bindings Mode

2.1.1 65635 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 - 65635	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	6%	6%	21.34 s	383.81 Mb
Source:10MB Single Transfer	0%	5%	6%	10.79 s	759.32 Mb
Source:100MB SingleTransfer	1%	4%	6%	10.06 s	814.15 Mb
Source:1MB Multi Transfer	1%	9%	8%	8.84 s	927.05 Mb
Source:10MB Multi Transfer	1%	7%	8%	8.21 s	998.09 Mb
Source:100MB Multi Transfer	1%	6%	8%	8.25 s	992.85 Mb

Table 1 Linux64 64KB chunk size for Single and Multiple instance transfers

Linux64 with MD5 – 65635	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB SingleTransfer	1%	5%	5%	33.01 s	248.13 Mb
Source:10MB SingleTransfer	1%	4%	5%	20.70 s	395.66 Mb
Source:100MB SingleTransfer	1%	4%	5%	19.60 s	418.04 Mb
Source:1MB MultiTransfer	0%	9%	10%	14.64 s	559.56 Mb
Source:10MB MultiTransfer	1%	9%	11%	8.95 s	914.97 Mb
Source:100MB MultiTransfer	1%	9%	10%	9.18 s	892.54 Mb

Table 2 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers

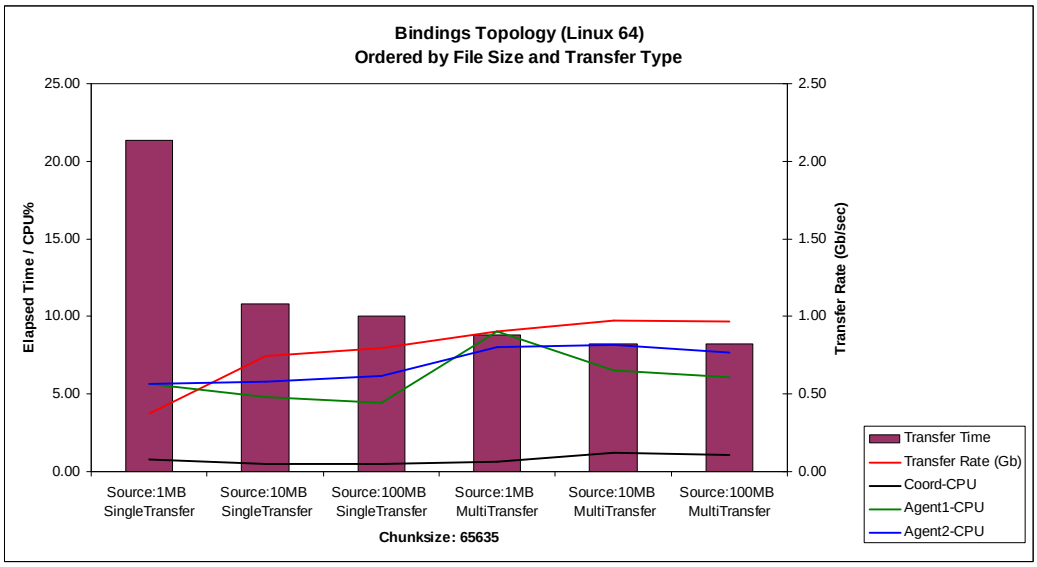

Figure 1 Linux64 64KB chunk size for Single and Multiple instance transfers

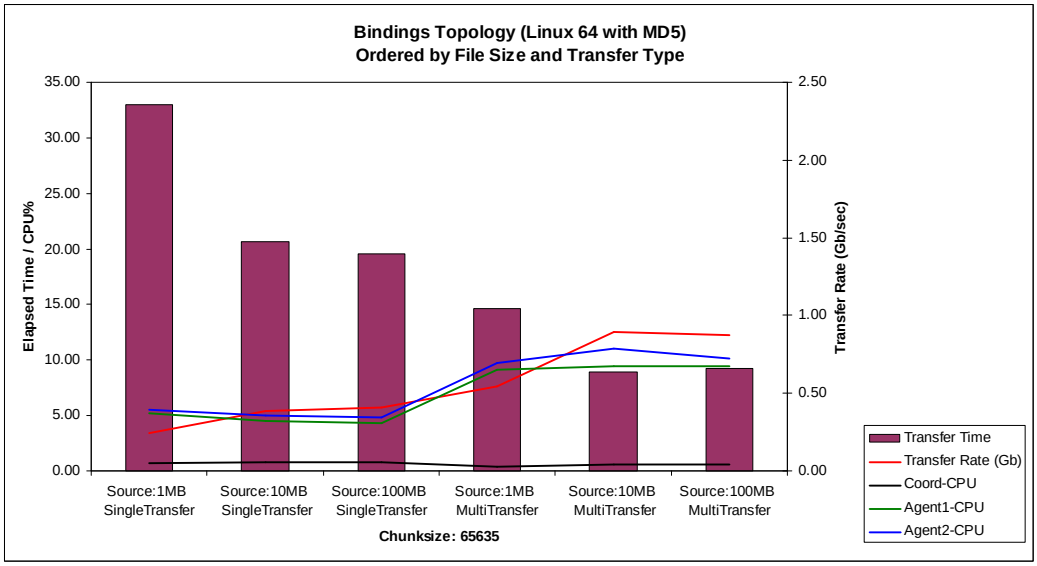

Figure 2 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers

2.1.2 131072 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 131072	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	6%	5%	21.18 s	386.76 Mb
Source:10MB Single Transfer	0%	5%	6%	10.85 s	755.21 Mb
Source:100MB Single Transfer	0%	4%	6%	10.21 s	802.25 Mb
Source:1MB Multi Transfer	1%	6%	6%	8.57 s	956.26 Mb
Source:10MB Multi Transfer	1%	6%	8%	8.17 s	1,003.31 Mb
Source:100MB Multi Transfer	1%	6%	8%	8.23 s	995.91 Mb

Table 3 Linux64 128KB chunk size for Single and Multiple instance transfers

Linux64 MD5 – 131072	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	5%	5%	32.76 s	250.04 Mb
Source:10MB Single Transfer	1%	5%	5%	20.82 s	393.56 Mb
Source:100MB Single Transfer	0%	4%	5%	19.68 s	416.27 Mb
Source:1MB Multi Transfer	0%	10%	10%	13.17 s	622.19 Mb
Source:10MB Multi Transfer	1%	9%	10%	9.82 s	834.10 Mb
Source:100MB Multi Transfer	1%	10%	10%	9.08 s	902.67 Mb

