

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

# A Scalable Application Server On Beowulf Clusters

A thesis presented in partial fulfilment of the requirement  
for the degree of

**Master of Information Science**

At Albany, Auckland, Massey University  
New Zealand

Supervised by: **Dr. Chris Messom**

**Michael Zhiyong Zhou**

2005

## **Abstract**

Application performance and scalability of a large distributed multi-tiered application is a core requirement for most of today's critical business applications.

I have investigated the scalability of a J2EE application server using the standard ECperf benchmark application in the Massey Beowulf Clusters namely the Sisters and the Helix. My testing environment consists of Open Source software: The integrated JBoss-Tomcat as the application server and the web server, along with PostgreSQL as the database. My testing programs were run on the clustered application server, which provide replication of the Enterprise Java Bean (EJB) objects.

I have completed various centralized and distributed tests using the JBoss Cluster. I concluded that clustering of the application server and web server will effectively increase the performance of the application running on them given sufficient system resources. The application performance will scale to a point where a bottleneck has occurred in the testing system, the bottleneck could be any resources included in the testing environment: the hardware, software, network and the application that is running.

Performance tuning for a large-scale J2EE application is a complicated issue, which is related to the resources available. However, by carefully identifying the performance bottleneck in the system with hardware, software, network, operating system and application configuration, I can improve the performance of the J2EE applications running in a Beowulf Cluster. The software bottleneck can be solved by changing the default settings, on the other hand, hardware bottlenecks are harder unless more investment are made to purchase higher speed and capacity hardware.

## **Acknowledgement**

My greatest appreciation goes to my supervisor: Dr Chris Messom, who has provided endless support both technically and morally. His broad and insightful understanding of the technology involved, his willingness to research emerging technology, his concise yet plain explanation to technically challenging concepts has made my learning process a pleasant experience.

Thanks to Dr Martin Johnson, Mr Andre Barczak and other people who have contributed greatly developing the Massey Beowulf computers: the Sisters and the Helix (the supercomputer) and made it available to us. Of course, thanks to the Allan Wilson Centre for Molecular Ecology and Evolution (AWC) for funding the development of the supercomputer.

Thanks to support from the system administrator James and the manager Lorri.

Also thanks to Massey University for the award of a Vice-Chancellor's Masterate Scholarship to support my study financially.

I must also thank the people working in the Open Source communities that have provided high quality open source software for my study, especially to the people working on Linux operating system, JBoss Group, Apache Tomcat Project and PostgreSQL database.

Finally, thanks to my family for their understanding and support. They are my source of energy and pleasure.

## Table of Contents

ABSTRACT .....	II
ACKNOWLEDGEMENT .....	III
TABLE OF CONTENTS.....	IV
TABLE OF TABLES.....	VII
TABLE OF FIGURES.....	IX
TABLE OF ABBREVIATIONS.....	XI
<b>CHAPTER 1: INTRODUCTION .....</b>	<b>1</b>
1.1 INTRODUCTION .....	1
1.2 MOTIVATION OF SCALABILITY STUDY FOR DISTRIBUTED APPLICATIONS.....	1
1.3 TODAY’S TECHNOLOGY SUPPORT FOR SCALABLE APPLICATION.....	2
1.4 SIGNIFICANCE OF MY STUDY.....	3
1.5 OVERALL STRUCTURE OF THE THESIS .....	4
1.6 SUMMARY .....	6
<b>CHAPTER 2: BACKGROUND KNOWLEDGE .....</b>	<b>7</b>
2.1 INTRODUCTION .....	7
2.2 COMPUTING TECHNOLOGY FOR A DISTRIBUTED SYSTEM .....	7
2.3 COMPARING J2EE WITH THE COMPETING TECHNOLOGIES .....	9
2.3.1 CORBA.....	9
2.3.2 J2EE.....	11
2.3.3 .NET.....	12
2.3.4 Comparing J2EE with CORBA.....	14
2.3.5 Comparing J2EE with .NET.....	16
2.4 CURRENT STUDY OF THE J2EE APPLICATION SERVER PERFORMANCE.....	18
2.5 MY RESEARCH APPROACH .....	22
2.6 SUMMARY .....	23
<b>CHAPTER 3: HARDWARE FOR THE TEST.....</b>	<b>25</b>
3.1 INTRODUCTION .....	25
3.2 INTRODUCTION OF SUPERCOMPUTERS AND BEOWULF CLUSTERS.....	25
3.3 MASSEY BEOWULF CLUSTER .....	27
3.4 SUMMARY .....	29
<b>CHAPTER 4: SOFTWARE FOR THE TEST .....</b>	<b>30</b>
4.1 INTRODUCTION .....	30
4.2 WHY CHOOSE OPEN SOURCE .....	30
4.3 OVERVIEW OF OPEN SOURCE SOFTWARE FOR THE J2EE APPLICATION SEVER .....	31
4.4 JBOSS APPLICATION SERVER.....	33
4.4.1 JBoss structure based on JMX standard.....	33
4.4.2 JBoss Clustering & Naming service.....	34
4.5 CHOOSING THE DATABASE.....	36
4.6 CHOOSING ECperf AS THE TESTING TOOL KITS.....	38
4.6.1 Why choose ECperf.....	39

