2005-03-10

A Statistical Performance Model of Homogeneous Raidb Clusters

Brandon Lamar Rogers
Brigham Young University - Provo

A STATISTICAL PERFORMANCE MODEL OF HOMOGENEOUS RAIDb CLUSTERS

by

Brandon L. Rogers

A thesis submitted to the faculty of

Brigham Young University

in partial fulfillment of the requirements for the degree of

Master of Science

School of Technology

Brigham Young University

April 2005

BRIGHAM YOUNG UNIVERSITY

GRADUATE COMMITTEE APPROVAL

of a thesis submitted by

Brandon L. Rogers

This thesis has been read by each member of the following graduate committee and by majority vote has been found to be satisfactory.

Date	Joseph J. Ekstrom, Chair
Date	Michael G. Bailey
Date	Ronald F. Gonzales

BRIGHAM YOUNG UNIVERSITY

As chair of the candidate's graduate committee, I have read the thesis of Brandon L. Rogers in its final form and have found that (1) its format, citations, and bibliographical style are consistent and acceptable and fulfill university and department style requirements; (2) its illustrative materials including figures, tables, and charts are in place; and (3) the final manuscript is satisfactory to the graduate committee and is ready for submission to the university library.

_____ _____

Date

Joseph J. Ekstrom
Chair, Graduate Committee

Accepted for Department

_____

Thomas L. Erekson
Director, School of Technology

Accepted for the College

_____

Douglas M. Chabries
Dean, Ira A. Fulton College of
Engineering and Technology

ABSTRACT

A STATISTICAL PERFORMANCE MODEL OF

HOMOGENEOUS RAIDb CLUSTERS

Brandon L. Rogers

School of Technology

Master of Science

The continual growth of the Internet and e-commerce is driving demand for speed, reliability and processing power. With the rapid development and maturation of e-commerce, the need for a quick access to large amounts of information is steadily rising. Traditionally, database systems have been used for information storage and retrieval. However, with online auctions, rapid Internet searches, and data archival, the need for more powerful database systems is also growing.

One type of distributed database is called *Redundant Arrays of Inexpensive Databases* (RAIDb). RAIDb clusters are middleware-driven to promote interoperability and portability. RAIDb clusters allow for multiple levels of data replication and publish the clustered system as a single, coherent database system.

In this thesis, performance models are created for RAIDb level 1 and level 2 clusters. A statistical three-factor, two-level factorial design is used to evaluate the

significance of several factors in a RAIDb cluster. These significant factors are then used to create a regression analysis, and eventually a regression equation, that can be used to predict the performance of RAIDb clusters. This performance model should be a useful predictive tool since the results have a 99% confidence interval.

ACKNOWLEGEMENTS

There are many people who have given me endless support during the writing of this thesis. I would like to thank my graduate chair, Dr. Joseph Ekstrom, for the confidence he holds in my ability and the time he has freely provided. To each of the other members of my committee, Dr. Michael Bailey and Dr. Ronald Gonzales, I thank for your support and your examples of professionalism. I am infinitely grateful to my wife, Kathryn, for the encouragement and support given on her behalf during late nights, early mornings, and 48-hour days. Finally, I would like to thank my parents, my sister, and my brothers for the unending love and support offered in my behalf.

Table of Contents

List of Tables

List of Figures

Chapter 1

1 INTRODUCTION

1.1 Background

The continual growth of the Internet and e-commerce is driving demand for speed, reliability and processing power in backend server systems. With the rapid development and maturation of e-commerce, the need for a quick access to large amounts of information is steadily rising. Traditionally, database systems have been used for information storage and retrieval. However, with online auctions, rapid Internet searches, and data archival, the need for more powerful database systems is also growing.

Many types of database architectures exist, such as single-node databases, hot-standby databases, and distributed databases. Single-node databases are common and easy to configure; hot-standby databases provide high availability; distributed databases provide a mechanism to store large amounts of data as well as provide high throughput and availability. Distributed databases have grown in popularity because of their characteristic performance enhancements (Ault, 2003) and their natural similarities to parallel processing systems, which have been used historically for implementation of very large search systems (Nielsen, 2003).

One emerging type of distributed database system is *Redundant Arrays of Inexpensive Databases* (RAIDb). RAIDb promises to limit cost and provide increased

system capability, RAIDb was proposed in September of 2003 at INRIA (Institut National De Recherche En Informatique Et En Automatique) by Emmanuel Cecchet, Julie Marguerite, and Willy Zwaenepoel (2003). RAIDb is, by specification, a middleware software layer responsible for the coordination of several database computers to provide a coherent, distributed database solution.

The concept of RAIDb is closely related to other clustering solutions. RAIDb includes a collection of independent nodes which are coordinated through a message layer to appear as a single, autonomous system. RAIDb clusters take advantage of each node's limited processing capability to provide power sufficient to accomplish large processing tasks. Hundreds of nodes in a clustered computer may appear to be a single, very powerful machine (Nielsen, 2003).

RAIDb clusters are distinguished by their price/performance ratio and their ability to mix node types and configurations. This heterogeneous nature, combined with the parallel processing performance improvements, is particularly remarkable because clusters can be built with inexpensive, off-the-shelf computers and networking hardware (Cecchet, 2003).

The RAIDb specification has been implemented and released as a Java-based software product called Clustered Java Database Connectivity (C-JDBC). As a middleware messaging layer, C-JDBC is different than many clustered computing environments. C-JDBC nodes can consist of machines running differing operating systems, hardware configurations, and database management systems. Also, C-JDBC does not need operating system-specific drivers, algorithms, or message-passing libraries

to function. Rather, all communications libraries and drivers are implemented as middleware. As middleware, the C-JDBC drivers and libraries are platform-independent.