Table 4 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers

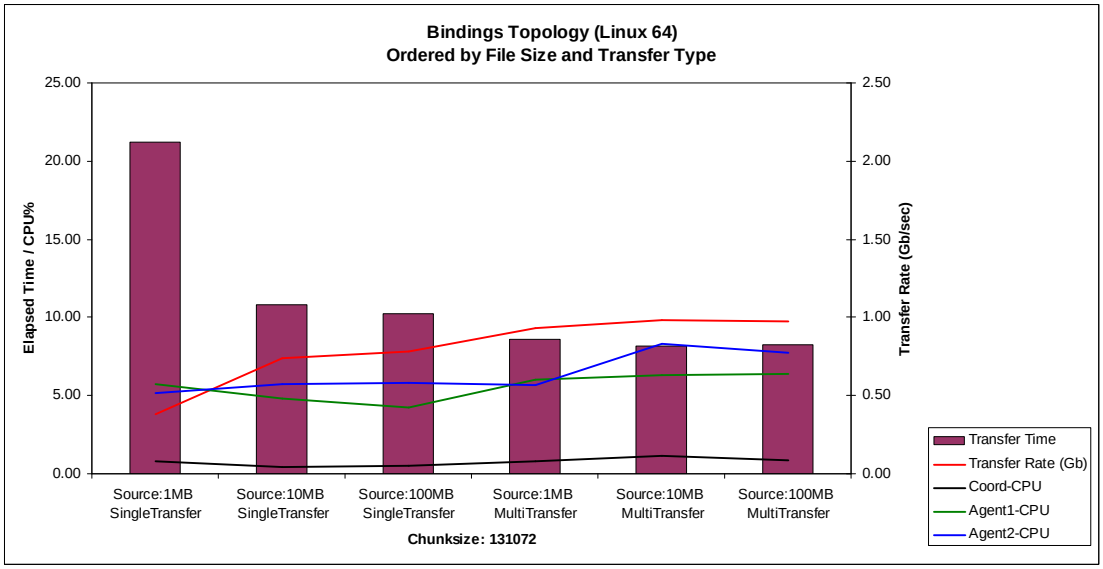

Figure 3 Linux64 128KB chunk size for Single and Multiple instance transfers

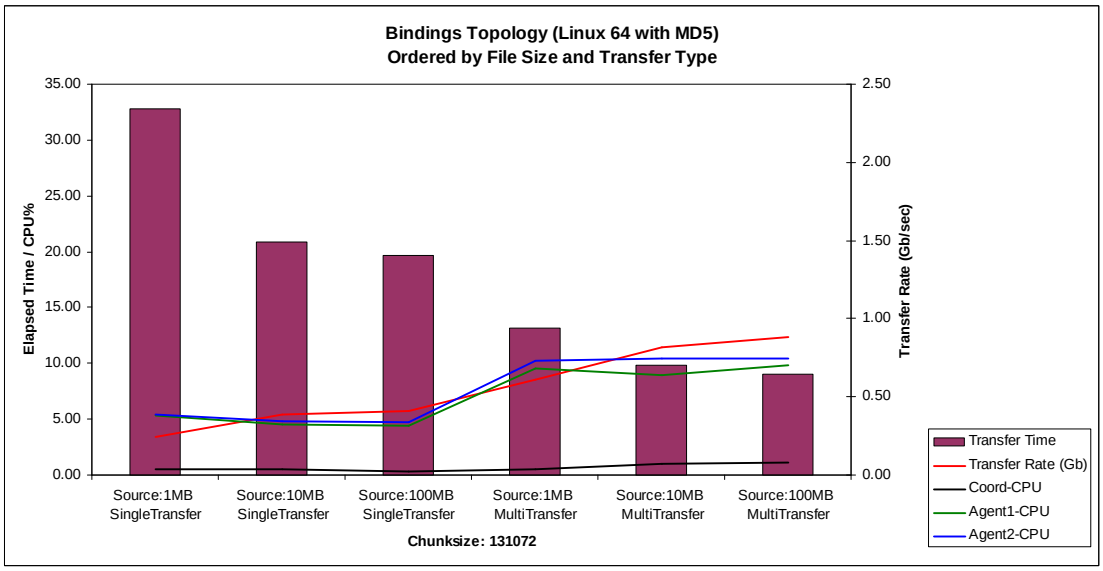

Figure 4 Linux64 MD5 128KB chunk size for Single and Multiple instance transfers

2.1.3 262144 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 256KB	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	6%	5%	21.03 s	389.58 Mb
Source:10MB Single Transfer	0%	5%	6%	11.05 s	741.54 Mb
Source:100MB Single Transfer	1%	4%	6%	10.29 s	796.32 Mb
Source:1MB Multi Transfer	1%	9%	8%	8.62 s	950.42 Mb
Source:10MB Multi Transfer	1%	6%	8%	8.08 s	1,013.90 Mb
Source:100MB Multi Transfer	1%	7%	8%	8.12 s	1,008.54 Mb

Table 5 Linux64 256KB chunk size for Single and Multiple instance transfers

Linux64 with MD5	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	5%	6%	32.51 s	251.95 Mb
Source:10MB Single Transfer	1%	5%	5%	20.79 s	394.10 Mb
Source:100MB Single Transfer	1%	4%	5%	19.71 s	415.61 Mb
Source:1MB Multi Transfer	1%	8%	9%	15.89 s	515.47 Mb
Source:10MB Multi Transfer	1%	10%	11%	8.94 s	915.89 Mb
Source:100MB Multi Transfer	1%	7%	7%	9.59 s	854.02 Mb