4.7 SUMMARY .....	42
<b>CHAPTER 5: TEST DESIGN .....</b>	<b>43</b>
5.1 INTRODUCTION .....	43
5.2 TWO TYPES OF THE TESTING ARCHITECTURE .....	43
5.2.1 <i>Centralized workload architecture</i> .....	43
5.2.2 <i>Distributed workload architecture</i> .....	44
5.3 TESTING PROGRAMS AND RELATED CONFIGURATION .....	45
5.4 TEST DESIGN FOR SISTERS .....	46
5.4.1 <i>Type of planned test</i> .....	46
5.4.2 <i>Two hardware architectures for testing the sisters</i> .....	47
5.4.3 <i>Test design for JVM test</i> .....	49
5.4.4 <i>Test design for Clustering of Session Beans</i> .....	52
5.4.5 <i>Test design for Cluster of all EJB</i> .....	53
5.4.6 <i>Test design for two databases</i> .....	54
5.4.7 <i>Test with Two partitions and two databases</i> .....	56
5.5 TEST DESIGN FOR HELIX.....	58
5.5.1 <i>Type of planned test</i> .....	58
5.5.2 <i>Test design for JVM test</i> .....	59
5.5.3 <i>Test design for using the default DB pooling in JBoss</i> .....	60
5.5.4 <i>Test design using optimised database pooling in JBoss cluster</i> .....	61
5.6 CONCLUSION .....	61
<b>CHAPTER 6: TEST ON SISTERS .....</b>	<b>63</b>
6.1 INTRODUCTION .....	63
6.2 TEST WITH DIFFERENT JVM HEAP SIZE VALUE .....	63
6.3 TEST WITH CLUSTERING OF ONLY SESSION BEANS .....	68
6.3.1 <i>Preliminary tests using the default connections in PostgreSQL</i> .....	68
6.3.2 <i>Test by using optimised connections in PostgreSQL</i> .....	70
6.4 TEST WITH CLUSTERING OF ALL EJB .....	74
6.5 TEST WITH FIRST-AVAILABLE LOAD POLICY FOR CLUSTERING OF ALL EJB .....	76
6.6 TEST WITH TWO DATABASES .....	80
6.7 TEST WITH TWO PARTITIONS AND TWO DATABASES .....	82
6.8 TEST BY DISABLING THE LOG FILES WITH TWO PARTITION AND TWO DATABASES	85
6.9 DISCUSSION OF THE SISTERS RESULT .....	87
6.10 SUMMARY .....	89
<b>CHAPTER 7: TEST ON HELIX – THE SUPERCOMPUTER.....</b>	<b>90</b>
7.1 INTRODUCTION .....	90
7.2 JVM HEAP SIZE TEST .....	90
7.3 TEST USING THE DEFAULT DATA SOURCE POOLING VALUES .....	94
7.4 TEST USING OPTIMISED DATABASE POOLING .....	96
7.5 SUMMARY .....	98
<b>CHAPTER 8: PERFORMANCE ANALYSIS AND DISCUSSION.....</b>	<b>100</b>
8.1 INTRODUCTION .....	100
8.2 SCALABILITY ANALYSIS FOR SISTERS AND HELIX .....	100
8.3 BOTTLENECK ANALYSIS.....	104
8.4 PERFORMANCE TUNING FOR THE CURRENT TESTING SYSTEM .....	106
8.5 FURTHER PERFORMANCE IMPROVEMENT DISCUSSION .....	110
8.5.1 <i>Scaling the database</i> .....	110

