

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.



J2EE Application for Clustered Servers

--- Focus on balancing workloads among clustered servers

A thesis presented in partial fulfillment of the requirements for the degree of

Master of Information Science

in

Computer Science

at Massey University, Albany,

New Zealand.

Supervised By: **Dr Chris Messom**

Xi Chen

[2006]

Acknowledgment

Special thanks go to my supervisor **Dr. Chris Messom** for his guidance, enthusiasm and technical support during the year. He impressed me with his broad knowledge, background and attitude on every issue. He teaches me not only the latest technologies but also solving problem skills. Thanks a lot.

I also would like to thank the system administrator James Chai for his support to make my project run smoothly, e.g. increasing shared memory size and resetting network routes etc.

Thanks also go to all the developers who have contributed to open source software.

Meanwhile, I appreciate my parents who gave me full support and took care of my sons during my study.

Abstract

J2EE has become a de facto platform for developing enterprise applications not only by its standard based methodology but also by reducing the cost and complexity of developing multi-tier enterprise applications. J2EE based application servers keep business logic separate from the front-end applications (client-side) and back-end database servers. The standardized components and containers simplify J2EE application design. The containers automatically manage the fundamental system level services for its components, which enable the components design to focus on the business requirement and business logic.

This study applies the latest J2EE technologies to configure an online benchmark enterprise application – MGProject. The application focuses on three types of components design including Servlet, entity bean and session bean. Servlets run on the web server Tomcat, EJB components, session beans and entity beans run on the application server JBoss and the database runs on the database server PostgreSQL. This benchmark application is used for testing the performance of clustered JBoss due to various load-balancing policies applied at the EJB level.

This research also focuses on studying the various load-balancing policies effect on the performance of clustered JBoss. As well as the four built-in load-balancing policies i.e. *FirstAvailable*, *FirstAvailableIdenticalAllProxies*, *RandomRobin* and *RoundRobin*, the study also extend the JBoss *LoadbalancePolicy* interface to design two dynamic load-balancing policies. They are *dynamic* and *dynamic weight-based* load-balancing policies.

The purpose of dynamic load-balancing policies design is to ensure minimal response time and obtain better performance by dispatching incoming requests to the appropriate server. However, a more accurate policy usually means more communications and calculations, which give an extra burden to a heavily loaded application server that can lead to drops in the performance.

Table of Contents

ACKNOWLEDGMENT	I
ABSTRACT.....	II
TABLE OF CONTENTS	III
LIST OF TABLES.....	VI
LIST OF FIGURES.....	VII
CHAPTER 1: INTRODUCTION.....	1
1.1 WHY CHOOSE J2EE.....	1
1.2 THE PURPOSE OF THE STUDY	2
1.3 HOW THE THESIS IS ORGANIZED	3
CHAPTER 2: BACKGROUND OVERVIEW.....	5
2.1 J2EE	5
2.1.1 <i>J2EE Platform</i>	5
2.1.2 <i>The J2EE Containers and APIs</i>	7
2.1.2.1 EJB 2.1	8
2.1.2.2 Servlet 2.4	10
2.1.3 <i>J2EE Deployment Structure</i>	11
2.2 BENCHMARK J2EE APPLICATION	13
2.2.1 <i>ECperf</i>	13
2.2.2 <i>The ECperf Design</i>	13
2.3 LOAD-BALANCING POLICIES IN DISTRIBUTED SYSTEM	15
2.3.1 <i>Distributed System</i>	15
2.3.1.1 Distributed System with N-tier Architecture	16
2.3.1.2 Clustered Distributed System	16
2.3.2 <i>Common Load-balancing Policies</i>	18
2.3.2.1 Hardware Load Balancer	18
2.3.2.2 Software Load Balancer.....	19
CHAPTER 3: HARDWARE & SOFTWARE FOR THE STUDY.....	21
3.1 HARDWARE.....	21
3.2 SOFTWARE.....	22
3.2.1 <i>IDE – NetBeans 4.1</i>	22
3.2.2 <i>Database – PostgreSQL 7.4</i>	23
3.2.3 <i>Web Server – Tomcat 5.5</i>	24
3.2.3.1 Tomcat Clustering	25
3.2.4 <i>J2EE Application Server – JBoss 4.0</i>	26
3.2.4.1 Clustering in JBoss	26
3.2.4.2 Load-balancing Policy in JBoss.....	27
3.2.4.3 Transaction Commit Options in JBoss	29
3.2.5 <i>Performance Test Tool – Jmeter 2.0.3</i>	30
CHAPTER 4: BENCHMARK J2EE APPLICATION	32
4.1 CASE STUDY	32
4.2 DATABASE DESIGN.....	35
4.2.1 <i>Database Implementation</i>	37