Table 6 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers

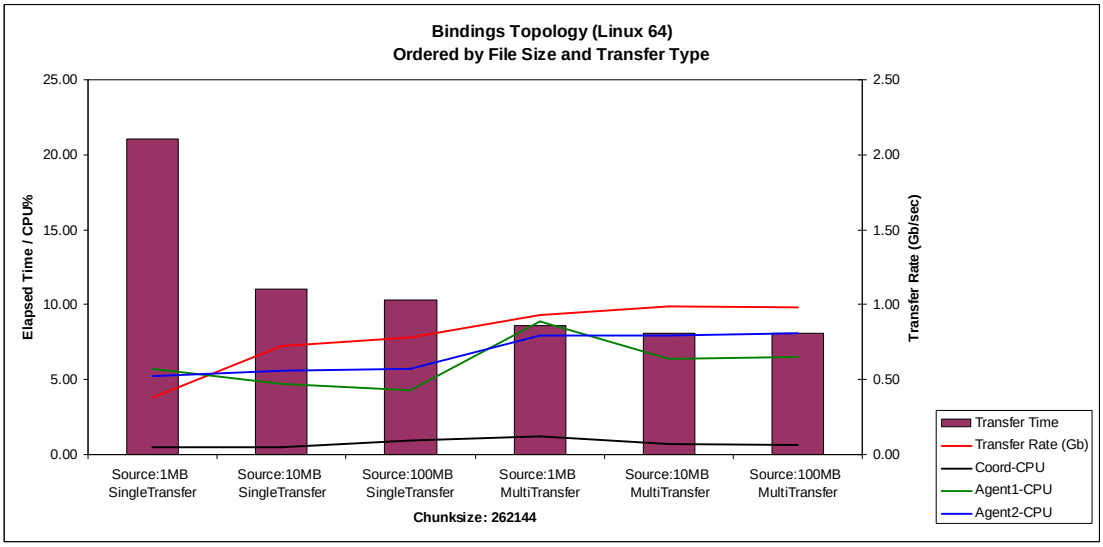

Figure 5 Linux64 256KB chunk size for Single and Multiple instance transfers

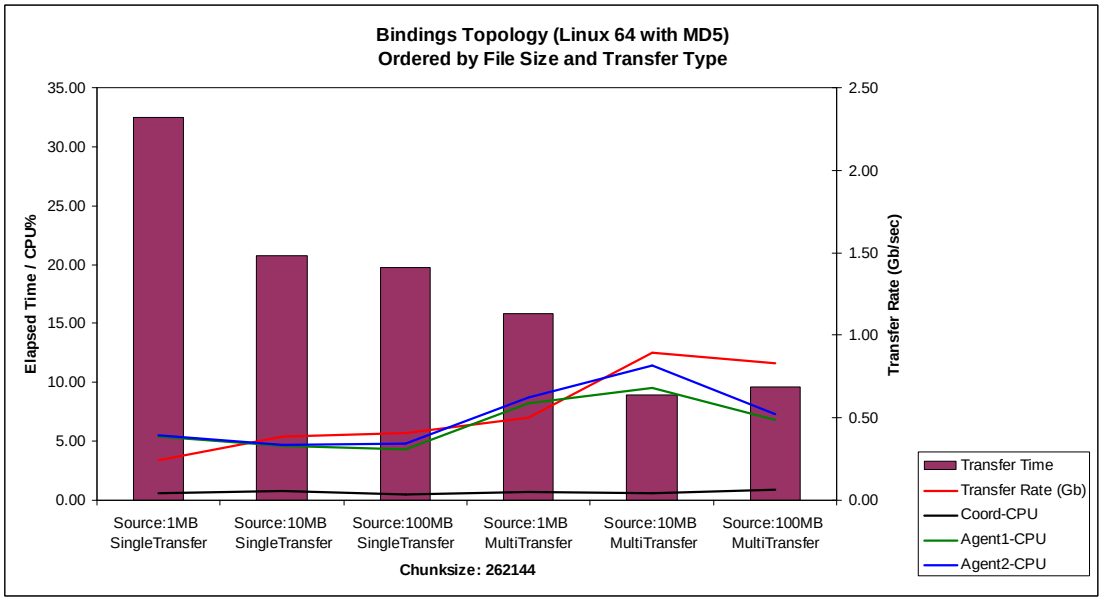

Figure 6 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers

2.1.4 524228 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 524228	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	6%	5%	16.76 s	488.89 Mb
Source:10MB Single Transfer	1%	3%	4%	10.45 s	783.87 Mb
Source:100MB Single Transfer	1%	4%	6%	9.33 s	877.68 Mb
Source:1MB Multi Transfer	1%	7%	7%	9.62 s	851.18 Mb
Source:10MB Multi Transfer	1%	5%	7%	9.27 s	883.74 Mb
Source:100MB Multi Transfer	1%	6%	8%	9.05 s	905.66 Mb

Table 7 Linux64 512KB chunk size for Single and Multiple instance transfers

Linux64 with MD5 – 524228	Coord-CPU	Agent1-CPU	Agent2-CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	1%	5%	5%	32.37 s	253.08 Mb
Source:10MB Single Transfer	1%	4%	5%	20.97 s	390.70 Mb
Source:100MB Single Transfer	1%	4%	5%	18.89 s	433.65 Mb
Source:1MB Multi Transfer	1%	6%	8%	30.34 s	270.02 Mb
Source:10MB Multi Transfer	1%	9%	11%	9.18 s	892.86 Mb
Source:100MB Multi Transfer	1%	9%	11%	9.37 s	873.91 Mb