8.5.2 <i>Scaling the JBoss application server</i> .....	112
8.6 POSSIBLE USE FOR COMMERCIAL APPLICATION .....	115
8.7 SUMMARY .....	117
<b>CHAPTER 9: CONCLUSION</b> .....	<b>119</b>
9.1 INTRODUCTION .....	119
9.2 CONCLUSION .....	119
9.2.1 <i>Contributions</i> .....	119
9.2.2 <i>Conclusion</i> .....	120
9.3 FUTURE WORK .....	121
<b>REFERENCE:</b> .....	<b>124</b>

## Table of Tables

Table 2.1: Comparison of basic features of J2EE and .NET

Table 2.2: Comparison of more critical features of J2EE and .NET

Table 6.1: Throughput of ECperf as function of the txRate

Table 6.2: JVM Heap Value vs. Maximum TPS Test

Table 6.3: Final Transaction output VS. txRate. (JBoss-Number = 1)

Table 6.4: Final Transaction output VS. txRate. (JBoss-Number = 1)

Table 6.5: Testing result of cluster all EJB on Sisters

Table 6.6: testing result of cluster all EJB using First Available load-balancing

Table 6.7: test result with 2 databases

Table 6.8: Transaction output with 2 Partitions and 2 Database systems

Table 6.9 is the transaction output with 2 partitions and 2 databases and disabled log

Table 7.1: JVM Heap Value VS. TPS Test

Table 7.2: Transaction output VS. JBoss Number

Table 7.3: Transaction Output VS. JBoss number in Helix

Table 8.1: Transaction Output in Sisters and Helix (JBoss=1)

Table 8.2: Transaction Output in Sisters and Helix (JBoss=2)



## Table of Figures

Figure 2.1: The CORBA object invocation mechanism

Figure 2.2: the J2EE application architecture

Figure 2.3: .NET Framework

Figure 4.1: The ECperf Architecture

Figure 5.1: Example configuration for the Centralized Workload

Figure 5.2: Example configuration for the Distributed Workload

Figure 5.3: Architecture for Centralized workload using Sisters

Figure 5.4: The architecture for distributed workload using Sisters

Figure 5.5 The hardware architecture for JVM heap value test in Sisters

Figure 5.6: Cluster of only Session Beans

Figure 5.7: Cluster of all Enterprise Beans

Figure 5.8: Clustering all EJB with two databases

Figure 5.9: Test with two partitions and two databases

Figure 5.10: EJB Replication with 2 JBoss Partition & 2 Databases

Figure 5.11: The hardware architecture for JVM heap value test in Helix

Figure 5.12: The hardware architecture for ECperf Test in Helix

Figure 6.1: Throughput as a function of the txRate (-Xmx = 180MB)

Figure 6.2: Maximum throughput as a function of the JVM Heap Size

Figure 6.3: The Transaction Output (TPS) VS. Client Number

Figure 6.4: The Transaction Output (TPS) VS. JBoss Number

Figure 6.5: Transaction Output VS. JBoss number when clustering all EJB

Figure 6.6: Transaction output VS. JBoss number

Figure 6.7: Transaction Output using the distributed architecture

Figure 6.8 Transaction Output with 2Partitions and 2 Databases

Figure 6.9: Test with 2 partition, 2 DB & disable log file.