4.3 J2EE COMPONENTS DESIGN & IMPLEMENTATION	38
4.3.1 EJB Design	38
4.3.1.1 Entity Bean	39
4.3.1.2 Session Bean	45
4.3.1.2.1 Stateless Session Bean	47
4.3.1.2.2 Stateful Session Bean	50
4.3.2 EJB Implementation in the MGProject Application	51
4.3.2.1 Entity Bean	51
4.3.2.2 Session Bean	53
4.3.3 Web Component Design	54
4.3.3.1 Servlets Design	55
4.3.3.2 User Interface Design - Servlets Implementation	58
CHAPTER 5: LOAD-BALANCING POLICIES DESIGN.....	67
5.1 DESIGN CONSTRUCTION – USING EJB TIMER SERVICE & JBOSS TREECACHE.....	67
5.1.1 The EJB Timer Service	68
5.1.2 The JBoss TreeCache.....	68
5.2 DYNAMIC LOAD-BALANCING POLICY DESIGN.....	69
5.2.1 The JBoss LoadBalancePolicy Interface	70
5.2.2 Pseudo Coding for Dynamic Load-Balancing Policy.....	71
5.3 DYNAMIC WEIGHT-BASED LOAD-BALANCING POLICY DESIGN	72
5.3.1 Implementation for Dynamic Weight-based Load-Balancing Policy.....	72
CHAPTER 6: TEST PLAN	74
6.1 TEST PLAN IN JMETER.....	74
6.2 STRESS TEST PLAN	76
6.2.1 Plan for Jmeter	77
6.2.2 Plan for Tomcat	78
6.2.3 Plan for JBoss.....	79
6.2.4 Plan for PostgreSQL.....	81
6.2.5 Test Plan Summary	82
CHAPTER 7: TEST RESULTS AND DISCUSSION.....	85
7.1 TO DETERMINE RAMP-UP PERIOD AND LOOPS VALUE	86
7.1.1 For the Test Plan with Thinking Time	86
7.1.2 For the Test Plan Without Thinking Time	87
7.1.3 Choose the Appropriate Version for the Performance Test	89
7.2 TO DETERMINE TOMCAT LEVEL CONFIGURATION FOR STRESS TEST ...	89
7.2.1 Single Tomcat	90
7.2.2 Clustering Two Tomcat with Apache.....	91
7.2.3 Directly Load Two Tomcat in Jmeter	92
7.3 SCALABILITY TEST	93
7.3.1 Performance Test for a Single JBoss.....	94
7.3.2 Performance Test for Clustered Two JBoss	95
7.3.3 Comparison and Discussion.....	97
7.4 TESTING ON EQUAL MACHINE LOAD	98
7.4.1 Performance Test for Two JBoss with Built-in Policies	99
7.4.2 Performance Test for Two JBoss with Dynamic Policies.....	101
7.4.3 Comparison and Discussion.....	102
7.5 TESTING ON UNEQUAL-LOAD MACHINES	104

7.5.1 Performance Test for Built-in Policies	105
7.5.2 Performance Test for Two JBoss with Dynamic Policies.....	106
7.5.3 Comparison and Discussion	108
CHAPTER 8: CONCLUSIONS	110
8.1 CONCLUSIONS FOR J2EE APPLICATION DESIGN	110
8.2 CONCLUSION FOR LOAD-BALANCING POLICIES.....	111
8.3 FURTHER WORK.....	112
REFERENCE.....	115
APPENDIX.....	119
A. IDE – NETBEANS	120
B. POSTGRESQL	121
C. JBOSS	123
D. APACHE & TOMCAT	125
E. JMETER	127
F. DATABASE - MG.SQL.....	128
G. SAMPLE BASH FILE.....	132