Table 8 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers

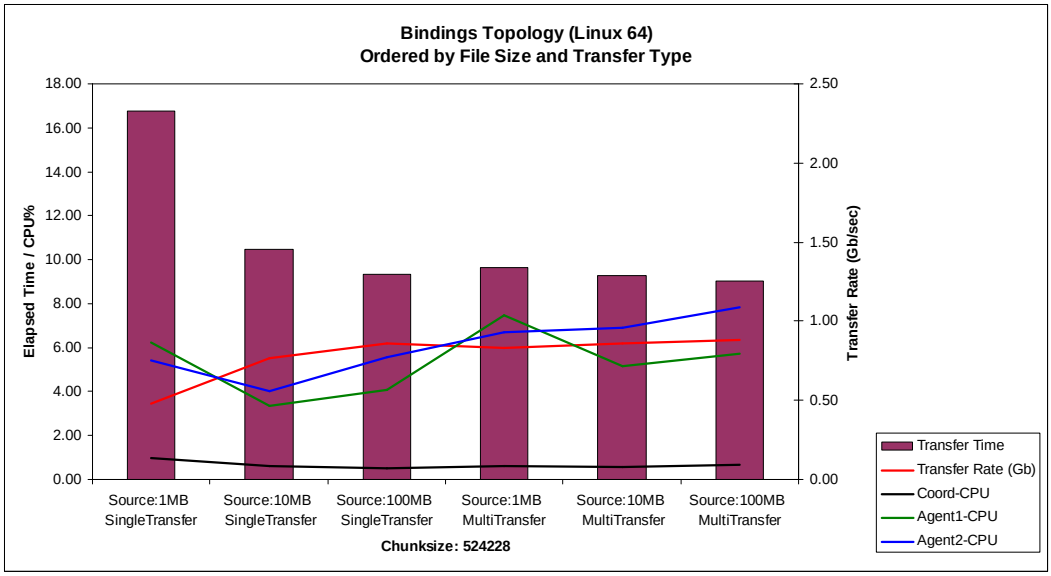

Figure 7 Linux64 512KB chunk size for Single and Multiple instance transfers

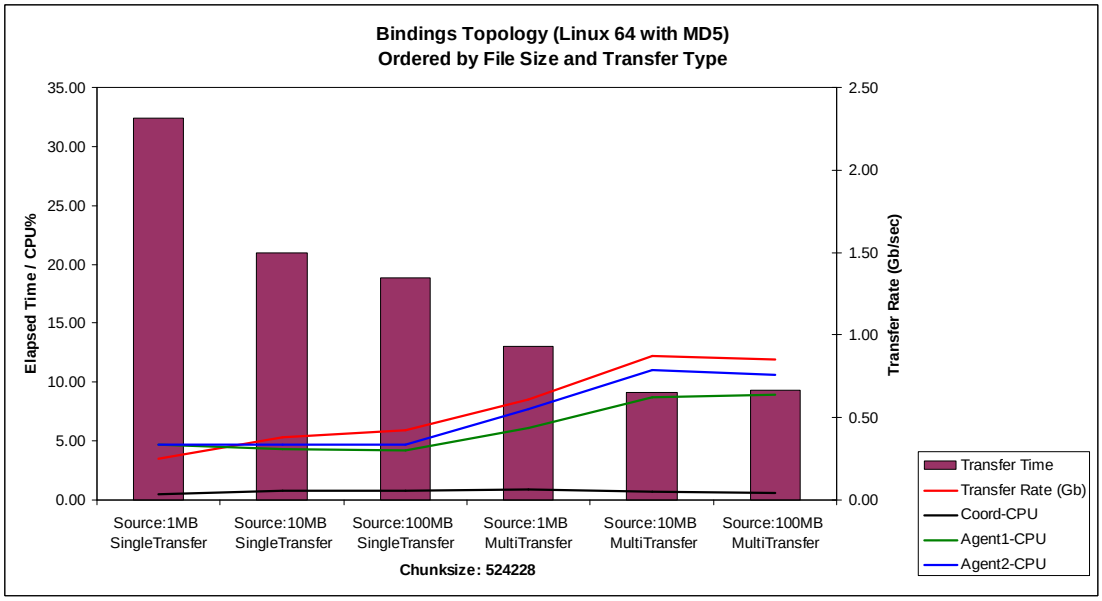

Figure 8 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers

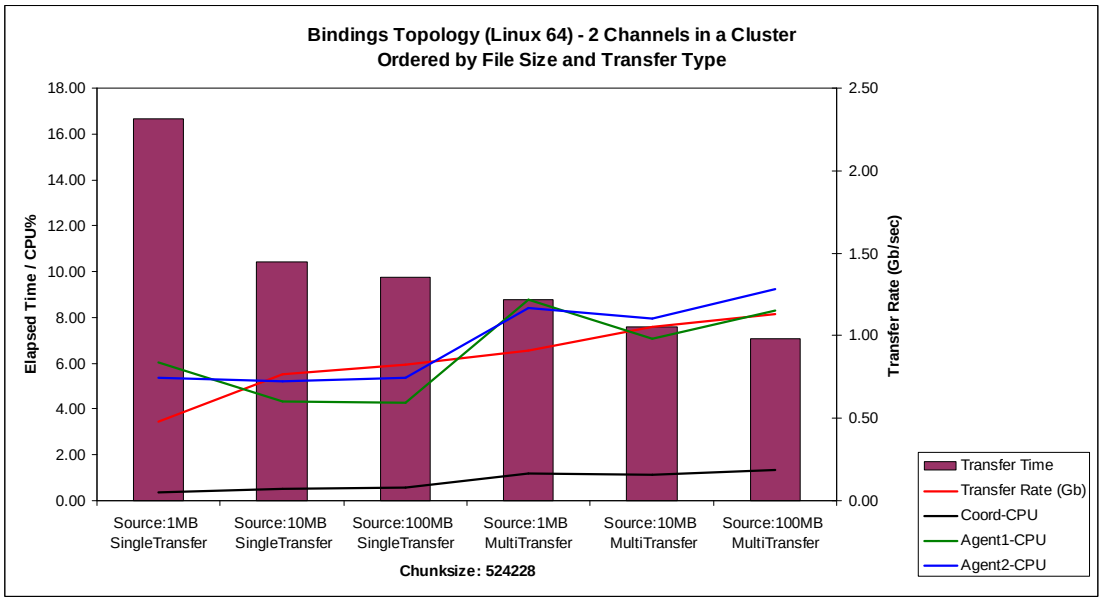

Figure 9 Linux64 512KB chunk size, Single & Multiple instance transfers, 2 channels

2.1.5 Test Summary

Looking across the results, the quickest transfers were attained at the following chunk sizes, file sizes and transfer types.

File Size and Transfer Type	Transfer Time	Chunk size
1MB Single Instance Transfer	16.76 s	524228
10MB Single Instance Transfer	10.45 s	524228
100MB Single Instance Transfer	9.33 s	524228
1MB Multi Instance Transfer	8.57 s	131072
10MB Multi Instance Transfer	8.08 s	262144
100MB Multi Instance Transfer	8.12 s	262144

Table 9 Linux64 Best transfer speeds for Single and Multiple instance transfers

File Size and Transfer Type	Transfer Time	Chunk size
1MB Single Instance Transfer	32.37 s	524228
10MB Single Instance Transfer	20.70 s	65635
100MB Single Instance Transfer	18.89 s	524228
1MB Multi Instance Transfer	13.17 s	131072
10MB Multi Instance Transfer	8.94 s	262144
100MB Multi Instance Transfer	9.08 s	131072

Table 10 Linux64-MD5 Best transfer speeds for Single & Multiple instance transfers

	Difference in Transfer Time as a percentage between a single and two channel cluster for different chunk sizes			
File Size and Transfer Type	65535	131072	262144	524228
1MB Single Instance Transfer	0%	0%	0%	0%
10MB Single Instance Transfer	0%	0%	0%	0%
100MB Single Instance Transfer	-5%	0%	-5%	-5%
1MB Multi Instance Transfer	-5%	-5%	-5%	+5%
10MB Multi Instance Transfer	+10%	+10%	+10%	+20%
100MB Multi Instance Transfer	+10%	+10%	+10%	+20%

Table 11 Linux64-2 Channel - Best transfer speeds for Single & Multiple instances

The table of results above (tables 9 and 10) show that generally, for single transfers the larger the chunk size, the quicker the test completed. This would not be expected to continue much beyond 512KB.
For multiple client transfers there is no optimum chunk size although a smaller chunk size seems to occur more often. What is significant better is the transfer time, but this does come at the cost of increased agent CPU which is more pronounced on larger file transfers.
Table 11 shows a throughput comparison between a single and multiple (two) channels as a percentage. It is clear from the table that for single instance transfers there is no performance gain from using more than one channel. In fact, with larger files sizes the throughput is slightly worse (around 5%). Similarly, for multiple instances with larger file sizes, it is obvious that using multiple channels has significant benefits with transfer rates improving by as much as 20% with larger chunk sizes. Further tests were made with more channels but no additional gains were seen so the results have not been included in this report (with four channels, for all transfers, there was a performance loss of up to 5% on the values in the table). The graph of the best results can be seen in Fig 9 – note that the transfer rates are up and the elapsed time is down for the multiple instance transfers.