Figure 6.10: All testing results for the sisters

Figure 7.1: Maximum throughput as a function of the JVM Heap Size

Figure 7.2: Transaction Output as a function of the JBoss number

Figure 7.3: Transaction Output VS. JBoss number in Helix

Figure 8.1: Helix and Sisters transaction output with one Jboss

Figure 8.2: Helix and Sisters transaction output with 2 Jboss

Figure 8.3: Example C-JDBC architecture

Figure 8.4: RAIDb-0 example

Figure 8.5 RAIDb-0-1 example

Figure 8.6 Architecture with JBoss and PostgreSQL cluster

## Table of Abbreviations

ANSI	American National Standards Institute
API	Application Programming Interface
BBop	Benchmark Business Operation
BSD	Berkeley Software Distribution
CCM	CORBA Component Model
CICS	Customer Information Control System
CLR	Common Language Runtime
CMP	Container-Managed Persistence
COM	Component Object Model
CORBA	Common Object Request Broker Architecture
COTS	Commercial off-the-shell
CPU	Central Processing Unit
CSIRO	Commonwealth Scientific & Industrial Research Organization
DBMS	Database Management System
DCOM	Distributed Component Object Model
EJB	Enterprise JavaBeans
HPC	High Performance Computing
HTTP	Hypertext Transfer Protocol
IMS-TM	Information Management System Transaction Manager
INRIA	French National Institute For Research In Computer Science And Control
J2EE	Java 2 Platform, Enterprise Edition
JDBC	Java Database Connectivity
JDK	Java Development Kit
JMS	Java Message Service
JMX	Java Management Extensions
JNDI	Java Naming and Directory Interface
JRE	Java Runtime Environment
JSP	Java Server Pages
JVM	Java Virtual Machines
LAN	Local Area Network
MPI	Message-passing Interface

MPP	Massively Parallel Processing
ODBC	Open Database Connectivity
OMA	Object Management Architecture
OMG	Object Management Group
ORB	Object Request Broker
PBS	Portable Batch System
PHP	Hypertext Preprocessor
PVM	Parallel Virtual Machine
RAIDb	Redundant Arrays of Inexpensive Database
RAM	Random Access Memory
RDBMS	Relational Database Management System
RMI-IIOP	Remote Method Invocation Over Internet Inter-Orb Protocol (Rmi Over Iiop)
SARs	Storage Area Network
SFTP	Secure File Transfer Protocol Message-passing Interface
SMP	Symmetric Multiprocessing
SOAP	Simple Object Access Protocol
SPI	service provider interface
SSH	Security Shell
SUT	System Under Test
UDDI	Universal Description Discovery And Integration
UNIX	Uniplexed Information and Computing System. (It was originally spelled "Unics.")
WAN	Large Area Network
WSDL	Web Services Description Language
XML	Extensible Markup Language

# Chapter 1: Introduction

## 1.1 Introduction

---

This thesis presents a study of the performance and scalability of J2EE applications. In particular, I concentrate on the application server, which is the core component of the J2EE architecture. I use a cluster of JBoss application servers to test how the scalability and performance of a J2EE application is effected.

In this introductory chapter, I start with the motivation of the scalability study and explain why it is important in the business world. Then I give some brief technical review about how scalability can be achieved using current available hardware and software. I explain why my particular study is useful and finally give an overview of the contents in each chapter.

## 1.2 Motivation of scalability study for distributed applications

---

Large-scale distributed systems are becoming increasingly important in the world, especially with online business activities. The Internet has greatly improved the accessibility to online businesses, and the increased accessibility has promoted ever-increasing e-commerce applications. The performance and scalability of an application is critical for a successful business, as a business application needs to have high performance to achieve competitive advantages over their competitors.

Performance can refer to many aspects, such as scalability, availability, fault tolerance and load balancing. I am particular interested in the scalability of an application. A scalable application has the capacity to serve additional users or transactions without fundamentally altering the application's architecture or program design. If an application is scalable, you can maintain steady performance as the load increases simply by adding additional resources such as servers, processors or memory.

The two most common types of scalability that can be applied to affect the overall application performance are:

- Horizontal scalability: Adding more servers (web, application or database servers) to improve performance.
- Vertical scalability: Adding more physical resources (memory, processors or network cards) to an existing server to improve performance.

The key point of scalability is to decide how well an application will perform when the size of the problem increases. Scalability is not only critical to maintain current system functionality in a changing workload, but also a key factor to guarantee the system can keep up with the growth potential and has the ability to scale to meet future user's demand.

### 1.3 Today's technology support for scalable application

---

Today's technology has provided high-quality hardware and software to support the development and deployment of applications with good scalability and high performance.

For the computer hardware, we have consistently increasing computing power with the CPU speed doubling every 18 months, while the price of a personal computer is gradually getting cheaper. Various architectures built on PCs have provided fundamental support for high performance computing.

Supercomputers, which are the most powerful computers in the world, are getting more powerful. Beowulf Clusters, which are built using the Commercial off-the-shelf (COTS) components such as PCs, are gradually becoming more important in the supercomputer field [7]. A major merit of a Beowulf Cluster is its significant cost advantages over traditional mainframe supercomputers with similar computing capacity.