1.2 Problem Statement

As database clusters increase in popularity, obtaining relevant performance evaluations of these clusters is extremely important. Some of the most common reasons for deploying a database cluster are to increase performance, to increase availability, and to reduce costs. Having an accurate measure of how these database systems will perform is essential to the business case for deployment and configuration of these clusters.

In 1998, Chung-Min Chen and associates published a report comparing the effectiveness of three distributed database management system (DBMS) approaches: the federated approach, the gateway approach, and the middleware approach. Chen explained how the federated and gateway approaches use inter-database communication to handle global queries. Chen stated the following concerning the middleware approach:

> The middleware, in contrast, is merely a software that coordinates the DBMS servers for collaborative activities. A middleware [approach] is *not* a fully loaded DBMS and must rely on the component DBMS servers to evaluate global queries. Middleware products are appealing to many [multidatabase] users as they are less expensive and more portable than the other two alternatives. However, there has been some doubt about the use of middleware due to a performance concern: the lack of an internal DBMS for efficient global query processing.

Although Chen casts doubt on the effectiveness of the middleware approach, recent developments in memory-based database servers, such as HyperSonic SQL and SQLMemTable, challenge Chen's claim that middleware approaches lack an internal DBMS for efficiency. With an internal, middleware database system, efficient global query processing and joining can be performed (Ceccet, 2003). Increased speed in

hardware and streamlined software also contribute to more efficient middleware solutions. Due to these improvements, Chen's claim for performance concerns in middleware approaches needs to be reevaluated.

In 2003, Emmanuel Cecchet and his research team challenged Chen's claim by heading an effort to analyze the performance of multiple levels of the RAIDb cluster. This analysis highlighted the simple load balancing and throughput enhancements available from a RAIDb cluster in a web-based transaction environment. The multiple levels of the RAIDb standard were tested using the TPC-W (Transaction Performance Council Web Transaction) benchmark against the C-JDBC software. The information published by Cecchet shows an increase in performance directly proportional to an increase in the number of nodes and the amount of redundancy employed in the cluster. This is an important step to show performance enhancements available with the C-JDBC system (Cecchet, 2003).

The report published by Cecchet, however, fails to include some important information about the tests and performance of the C-JDBC cluster. Details on the methods of randomization for each run were not provided. Without this information, verifying the unbiased nature of the test results becomes very difficult. The report also fails to provide information on the relative efficiency of RAIDb cluster configurations. This information is vital for setting up clusters that do not perfectly coincide with the configuration represented in the performance analysis. Cecchet's report additionally fails to include several important statistics in evaluating the results, such as confidence and significance levels. Without these levels, the extent to which these results can be trusted is not clear. Regression and statistical analysis is also lacking from Cecchet's report.

Without regression analysis, a model used to predict the performance of a cluster is unavailable. Furthermore, without statistical testing, the variables that contribute to the prediction model cannot be identified. Finally, Cecchet's research was largely performed as a proof of concept trial. During Cecchet's tests, the C-JDBC software was not yet publicly released and has been greatly changed since that time. The first release of the C-JDBC software contains many changes and enhancements that may alter the performance of the system.

Based on Cecchet's research, many companies, schools, and consumers began implementing RAIDb as a database solution (Depue, 2004). However, because Cecchet's performance report was incomplete, these organizations also have incomplete information regarding the performance of RAIDb. Without a sufficient performance model, users of RAIDb databases cannot predict cluster performance. Without a sufficient performance model, the full capabilities and effects of a RAIDb cluster are not clear.

1.3 Hypothesis

Cecchet's performance analysis attempted to predict the relative performance of a standard C-JDBC cluster (Cecchet, 2003). However, with the data provided, there was insufficient information to achieve this goal. Other research involving C-JDBC clusters showed improved performance (Depue, 2004), but failed to describe those results in a way that is statistically and scientifically decisive. A more conclusive prediction of C-JDBC cluster performance might have been produced if a statistical performance model for these clusters was established.

The following approach was used in this thesis to create a statistical performance model for RAIDb systems:

- Benchmarking experiments were executed by varying several key factors to determine optimal configurations and analyze the effect of each of the factors on the cluster's performance.
- The results of these benchmarks were used to create a mathematical regression which will be useful to describe and predict the performance of these clusters.
- The effects of all key factors were statistically evaluated to determine which significantly affect performance. This was accomplished by analyzing the effects of these key factors for differing configurations for each RAIDb level.

In addition to cluster performance levels, several vital tests were executed to determine the viability of RAIDb clusters as an effective database solution. These tests were used to evaluate several characteristics that are essential to data integrity, including Atomicity, Consistency, Isolation, and Durability (ACID).

By measuring and assessing C-JDBC cluster performance characteristics using ANOVA techniques, a statistically sound performance model was produced. By testing ACID-compliance, the viability of the C-JDBC clusters as a database solution was also evaluated.

1.4 Justification

The need for high performance and high capacity database systems, at low cost, is often the main reason for building a database cluster. When the performance of a single high-end database server is not sufficient for a given workload, a cluster of two or more

machines is built to alleviate stress on the server. This effectively decreases the work each machine is responsible for and provides more processing power to complete the task. When database clusters are necessary, it is imperative to understand the implications of using this type of system.

C-JDBC is an affordable, heterogeneous clustering solution. Even without an established performance model, many companies are beginning to use C-JDBC as a clustering solution (Depue, 2004). Many of these companies struggle to find optimal cluster configurations and to determine the factors that most strongly influence cluster performance.