List of Tables

Table 1: Hardware in Beowulf Cluster - Sisters.....	22
Table 2: Machines Allocation for Jmeter Parameter-Setting Test	86
Table 3: The Comparison of Performance for Different Version	89
Table 4: Stress Test with Only One Tomcat	90
Table 5: Performance Comparison among Various Web Level Configurations.....	92
Table 6: Stress Test with Starting Two Tomcat from Jmeter	93
Table 7: Machine Allocation for Scalability Test	94
Table 8: The Quantitative Analysis for Scalability Testing Systems	98
Table 9: The Throughput Comparison for Dynamic Policy.....	107
Table 10: The Throughput Comparison for Dynamic Weight-Based Policy.....	107

List of Figures

Figure 1: Overview of J2EE Application Architecture	6
Figure 2: Modules view of a J2EE Application	12
Figure 3: Four Domains Workloads in ECperf	14
Figure 4: Middleware View of Distributed System.....	15
Figure 5: Working Principle of Hardware Load Balancer	19
Figure 6: Tomcat Clustering.....	25
Figure 7: Load Balancer Implementation in JBoss	28
Figure 8: Use Case Diagram of the Benchmark Application	34
Figure 9: Class Diagram Showing Database Design.....	34
Figure 10: The Strict Communication in This Benchmark Application.....	39
Figure 11: Example Deployment Descriptor of Entity Bean	40
Figure 12: Details of CMP Bean.....	41
Figure 13: Example of Local Home Interface Design for Entity Bean	42
Figure 14: Example of Local Component Interface Design for Entity Bean.....	43
Figure 15: Example of Bean Class Design for Entity Bean.....	45
Figure 16: Example of Remote Home Interface for Stateless Session Bean	48
Figure 17: Example of Component Interface for Stateless Session Bean	48
Figure 18: Example of Bean Class for Stateless Session Bean	49
Figure 19: Set HTTP Session for Stateful Session Bean.....	50
Figure 20: Get HTTP Session for Stateful Session Bean	50
Figure 21: XDoclet Tags for Entity Bean	52
Figure 22: The Next Primary Key Generator System	53
Figure 23: Session Beans Design and the Referenced Entity Beans.....	54
Figure 24: Example Codes for a Servlet Design and Session Management.....	57
Figure 25: Interface Design.....	58
Figure 26: The Home Page of MGProject	60
Figure 27: New Order Servlet.....	61
Figure 28: The Interface of priceQuote Servlet	61
Figure 29: addToCartServlet Appending with checkCartServlet.....	62
Figure 30: The Interface of registerSuccessServlet	63
Figure 31: The Interface of orderList.....	64
Figure 32: The Interface of cancelSelect Servlet.....	65
Figure 33: Display the Status of Outstanding Orders for the Customer	65
Figure 34: The Interface of custDetail Servlet	66
Figure 35: EJB Timer Service Creates an Interval Timer.....	68
Figure 36: The Transactions Processed after Time Out.....	68
Figure 37: Using TreeCache as JBoss MBean Service	69
Figure 38: Pair of Information in JBoss TreeCache.....	70