2.2 Agents Connecting in Client Mode

2.2.1 65635 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 65635	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	22%	1%	3%	3%	31.91 s	256.72 Mb
Source:10MB Single Transfer	22%	1%	3%	4%	16.74 s	489.24 Mb
Source:100MB Single Transfer	20%	1%	2%	4%	16.77 s	488.37 Mb
Source:1MB Multi Transfer	23%	1%	3%	4%	22.46 s	364.71 Mb
Source:10MB Multi Transfer	23%	1%	2%	3%	21.46 s	381.80 Mb
Source:100MB Multi Transfer	26%	1%	2%	4%	18.86 s	434.47 Mb

Table 12 Linux64 64KB chunk size for Single and Multiple instance transfers

Linux64-MD5 – 65635	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	17%	0%	4%	4%	38.45 s	213.03 Mb
Source:10MB Single Transfer	15%	1%	4%	5%	21.16 s	387.24 Mb
Source:100MB Single Transfer	15%	1%	4%	5%	19.81 s	413.44 Mb
Source:1MB Multi Transfer	24%	1%	6%	7%	21.30 s	384.57 Mb
Source:10MB Multi Transfer	25%	0%	5%	6%	19.53 s	419.40 Mb
Source:100MB Multi Transfer	25%	1%	5%	6%	19.61 s	417.65 Mb

Table 13 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers

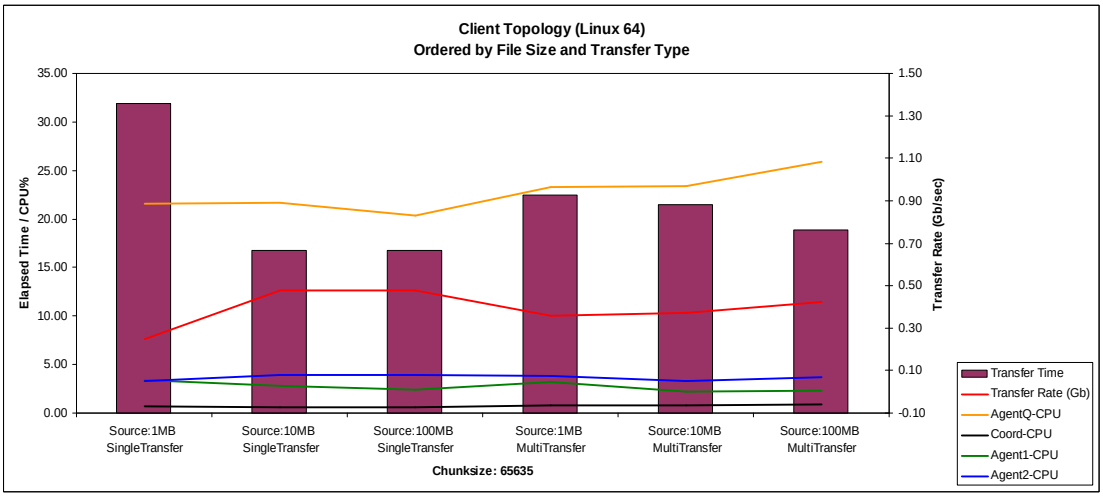

Figure 10 Linux64 64KB chunk size for Single and Multiple instance transfers

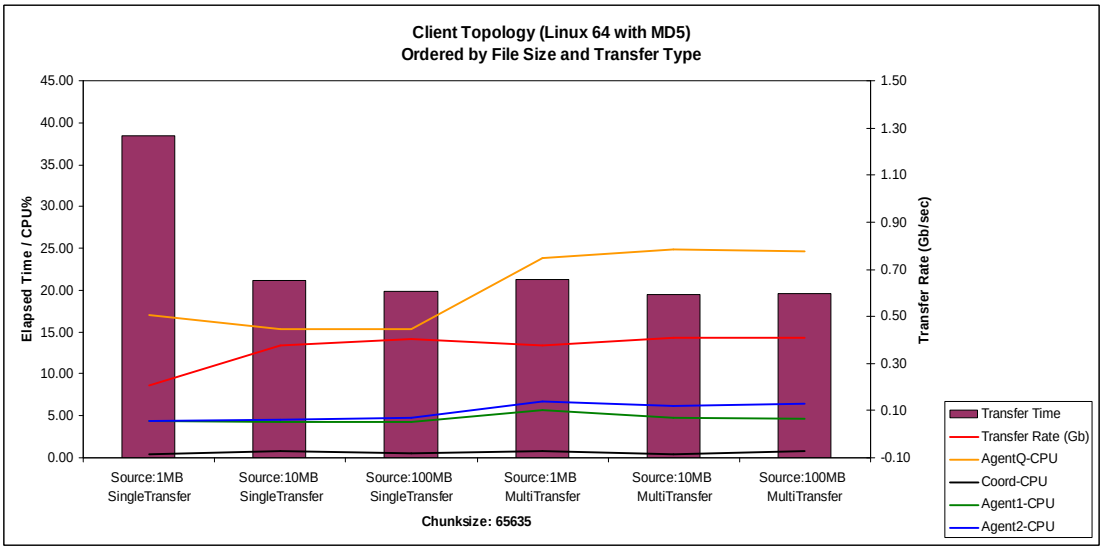

Figure 11 Linux64-MD5 64KB chunk size for Single and Multiple instance transfers

2.2.2 131072 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 131072	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	23%	1%	3%	3%	32.97 s	248.44 Mb
Source:10MB Single Transfer	22%	1%	3%	4%	16.84 s	486.53 Mb
Source:100MB Single Transfer	20%	1%	2%	4%	16.83 s	486.68 Mb
Source:1MB Multi Transfer	24%	1%	3%	4%	21.25 s	385.49 Mb
Source:10MB Multi Transfer	24%	0%	2%	3%	20.52 s	399.21 Mb
Source:100MB Multi Transfer	24%	0%	2%	4%	19.35 s	423.38 Mb

Table 14 Linux64 128KB chunk size for Single and Multiple instance transfers

Linux64-MD5 – 131072	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	16%	1%	5%	4%	36.65 s	223.52 Mb
Source:10MB Single Transfer	15%	0%	4%	5%	21.42 s	382.45 Mb
Source:100MB Single Transfer	15%	1%	4%	5%	19.93 s	411.04 Mb
Source:1MB Multi Transfer	23%	0%	6%	6%	22.01 s	372.16 Mb
Source:10MB Multi Transfer	25%	1%	5%	6%	19.77 s	414.37 Mb
Source:100MB Multi Transfer	24%	1%	5%	6%	19.98 s	409.94 Mb