This thesis provides a model with which the performance of a C-JDBC cluster can be predicted. Using this model, the implications and outcomes of a C-JDBC cluster can be anticipated. This model also describes the effect that several factors have on C-JDBC clusters.

1.5 Assumptions

This research assumes that the relative change in cluster performance is proportionally equal across all CPU speeds (Cecchet, 2003 and Rogers, 2004). This assumption allows a performance model to be built and cluster configurations to be evaluated despite differences in processor speed.

1.6 Delimitations

In order to obtain an accurate performance model for C-JDBC clusters, it was necessary to limit the database nodes to established hardware and software

configurations. This research was strictly limited to the TPC-H database benchmarking suite. The cluster's network speed was held at 100 Mbps. Each node contained a 350 MHz Pentium II processor and a 20 GB hard drive. All hardware architecture was identical throughout the cluster. Clusters from two to six nodes were used for this research.

Chapter 2

2 REVIEW OF LITERATURE

2.1 Databases

Database Management Systems (DBMS) are now an essential component of the business marketplace. The role of the database is to provide a mechanism for easy, electronic storage and retrieval of information. While methods of storing and retrieving data have existed for many years, the first successful electronic database was SABRE, produced by International Business Machines (IBM) and American Airlines in the early 1960s (Brewer, 2004). The SABRE system resided on a single mainframe machine in Briarcliff Manor, N.Y from 1960 to 1972 when it was moved to Tulsa, Oklahoma. The SABRE system was built to automate airline reservations by electronically filing requests and tickets. The SABRE system, however, was not engineered as a relational database, requiring users to understand the structure and construction of the database model in great detail to use it (Vaughn, 2003).

2.1.1 Relational Databases

In June 1970 Dr. Edgar Frank Codd, a researcher for IBM, published an article that describes an area of study called relational mathematics. Codd outlined a way to use relational calculus and algebra to allow for the storage and retrieval of large amounts of

information. Codd envisioned a system where the user would be able to access information with English-like commands, and where information would be stored in tables (Codd, 1970). Codd's article was applied to a research project started by IBM known as "System-R." The purpose of the System-R project was to create a database system which relied on the relational mathematics and set theory that Codd proposed. As a part of this system, information would be stored and retrieved through the use of Codd's relational algebra. The System-R project also created a language of English-like commands to query, modify, and retrieve the information in the database system. System-R called this the Structured Query Language (SQL). A standardized version of SQL, based on the System-R project, is the query language most often used in databases today. The System-R database system also provided a standard for Relational Database Management Systems (RDBMS) and their design. The System-R project eventually evolved into the IBM's SQL/DS and DB2, Oracle, HP's Allbase, Tandem's Non-Stop SQL and others (Brewer, 2004).

Although IBM was the primary inventor of relational database systems, it was not the only product to market. During the time that System-R was being developed, another database project, called Ingres, was being completed. While based on relational mathematics, Ingres used a language called QUEL, rather than SQL, to retrieve information from the relational tables. The development of the Ingres database system also led to other database projects including Sybase, MS SQL Server, Britton-Lee, PACE, and others (Brewer, 2004).

The first widely-available relational database system was marketed by Oracle with the release of Oracle 2 in the early 1980s. Oracle 2 was also the first publicly-

available database system developed on the SQL standard (Ault, 2003). Soon after, IBM's first release of SQL/DS was also made available. Throughout the 1980s and 1990s, database standards and packages were refined. Through time, more providers of database services began their work to create a superior product. These efforts resulted in many database systems, including mySQL, PostgreSQL, Microsoft Access, Microsoft SQL Server, Oracle, and many more. Currently, most database models available are relational in nature (Vaughn, 2003).

2.1.2 Post-Relational Databases

First developed in the late 1980s, post-relational databases are built on object-based architectures that started gaining popularity around this time. Object-oriented database systems move away from the relationship table storage model and store all information as objects (Matisse, 2003). While most of these post-relational databases can still be accessed using a modified version of the SQL standard, objects are mapped to relational data sets in order to use relational math to retrieve information (Shen, 1995).

2.1.3 Database System Configuration Types

Due to the many performance needs in modern database systems, many distinct classifications of system configurations exist. Databases may be classified based on redundancy, high availability, speed, or throughput. Many relational database systems, such as Oracle, Microsoft SQL Server, and IBM DB2, employ varied means, from distributing database tables among computing nodes to database redundancy, in order to provide high performance and availability. Four general types of database system

configurations include single-node databases, hot stand-by databases, database federations, and distributed databases.

2.1.3.1 Single-node Databases

Almost all modern single-node database systems, from IBM's SQL/DS to mySQL, require only one computer node, or machine, to operate. An entire database structure is housed and confined on a single machine. This allows for easy administration, yet can become a limiting factor in high volume environments. The only method available to improve performance on a single-node database system is to upgrade system hardware. Single-node database systems provide the necessary foundation for all other system types as other database systems extend the capabilities provided by the single-node system (Ault, 2003).

2.1.3.2 Hot Stand-by Databases

Often the need for high availability drives the need for a database platform. One method to satisfy this requisite is called a hot stand-by database cluster. Hot stand-by solutions usually involve two machines that are linked with a heartbeat cable (Figure 2.1). The heartbeat cable is used to transmit information about the state of each machine (Austin, 1999). During operation, only one of the two machines is active and hosting the database. If this primary node is damaged or becomes unavailable for any reason, the heartbeat signal from the primary node is no longer transmitted. The missing heartbeat signal communicates the failure of the primary node to the secondary node. The master

node then deactivates and the secondary node activates, providing a fail-over solution and providing high availability (Chen, 2003).