Figure 39: Overloaded Method chooseTarget in JBoss <i>LoadBalancePolicy</i> Interface.....	70
Figure 40: Downloaded Information in Client Stub.....	71
Figure 41: Pseudo Code for Dynamic Load-Balancing Policy	71
Figure 42: Calculation the Ration for the Weight-Based Policy	72
Figure 43: Design a serializable object cRatio	73
Figure 44: The Interface of Jmeter.....	74
Figure 45: An Example of <code>__regexFunction</code> using in Jmeter	75
Figure 46: Clustered JBoss with Reside Tomcat for Web Application	78
Figure 47: Define Read-Only Methods via XDoclet.....	79
Figure 48: Read-Only Methods in the <code>jboss.xml</code>	80
Figure 49: Define EJB Container Configuration for EJB	80
Figure 50: Clustering Stateless Session Beans with Load-balancing Policy via XDoclet.....	81
Figure 51: Clustering Stateful Session Beans with Load-balancing Policy via XDoclet	81
Figure 52: Row Level Locking for SeqidBean Entity Bean	82
Figure 53: The Scenario of Testing Configuration.....	83
Figure 54: Ramp-up effect on Throughput	Figure 55: Ramp-up effect on Average Delay
Figure 56: Ramp-up effect on Throughput	Figure 57: Ramp-up effect on Average Delay
Figure 58: Loops Effect on Throughput	Figure 59: Loops Effect on Average Delay.....
Figure 60: The Architecture of Stress Testing System.....	93
Figure 61: Throughput from Single JBoss	Figure 62: Average Delay from Single JBoss
Figure 63: Throughput from Two JBoss	Figure 64: Average Delay from two JBoss
Figure 65: Throughput Comparison	Figure 66: Average Delay Comparison
Figure 67: Policies Effect on Throughput	99
Figure 68: Policies Effect on Average Delay	100
Figure 69: Policies Effect on the Throughput	Figure 70: Policies Effect on Average Delay.....
Figure 71: The Policies Effect on the Throughput without Extra Load.....	102
Figure 72: The Policies Effect on the Average Delay without Extra Load.....	103
Figure 73: Throughput of Built-In Policies.....	105
Figure 74: Average Delay for Built-In Policies.....	105
Figure 75: Throughput of Various Policies.....	108
Figure 76: Average Delay of Various Policies.....	109
Figure 77: Adding <code>check-dirty-after-get</code> to JBoss Deployment Descriptors	113
Figure 78: Try to Create Entity with <code>unknown-pk</code> for J2EE Application	114

Chapter 1: Introduction

Summary: The chapter describes some brief reasons why the J2EE platform was selected. And then presents the purpose of the study and aims of the study. Finally, the overall structure of the thesis and main contents in each chapter are listed.

1.1 Why Choose J2EE

Since the 1990s, middleware technology i.e. middle-tier software has developed to simplify distributed assembly of components, which is a collection of services including managing communication, security and threads etc. to enable multiple processes on different working machines to interact across the network. For a distributed enterprise application, usually we use middleware to connect separate applications such as linking between web applications to a database system, the technique provides an abstraction capability to simplify the construction of a distributed enterprise system and allows the application developers to only focus on business logic [52]. The most attractive middleware today are CORBA, J2EE and .NET. All of them are deployed as standard components / objects. But CORBA (Common Object Request Broker) provides only middleware techniques to model a standard and consistent component architecture framework such as clients' request for services from servers via well-defined interfaces across network. J2EE and .NET are the two most popular software development platforms to design server-side enterprise applications.

Both J2EE and .NET are emerging and competing platforms that contribute to simplifying writing enterprise applications. .NET is restricted to MS Windows-based platforms although some limited open source multi-platform .NET systems are in development, such as the Mono Project [16]. This research will utilize Massey Beowulf Clustered computers that are running the Linux operating system. In addition, the advantage of J2EE is that it is programmed in Java that can be deployed cross-platform. Moreover, this research extends a previous study, Zhou's Master's research that used SUN ECPe --- a benchmark J2EE application as the testing application to study "A scalable application server on Beowulf Cluster" [1]. As a result, the J2EE platform is a natural choice for the study.

1.2 The Purpose of the Study

This research addresses techniques that support the development of distributed enterprise applications particularly its non-functional requirements including security, scalability, fault tolerance and load-balancing etc but focuses specifically on load-balancing policies. Due to the dramatically dropping hardware price, using more than one server has become affordable for more and more companies, including most small companies. Moreover, the Internet is providing a potential e-market for enterprises since the network has effectively shortened distances for international or national trade. The quality of e-commerce services, such as securing customer information, ensuring minimal online waiting time to increase returning customers has been a priority. Improving the performance of servers, particularly clustered servers are now being addressed. Load balancing is one of the key technologies behind clustering that affects the performance of the clustered servers when the servers face heavy loads.