Table 15 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers

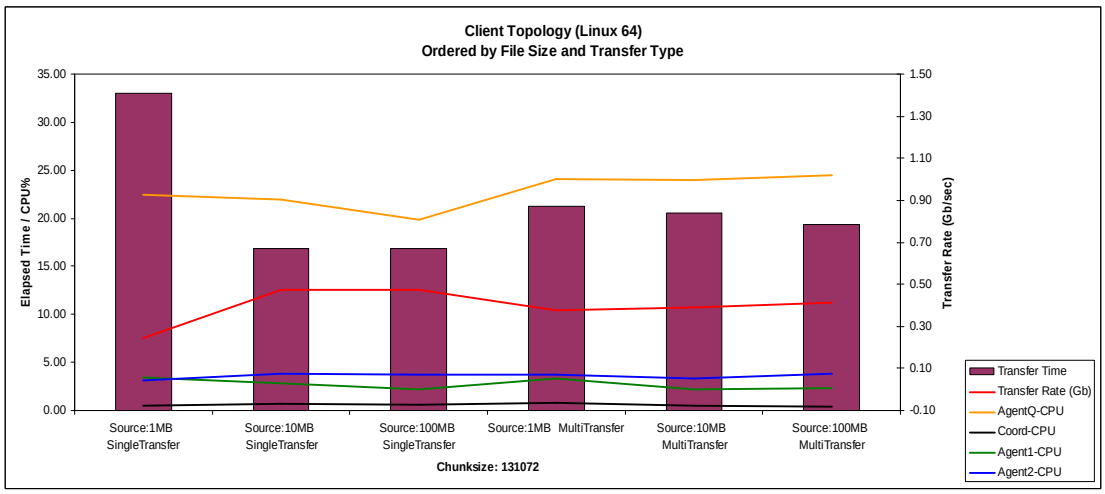

Figure 12 Linux64 128KB chunk size for Single and Multiple instance transfers

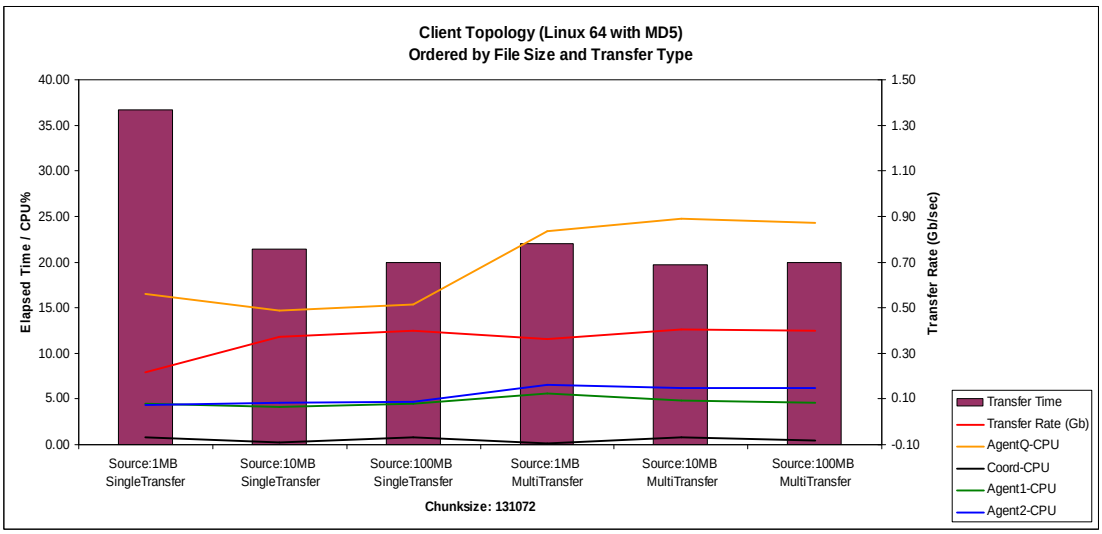

Figure 13 Linux64-MD5 128KB chunk size for Single and Multiple instance transfers

2.2.3 262144 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 262144	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	22%	1%	3%	3%	32.69 s	250.62 Mb
Source:10MB Single Transfer	21%	0%	3%	4%	17.72 s	462.22 Mb
Source:100MB Single Transfer	21%	1%	2%	4%	15.80 s	518.39 Mb
Source:1MB Multi Transfer	25%	1%	3%	4%	21.42 s	382.37 Mb
Source:10MB Multi Transfer	24%	1%	2%	3%	20.72 s	395.35 Mb
Source:100MB Multi Transfer	25%	1%	2%	4%	19.39 s	422.42 Mb

Table 16 Linux64 256KB chunk size for Single and Multiple instance transfers

Linux64-MD5 – 262144	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB SingleTransfer	17%	0%	4%	4%	38.06 s	215.24 Mb
Source:10MB SingleTransfer	15%	1%	4%	5%	21.09 s	388.47 Mb
Source:100MB SingleTransfer	14%	1%	4%	4%	19.83 s	413.17 Mb
Source:1MB MultiTransfer	25%	1%	6%	7%	21.27 s	385.16 Mb
Source:10MB MultiTransfer	24%	1%	5%	6%	20.36 s	402.38 Mb
Source:100MB MultiTransfer	25%	1%	5%	6%	19.62 s	417.58 Mb