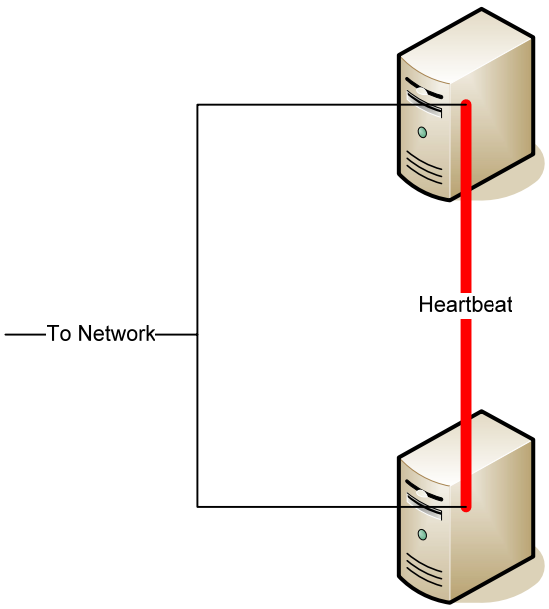

Figure 2.1. Hot stand-by database cluster with heartbeat cable.

2.1.3.3 Database Federations

Database federations attempt to unite many database platforms at one central point of access. Unlike hot-standby solutions, all nodes that are members of a database federation are active. One machine is used as a control server that dispatches requests to the correct data source (Hemibigner, 1985). Also, database system implementations are not required to be homogenous. The information on federation nodes is generally not replicated, which creates a large distributed information store. The goal of most database federations is to make a uniform system from many distinct and unrelated database machines (Hsiao, 1992). A simple federated database topology can be found in Figure 2.2.

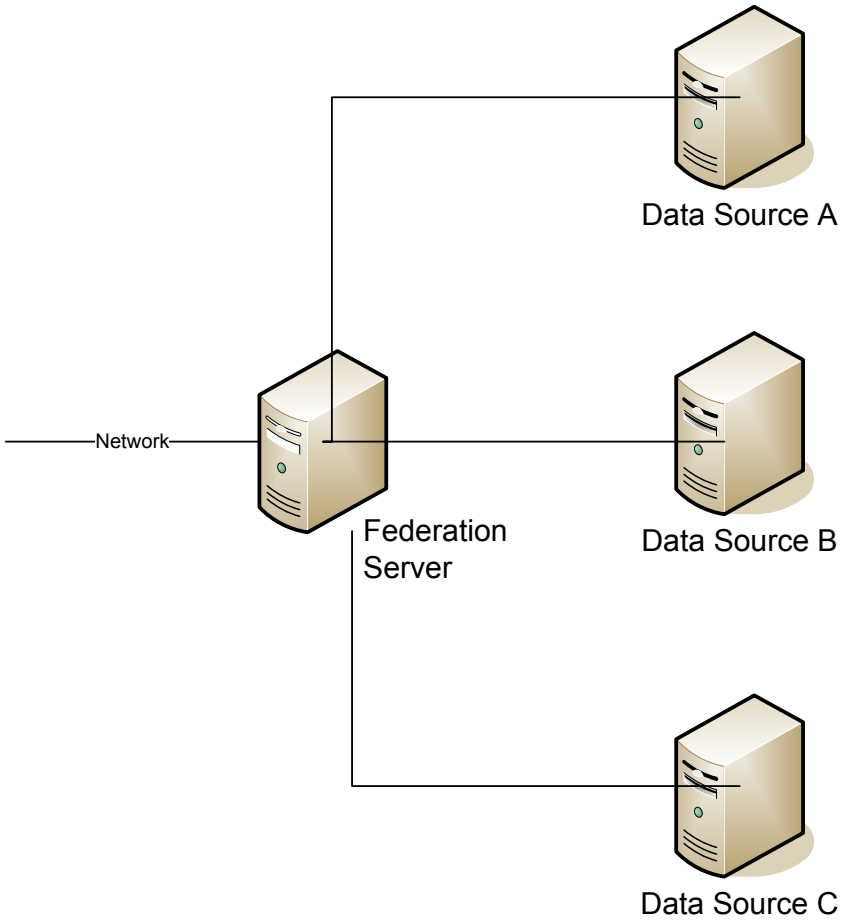

Figure 2.2. Federated database cluster with three heterogeneous data sources.

2.1.3.4 Distributed Databases

Much like database federations, distributed databases partition one or more large databases across several backend nodes. A master node, also called a transaction manager, usually coordinates requests and transactions to the backend nodes (Ault, 2003). Distributed databases can be non-replicated (Figure 2.3) or replicated (Figure 2.4) and are usually homogeneous. Non-replicated distributed databases do not duplicate information across backends, while replicated distributed databases store multiple copies of the data. Replication can provide many performance improvements derived from load balancing, redundancy, and high availability (Austin, 1999).