The goal of load balancing among clustered servers is to ensure minimal response time and obtain better performance by dispatching incoming requests to the appropriate server. When Zhou [1] studied the scalability of Beowulf clustered servers using JBoss, he found the different load-balancing policies could get different performance results. His study indicated that a better scalability result could be achieved by using First Available Load-Balancing policy than using default Round Robin Load-balancing policy, but both of them became a bottleneck under heavy workloads [1]. In order to improve this situation, this research will focus on studying the load-balancing mechanism. The study will extend the JBoss *LoadbalancePolicy* interface to design a more appropriate load-balancing policy --- dynamic and dynamic weight-based load-balancing policies.

In order to test the performance of clustered servers using various load-balancing policies, this study also includes building a benchmark J2EE application. It consists of:

- Client side: client view of dynamically generated HTML pages using server side technology Servlets that is implemented on Tomcat, a servlet container.
- Server-side: refers to both web server and application server. The web server (Tomcat) focuses on user interface design via web component Servlets and the application server (JBoss) which handles business processes via designing J2EE enterprise Java beans components entity beans and session beans
- Database: used for store business persistent data that is implemented on the open

source database application PostgreSQL

1.3 How the Thesis Is Organized

The thesis focuses on the J2EE application and load-balancing policies design. The entire thesis will be focused on these two main topics and organized into eight chapters. The overall structure is as follows:

- **Chapter 1** describes some brief reasons why the J2EE platform was selected. And then presents the purpose of the study and aims of the study. Finally, the overall structure of the thesis and main contents in each chapter are listed.
- **Chapter 2** overviews the knowledge background and terminology needed to understand this thesis. J2EE is introduced including the architecture of the J2EE platform, containers and APIs and the structure of deployment of J2EE applications. Secondly the well-known benchmark J2EE application – ECperf design is introduced. And finally the n-tier clustered distributed system and common load-balancing policies in use are presented.
- **Chapter 3** covers the hardware system for this study - the clustered systems built in Massey University Albany campus. The detail of the open source software chosen for this study including J2EE application design IDE, web server Tomcat, application server JBoss, database PostgreSQL and performance test tool Jmeter are discussed.
- **Chapter 4** details the entire benchmark J2EE application – MGProject design, including database, EJB and web components design. A brief introduction of the EJB and Servlet design, including the life cycle of the EJBs and Servlets, including the main concerns on designing these components is given. In addition, example coding of the design to explain the components are given. Moreover, how the web component and EJB components interact with each other in the application and how to use XDoclet in the NetBeans IDE to generate JBoss specific deployment descriptors is presented.
- **Chapter 5** covers the detailed design of the dynamic and dynamic weight-based load-balancing policies for clustered JBoss. The policies will extend JBoss *LoadBalancePolicy* interface and utilizes EJB timer service and JBoss *TreeCache*. The *LoadBalancePolicy* interface exposes the cluster members' information. The EJB timer service performs tasks in a regular period and JBoss

TreeCache provides shared storage for all timer session beans located in different machines. In addition, the chapter also lists the pseudo code design.

- **Chapter 6** The test plan in this chapter has two meanings. One refers to establishing a test plan in Jmeter. The plan will simulate a large number of independent clients performing online actions to interact with back-end servers. Another means to plan a test procedure for testing the performance of JBoss under various load-balancing policies. This plan should provide a sequence of testing steps to achieve the testing goal. Meanwhile, techniques to ensure the benchmark application runs as smoothly as possible under the very heavy loads are discussed.
- **Chapter 7** presents the details of the tests implementation on the Sisters cluster. The main tests are classified into two kinds – scalability tests and load-balancing policy comparison tests. The scalability refers to a single JBoss and clustering two or more JBoss performance test. The load-balancing policy tests will apply various load-balancing policies on the EJB level to identify if each of them affects the final performance of clustering JBoss.
- **Chapter 8** gives the final conclusions of the thesis. The conclusions include a critique of the J2EE application design and the effect of the various load-balancing policies on the system performance. Finally, the chapter also presents potential further work.