Table 17 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers

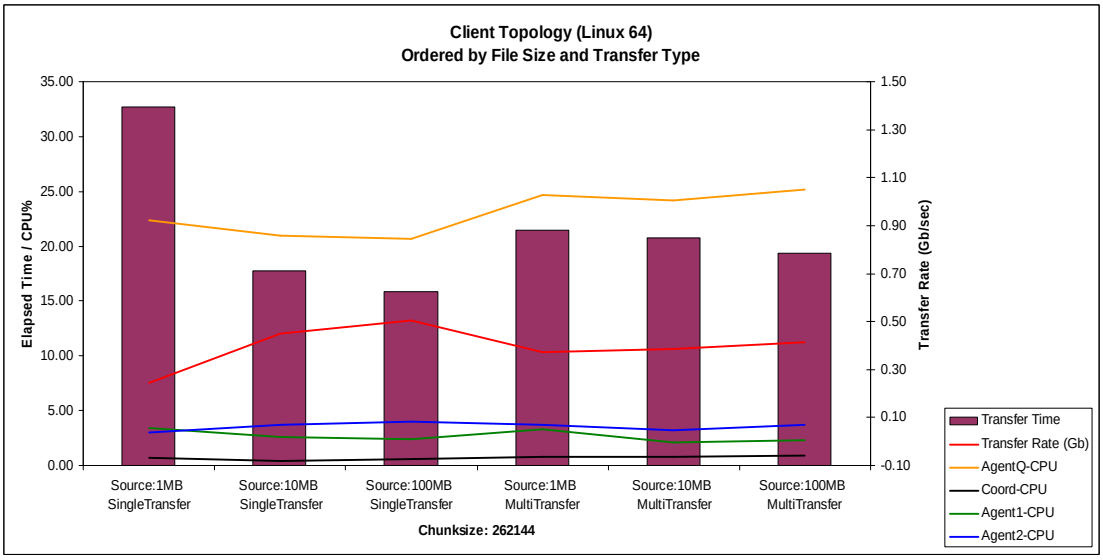

Figure 14 Linux64 256KB chunk size for Single and Multiple instance transfers

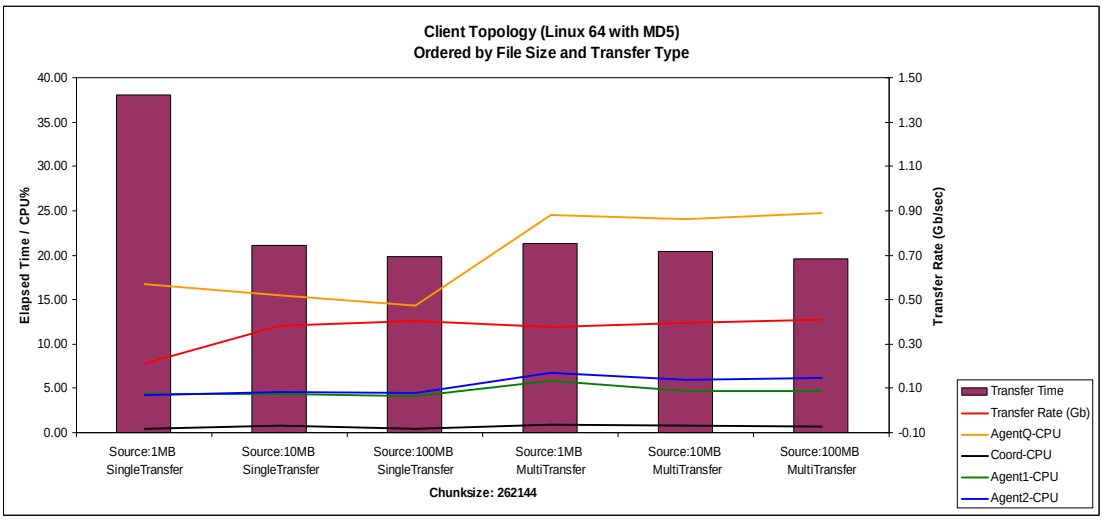

Figure 15 Linux64-MD5 256KB chunk size for Single and Multiple instance transfers

2.2.4 524228 ChunkSize

The table and chart below shows the relevant times and CPU utilisation for single and multi-application transfer with three different file sizes.

Linux64 – 524288	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	23%	0%	4%	4%	23.36 s	350.72 Mb
Source:10MB Single Transfer	21%	1%	2%	3%	18.78 s	436.13 Mb
Source:100MB Single Transfer	22%	0%	2%	3%	17.59 s	465.62 Mb
Source:1MB Multi Transfer	31%	1%	2%	3%	28.64 s	286.01 Mb
Source:10MB Multi Transfer	34%	1%	1%	2%	32.42 s	252.70 Mb
Source:100MB Multi Transfer	33%	1%	2%	3%	28.10 s	291.50 Mb

Table 18 Linux64 512KB chunk size for Single and Multiple instance transfers

Linux64-MD5 – 524288	AgentQ CPU	Coord CPU	Agent1 CPU	Agent2 CPU	Transfer Time	Transfer Rate
Source:1MB Single Transfer	14%	1%	4%	5%	32.08 s	255.34 Mb
Source:10MB Single Transfer	14%	1%	4%	4%	20.32 s	403.16 Mb
Source:100MB Single Transfer	17%	1%	4%	5%	18.50 s	442.84 Mb
Source:1MB Multi Transfer	30%	1%	4%	5%	28.56 s	286.88 Mb
Source:10MB Multi Transfer	32%	1%	3%	4%	30.53 s	268.31 Mb
Source:100MB Multi Transfer	32%	1%	3%	4%	28.07 s	291.88 Mb

Table 19 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers

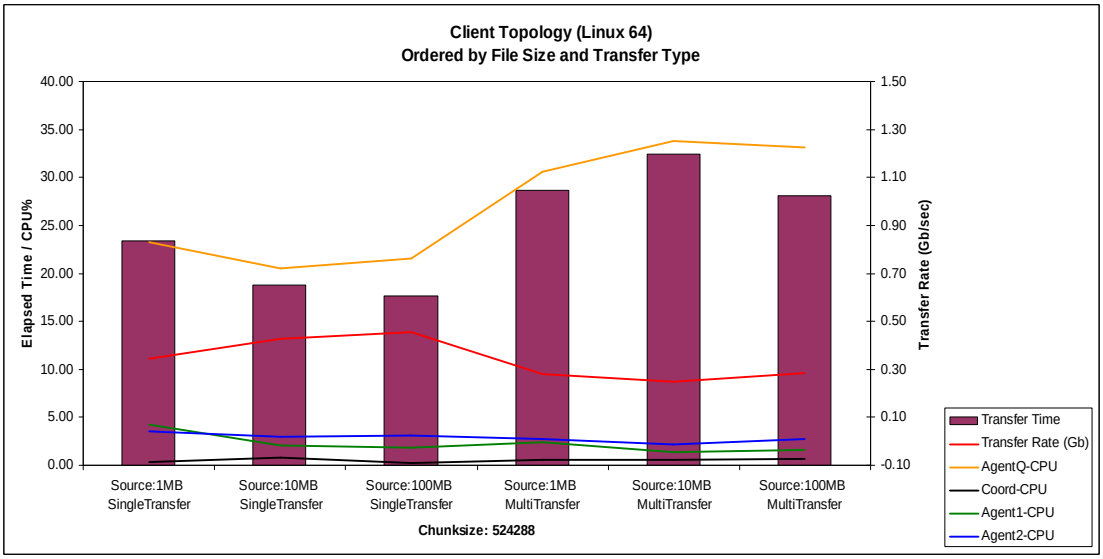

Figure 16 Linux64 512KB chunk size for Single and Multiple instance transfers

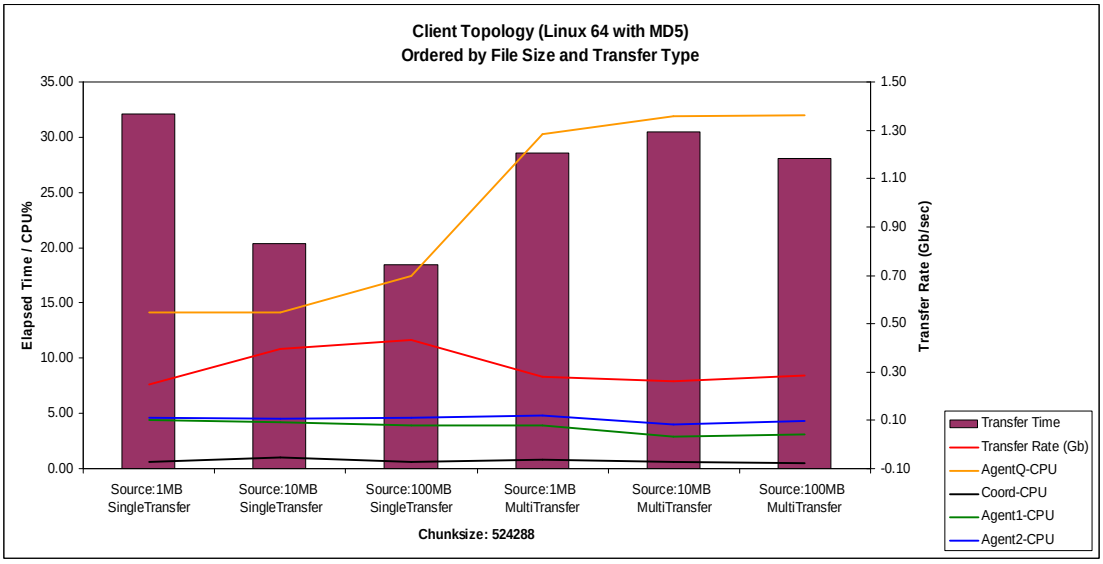

Figure 17 Linux64-MD5 512KB chunk size for Single and Multiple instance transfers

2.2.5 Test Summary

Looking across the results, the quickest transfers were attained at the following chunk sizes, file sizes and transfer types.