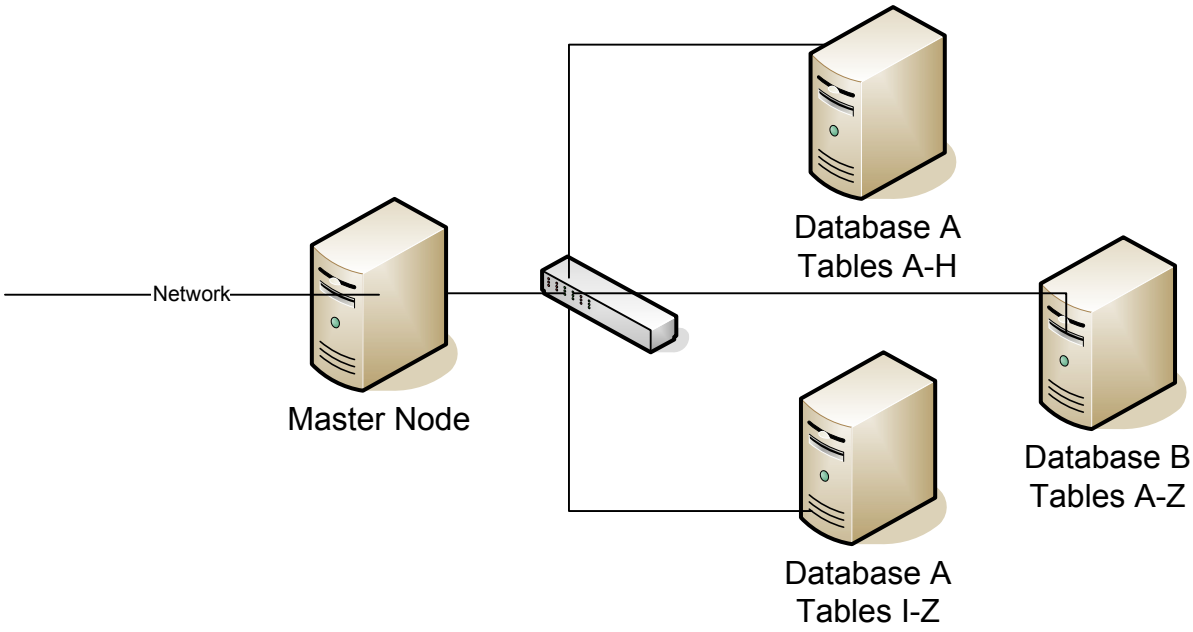

Figure 2.3. Homogenous distributed database cluster without data redundancy

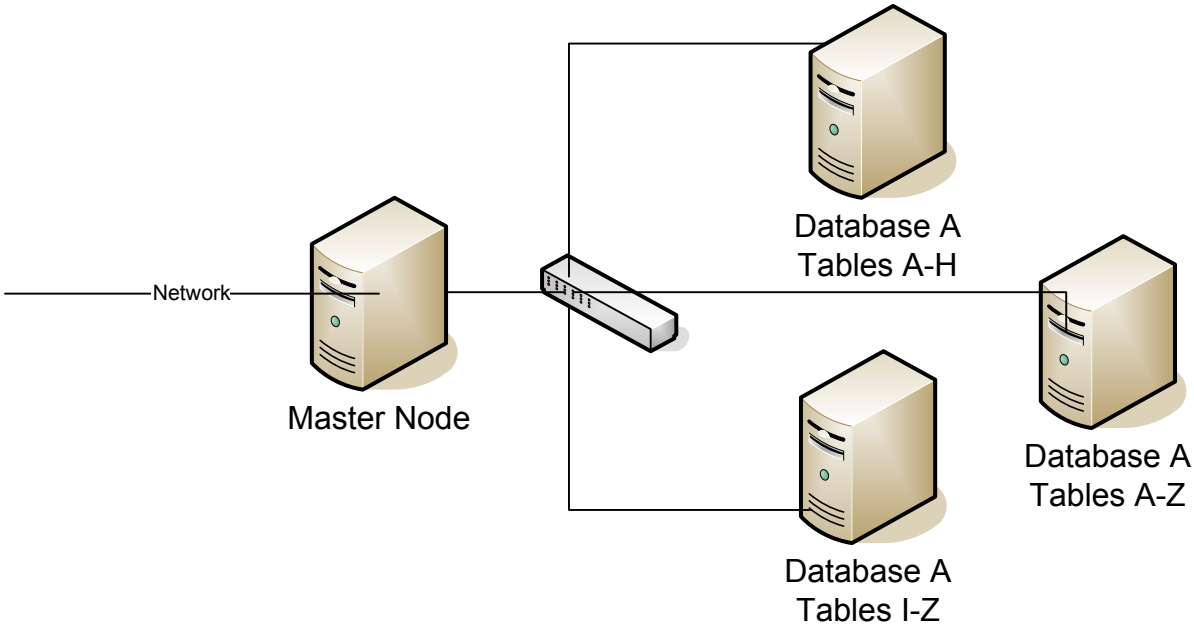

Figure 2.4. Homogenous distributed database cluster with data redundancy

While the concept of distributed databases is simple, implementation has long been a struggle. Oracle Corporation has led the industry in the development of database distributions. In the early 1990s, Oracle introduced its first database system founded on

the principles of distributed databases. In 1991, Oracle Corporation released Parallel Server for Oracle 6. By 1997, this evolved into support for Very Large Databases (VLDB). By 2000, full distributed and federated database capabilities were available as Real Application Clusters, or Oracle RAC (Ault, 2003).

2.1.4 Distributed Systems

Distributed databases and database federations can be both homogenous and heterogeneous distributed computing systems (Figure 2.3 and Figure 2.4). A distributed system has been defined by Nielsen (2003) as, "a collection of independent computers that appears to its users as a single coherent system." There are many important requirements that must be considered when implementing a distributed system that satisfies this definition. The details of how these requirements are satisfied create the diversity observable in available distributed systems.