The wide adoptions of fast network connections, for either local or large area networks, as well as the Internet technology have enabled reliable communication facilities to support high performance applications. Combined with the supercomputer and the reliable Internet connections, it is much more practical to build a GRID [20], a network of supercomputers using today's technology.

Software has been developed to take advantage of the hardware architecture to achieve high performance and scalability. Using a cluster of application servers for a J2EE application, the application server components such as an EJB can be replicated across a cluster of application server machine. By load balancing the client request to members in the application server cluster, each client can interact concurrently with local copies of the same EJB component. This results in increased accessibility to computing power, thus, an increased application performance and scalability can be achieved.

I am going to investigate the J2EE application performance using open source software. JBoss, the leading open source application server has recently introduced cluster support, which I will use for my study.

#### 1.4 Significance of my study

---

I am going to investigate application scalability using open source software running in the Beowulf Cluster. The advantage of this approach is that I have total control of the resource, because the Beowulf Cluster was built and maintained by our department in the Massey University, and the open source software can be used free of charge with access to the source code.

From a business point of view, I am using one of the most cost effective combined hardware and software for running J2EE applications. Building a Beowulf Cluster cost only 5% to 20% of total cost compared with traditional mainframe supercomputers with the same computing power [21]. The JBoss application server is

free of charge but with most of the features a leading commercial application server provides. PostgreSQL is the most advanced open source database. By running J2EE applications using a Beowulf cluster as the hardware, the JBoss cluster, PostgreSQL as software, I can expect good scalability results. A good scalability result means that the system could be very useful for developing and deploying cost effective commercial applications.

## 1.5 Overall structure of the thesis

---

The thesis is organised into the following nine chapters:

**Chapter 1** introduces the overall structure of the thesis. I start with the motivation for the scalability study, followed by current hardware and software technology that can be used to build scalable applications. I then give reasons why my particular approach is useful, and finish with the overall thesis structure.

**Chapter 2** presents some of the background knowledge necessary for understanding my study. Three of the most important architectures for building large-scale distributed applications are presented and compared, this information helps to identify why I chose the J2EE architecture for my study. I give some of the related literature review and also state my research hypothesis.

**Chapter 3** gives a detailed description of the hardware architecture of my study. I use the Beowulf Cluster computers in Massey University for my performance study. Starting with the general architectures of various high performance supercomputers, the advantages of cluster-based system are discussed. At last, the helix and sisters clusters in Massey University are introduced in details.

**Chapter 4** covers the software used in the study. I have chosen all software from open source, which I am particularly interested in. I cover the software for running a distributed applications based on J2EE technology. To be more specific, I give some detail about why I choose the integrated JBoss-Tomcat as the application server and web server, the PostgreSQL as the database and the ECperf as my testing application.



**Chapter 5** describes the details of the test design for both the Sisters and Helix. I start with different types of hardware architecture I will use, followed by some detailed information about how to run various test programs in the system. The last part gives my preliminary design selections on the type of test, and briefly identifies the reason for that selection.

**Chapter 6** gives detailed testing procedures for my study in Sisters. I have done various tests based on different hardware architecture, software and the application configurations. For each type of test, I present with details about the test design and implementation procedures. Followed by a test result, Analysis and discussion on these results reveal several important conclusions.

**Chapter 7** gives detailed testing procedures for my study in Helix. Again, I have followed a similar approach used for the Sisters. But the Helix concentrates on some different aspect of my study and reveals some different results as compared with the Sisters.

**Chapter 8** covers the further analysis and discussions based on the results obtained on both Sisters and Helix. I have given a broader view on how to further improve the performance and scalability of my current system. A higher level of discussions about how to improve the current implementation and use better software features and hardware architecture are discussed.

**Chapter 9** gives the conclusion to my study. Based on the analysis of testing results in the previous chapters, I make my final conclusions and show some of the contributions made by my study. I also anticipated the future work in my research field.

## 1.6 Summary

---

I have introduced the overall structure of the thesis. Firstly, I gave the motivation for the scalability study, followed by the technical support that can be used for building a scalable application. I then show why my particular approach is useful. Finally, I listed the major topics of each chapter in the thesis.