File Size and Transfer Type	Transfer Time	Chunk size
1MB Single Instance Transfer	23.36 s	524288 Bytes
10MB Single Instance Transfer	16.74 s	65635 Bytes
100MB Single Instance Transfer	15.80 s	262144 Bytes
1MB Multi Instance Transfer	21.25 s	131072 Bytes
10MB Multi Instance Transfer	20.52 s	131072 Bytes
100MB Multi Instance Transfer	18.86 s	65635 Bytes

Table 20 Linux64 Best transfer speeds for Single and Multiple instance transfers

File Size and Transfer Type	Transfer Time	Chunk size
1MB Single Instance Transfer	32.08 s	524288 Bytes
10MB Single Instance Transfer	20.32 s	524288 Bytes
100MB Single Instance Transfer	18.50 s	524288 Bytes
1MB Multi Instance Transfer	21.27 s	262144 Bytes
10MB Multi Instance Transfer	19.53 s	65635 Bytes
100MB Multi Instance Transfer	19.61 s	65635 Bytes

Table 21 Linux64-MD5 Best transfer speeds for Single & Multiple instance transfers

For single instance transfer, as with agent bindings, agents connecting in client mode have the best transfer time for larger chunk sizes. The difference is that there is now an overwhelming advantage for larger chunk sizes for multiple instance transfers. It should also be noted that as the file size increases, the difference between single and multiple threaded transfer times becomes less significant. The addition of the off-board MQ queue manager instance altered the CPU utilisation as expected, but did not give a decrease in transfer time.

3 Tuning Recommendations

3.1 WebSphere MQ Setup

Readers of this performance guide should make themselves familiar with the WebSphere MQ Performance Supportpacs that are continually released. They can be found here: http://www.ibm.com/support/docview.wss?rs=171&uid=swg27007197#1. Of particular interest for Linux is Supportpac MPL3 for MQ Version 7.5.

For this performance report, advice was taken from the aforementioned (MPL3) and applied to the queue managers created accordingly. Queue managers were created using the following crtmqm command:

```
crtmqm –q –u SYSTEM.DEAD.LETTER.QUEUE –lp 16 –lf 16384 <QueueManagerName>
```

Once the queue manager was created, tuning parameters were added to the queue managers' qm.ini as follows:

```
Channels:
MQIBindType=FASTPATH

TuningParameters:
DefaultPQBufferSize=1045876
DefaultQBufferSize=1048576
```

Note that the qm.ini was updated before the queue manager was started (and therefore before the WebSphere MQ Managed File Transfer objects were created).

By increasing the amount of memory available to queues for persistent and non-persistent messages, you can help to avoid writing messages out to disk unnecessarily. Turning on FASTPATH for channels removes the channel process, and enables the channel to run within the main queue manager process. Please consult your documentation to understand what this means for your WebSphere MQ installation.

For more information on tuning a WebSphere MQ queue manager, please refer to the Supportpacs mentioned above.

The use of high performance disks (SAN for example) is recommended for a WebSphere MQ installation. Separating out your /var/mqm/log directory structure from the regular file system is a well documented best practice that helps to create a queue manager that responds well to high throughput scenarios.

3.2 WebSphere MQ Managed File Transfer Setup

When running agents for this performance report, the following environment property was used:

```
export FTE_JVM_PROPERTIES="Xmx2048M Xms2048M"
```

This property was set before starting an agent and sets the starting and maximum JVM heap size to be 2GB. These values were used to ensure that the agent had sufficient memory to allocate when running the multiple transfer scenarios.

As demonstrated in the results, altering the agentChunkSize can have a significant impact on both CPU utilisation and transfer time. There is another property agentWindowSize that can be used to control the amount of syncpoints committed, and the number of acknowledgements sent between two agents when transferring files. This property has a default value of 10. This means that for every 10 chunks of data sent over WebSphere MQ, the sending agent will take an internal checkpoint, and wait to receive an acknowledgement from the receiving agent before sending more data. The property's default value was determined after extensive performance work during the development of version 7.0.1. Increasing this property increases the amount of data that could potentially need to be retransmitted if a recovery is required, and is not recommended for unreliable networks.

3.3 WebSphere MQ Managed File Transfer: Transfer Recommendations

The following are a list of bullet pointed recommendations when planning your WebSphere MQ Managed File Transfer network.

- Send large numbers of files over multiple transfers, rather then a single large transfer. This will increase the efficiency of the I/O involved in transferring the files, which will ultimately decrease the transfer time.

- Test your typical transfers using a range of agentChunkSize parameters. Depending on the underlying hardware, you may find an optimum value for your setup.

- Multiple smaller files place the agent under strain due to the operating system open/close costs associated with more files. Where possible configure your file creation processes to generate archives of smaller files, enabling WebSphere MQ MFT to use less open/close calls.

- Reading and writing to physical disk is often going to be the performance bottleneck. For agents that will see a large number of incoming, and outgoing transfers it would be best if high performance disks were used to read data from and write data to. This is demonstrated by the multiple/client and multiple/bindings scenarios. Due to the use of internal disks for WebSphere MQ, Client connectivity actually outperformed Bindings. This behaviour can be explained because the Agents reading/writing to the physical disks at the same time as a local MQ instance, causing an I/O bottleneck.

- When configuring your MQ network, use the appropriate WebSphere MQ Performance Report to apply optimal settings for your platform.

- Ensure that you have sufficient RAM for your agent. The performance tests used 2GB of RAM, it is recommended that you read your Operating System guide on memory usage and plan accordingly.

4 Measurement Environment

4.1 Agents

- WebSphere MQ Managed File Transfer Version 7.5 was used for this report.
- Default properties were used for agents, except for agentChunkSize.
- Agents were reading/writing files to the local file system, not the SAN.

4.2 WebSphere MQ

- WebSphere MQ Version 7.5 was used for the coordination queue manager.
- Queue managers created in accordance with Performance report.
- /var/mqm and /var/mqm/log were mounted on SAN disks.

4.3 Operating System

- Red Hat Enterprise Linux 5.3 (Tikanga) 32bit.
- Red Hat Enterprise Linux Server 5.5 64bit.

4.4 Hardware

eServer x3850:	MQPERFD2, MQPERFD8
Machine Type:	8864-4RG
Processor:	Intel® Xeon™ @ 3.33GHz
Architecture:	4 CPU
Memory (RAM):	4 GB
Disk:	Internal battery backed, cached disks for measurements
Network:	1 Gb Ethernet Adapter (onboard)

eServer x3650M3:	MQPERFLC, MQPERFLD
Machine Type:	7945-L2G
Processor:	Intel® Xeon™ X5660 @ 2.80GHz
Architecture:	6 CPU
Memory (RAM):	32 GB
Disk:	Internal battery backed, cached disks for measurements
Network:	1 Gb Ethernet Adapter (onboard)