2.1.4.1 Transparency

Hiding the particulars of system component operation and allowing the user to interact with the distributed system uniformly is an important characteristic of distribution schemes. A common term used to describe these characteristics is transparency. There are many important types of transparency to understand when dealing with a distributed system. Table 2.1 illustrates the several types of transparency attributes and their meanings. A well-designed distributed system will employ a mechanism to handle each type of transparency (Tanenbaum, 2002).

Table 2.1. Distributed system transparency types and their descriptions.

Transparency	Description
Access	Hide differences in data representation and how a resource is accessed.
Location	Hide where a resource is located.
Migration	Hide that a resource may move to another location.
Relocation	Hide that a resource may be moved to another location while in use.
Replication	Hide that a resource is replicated.
Concurrency	Hide that a resource may be shared by several competing users.
Failure	Hide the failure and recovery resource.
Persistence	Hide whether a (software) resource is in memory or on disk.

2.1.4.2 Atomicity, Consistency, Isolation, and Durability

Another important requirement for distributed systems is to maintain data atomicity, consistency, isolation, and durability (ACID). ACID-compatibility is designed to ensure that data quality is preserved throughout all modifications, or transactions. Atomicity of information ensures that any alteration or usage of the data is uninterrupted and complete. If any atomic part of the transaction does fail, "a mechanism is provided that ensures the return of the system to its state prior to initiation of the transaction." (Microsoft Corporation, 2004) Atomicity implies that there is never an incomplete or interrupted change made to the information or data structure of a distributed system (Transaction Processing Performance Council, 2003).

Consistency of data is designed to ensure that information available within the distributed system remains consistent. This implies that data must move from one stable state to another while maintaining consistent state throughout the transaction. Lack of consistency may be caused by failure to update replicated nodes creating illegitimate copies of data. Communication or hardware failure also can result in modification to only a fraction of the system. As with atomicity, if any part of the transaction does not meet consistency requirements, the transaction is aborted and a mechanism is provided to

return the system to its previous state (Transaction Processing Performance Council, 2004).

Isolation refers to the requirement of a distributed system to keep all alterations made in a transaction invisible to other operations, processes or threads. Failures in isolation can be described as dirty write, dirty read, non-repeatable read, and phantom read. A description of each of these failures can be found in Table 2.2.

Table 2.2. Four types of isolation failures with accompanying descriptions (Transaction Processing Performance Council, 2004).

Isolation Failure	**Description of Failure**
Dirty Write	1. A primary transaction begins; data is opened and modified. 2. A second transaction modifies or deletes data opened in (1). 3. The primary transaction reloads data; some of the data previously open in (1) is now deleted or modified. Modifications made in (1) now may be inaccurate.
Dirty Read	1. A primary transaction begins and modifies data. 2. A second transaction opens the same, now modified, data as in (1). 3. The primary transaction fails and is rolled back; information read in (2) is now incorrect or incomplete.
Non-repeatable Read	1. A primary transaction begins; data is opened for reading. 2. A second transaction modifies or deletes the same data in (1). 3. The primary transaction reloads data; some of the data previously opened in (1) is now deleted and modified.
Phantom Read	1. A primary transaction opens data from reading. 2. A secondary transaction creates additional information that should have been opened in (1). 3. The primary transaction reloads the data in (1); the data set is different than in (1).

Each of these violations may cause the system to fail to isolate transactions. Each of these violations can be mitigated by hiding changes in data to the full system until the transaction is complete. If system data does not meet isolation requirements, or if data becomes inconsistent due to the isolation requirement, the distributed system is reverted back to a previous system state.

Phase locking is a two- or three-step process that is designed to guarantee isolation in transactions and handles transaction state should the constraints of isolation be violated. The first step in phase locking is for preparation. If any problem occurs in the preparation stage, the transaction is rolled back. Problems in the preparation stage could consist of attempts to lock records which are already locked, hardware or software failure, or any problem found with any of the ACID properties. The second step in phase locking is called the commit. If a third step is included in the phase-locking, a final check is made to verify the correct commit of the transaction (Sadoski, 2004). A complete description of each step in the phase locking process can be seen in Table 2.3.

Table 2.3. Three steps in phase locking and corresponding descriptions.

Step in Phase Locking	Description
Preparation	1. All affected data elements are locked. 2. Data is changed within the scope of the transaction. 3. All preparations to make the change permanent are made to the system. 4. The success of the preparation is evaluated on all involved nodes; if any problem exists on any of the nodes, the transaction is aborted and rolled back.
Commit	1. All changes are made permanent. 2. Modifications become visible to all other elements in the system. 3. All locks on the data elements are released.
Post-Commit	1. Success of the commit phase is evaluated.

Durability, the fourth element of ACID, is "the ability to preserve the effects of committed transactions," or to preserve the distributed system's state (Transaction Processing Performance Council, 2004). Durability can be accomplished by recording to an inherently non-volatile medium such as magnetic disk, magnetic tape, or optical disk. A volatile medium that will move data to a non-volatile storage medium before complete power failure is also a method to comply with the durability requirement (Transaction Processing Performance Council, 2003).

2.1.4.3 A Middleware Approach to Distributed Systems

In recent years, middleware solutions have entered virtually every computing market, from distributed industrial and real-time systems (Takegaki, 1995) to computing clusters (Nielsen, 2003) to large scale data distribution and replication (Shiroshita, 1996). Due to this overwhelming adaptation of middleware solutions, much research has been performed on their plausibility.

Middleware systems provide many benefits over other solutions. One large benefit of distributed systems based on middleware is that they are intrinsically heterogeneous (Balkanski, 2003). Unlike systems that are not middleware, compatibility between differing platforms is an essential element of the middleware design. The characteristic of designed heterogeneity allows the system to be more widely distributed as more system platforms are available for use. With the advent of network-enabled household computing devices, such as personal digital assistants, multimedia players, tablet computers, and game console systems, computing nodes available for inclusion in distributed systems has increased dramatically. The demand to interlink each of these

devices has also risen as these diverse computing platforms become commonplace (Schilit, 2004). Middleware distributed systems is a natural solution to accomplish networking between heterogenous nodes.

Another benefit derived from middleware distributed systems is modularity. Using middleware systems, platforms can be replaced without the need to reprogram code or redesign the existing network cluster. Since all communication is made using the middleware protocol, new service providers or systems need only to implement the middleware API (Brown, 1998).

2.2 Redundant Arrays of Inexpensive Databases

As defined in September 2003 at INRIA (Institut National De Recherche En Informatique Et En Automatique) by Emmanuel Cecchet, Julie Marguerite, Willy Zwaenepoel, Redundant Arrays of Inexpensive Databases, or RAIDb, is a middleware database system which provides federated and distributed solutions. The standard for RAIDb is derived from two well-established disciplines: Redundant Arrays of Inexpensive Disks (RAID) and computing clusters. Referring to RAIDb's resemblance to RAID arrays, Cecchet (2003) states:

> RAIDb is to databases what RAID is to disks. RAIDb aims at providing better performance and fault tolerance than a single database, at low cost, by combining multiple database instances into an array of databases. Like RAID, we define different RAIDb levels that provide various cost/performance/fault tolerance tradeoffs. RAIDb-0 features full partitioning, RAIDb-1 offers full replication and RAIDb-2 introduces an intermediate solution called partial replication, in which the user can define the degree of replication of each database table.

As with computing clusters, RAIDb also attempts to harness the benefit of distributing workload. Load balancing and caching techniques increase performance capability of the system.

2.2.1 RAIDb-0: Full Partitioning

RAIDb, level 0 (RAIDb-0), is best described as database striping, or distributing the tables in the database among several backend nodes (Figure 2.5). RAIDb-0 is similar to common distributed database systems, without replication, such as Oracle RAC, PostgreSQL Replication Project, and Emic Application Clusters for MySQL. Data stored in a RAIDb-0 system is completely distributed among nodes. No replication or duplication of information is performed. The RAIDb-0 standard also requires the ability to perform distributed joins across multiple nodes.

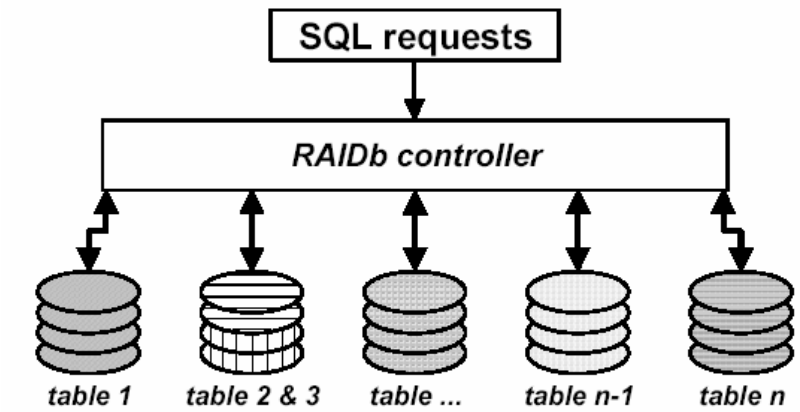

Figure 2.5. RAIDb-0 database cluster with full partitioning of n tables (Cecchet, 2003).

2.2.2 RAIDb-1: Full Replication

RAIDb, level 1 (RAIDb-1), performs data replication through database mirroring, or full replication (Figure 2.6). RAIDb-1 requires that each backend node have the

ability to handle the entire database. Using RAIDb-1, load balancing is also available as requests can be distributed over backend nodes.

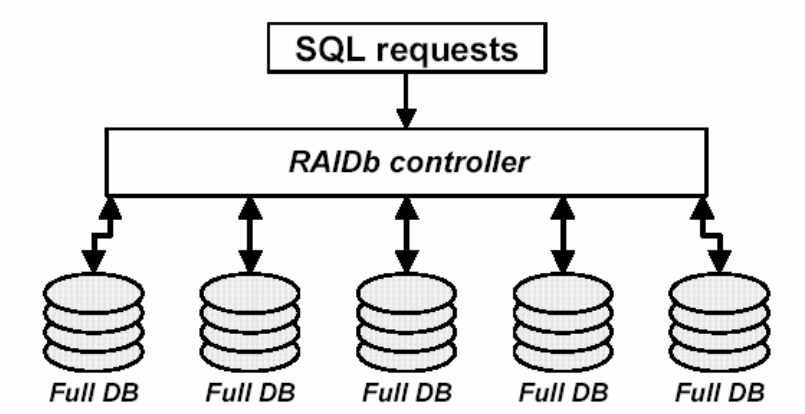

Figure 2.6. RAIDb-1 database cluster with full replication across five nodes (Cecchet, 2003).

2.2.3 RAIDb-2: Striped Redundancy

RAIDb, level 2 (RAIDb-2), features a combination of RAIDb-0 and RAIDb-1, or partial replication (Figure 2.7). Tables are distributed, like RAIDb-0, allowing for databases that are larger than the storage and performance capacity of a single node. However, using RAIDb-2, tables are also replicated, providing fault tolerance and load balanced clustering. The RAIDb-2 solution allows for continuous operation despite failed nodes.

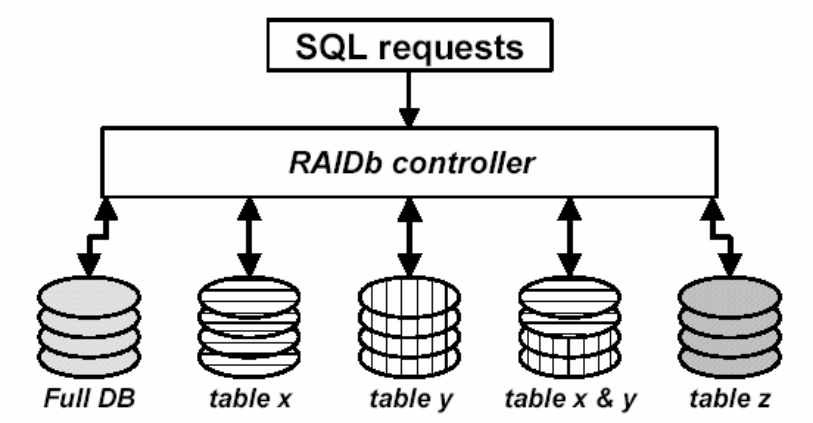

Figure 2.7. RAIDb-2 database cluster with redundancy and partial partitioning of tables across five nodes (Cecchet, 2003).

2.2.4 Additional RAIDb Levels

Two additional RAIDb levels are also detailed in the RAIDb standard, including error-checking levels for RAIDb levels 1 and 2. Error-checking RAIDb-1 (RAIDb-1ec) aims to discover and handle byzantine failures that can occur in a highly stressed cluster of PCs. Error-checked RAIDb-2 (RAIDb-2ec) aims to provide the same error checking as RAIDb-1ec to RAIDb-2 clusters.

2.2.5 RAIDb Controller

The RAIDb controller functions like the master node of a distributed database or the federation server of a federated database. The RAIDb controller has four main roles: message server, transaction coordinator, load balancer, and control server. As a message server, the RAIDb controller directs the messages used to control operations within the cluster. Messages from the controller are sent to database nodes to coordinate data requests and manage transactions. Inter-controller communication is used to coordinate multi-master operations. Other messages are also sent from the controller to signal node

regeneration and node failure. Additionally, all messages and requests from a database client are sent to the controller and forwarded to the intended nodes. As a transaction coordinator, the RAIDb controller manages transactions and ACID compliance. While the RAIDb controller allows most ACID assurances to be handled within each of the database nodes, deadlock, stale transactions and other hindrances to ACID compliance for the integral system is handled by the controller. The transaction management layer is also responsible for result set and query caching. As a load balancer, the RAIDb controller can use several algorithms to balance load among replicated nodes in RAIDb-0 or RAIDb-1 configurations. As a control server, the RAIDb controller offers node replication, dynamic cluster reconfiguration, and data backup (Cecchet, 2003).

2.2.6 Scalability

The RAIDb standard also provides a simple mechanism for horizontal and vertical scalability. Vertical scalability, as defined for RAIDb clusters by Cecchet (2003), is the ability to "compose [and combine] several RAIDb levels to build large-scale configurations." By nesting RAIDb controllers, vertical scalability can be easily obtained. RAIDb level hierarchy in the nested controllers is configurable. Configurations can be implemented as redundant RAIDb-0 clusters (Figure 2.8) or partitioned RAIDb-1 clusters (Figure 2.9) without significant change to the controller software or messaging libraries (Cecchet, 2004).

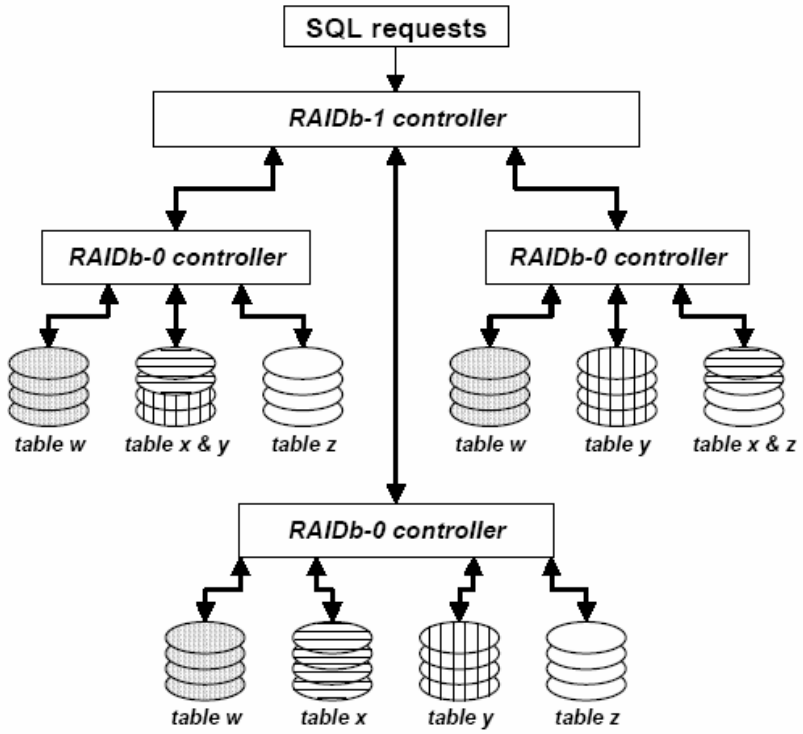

Figure 2.8. Redundant RAIDb-0 clusters showing vertical scalability (Cecchet, 2003).

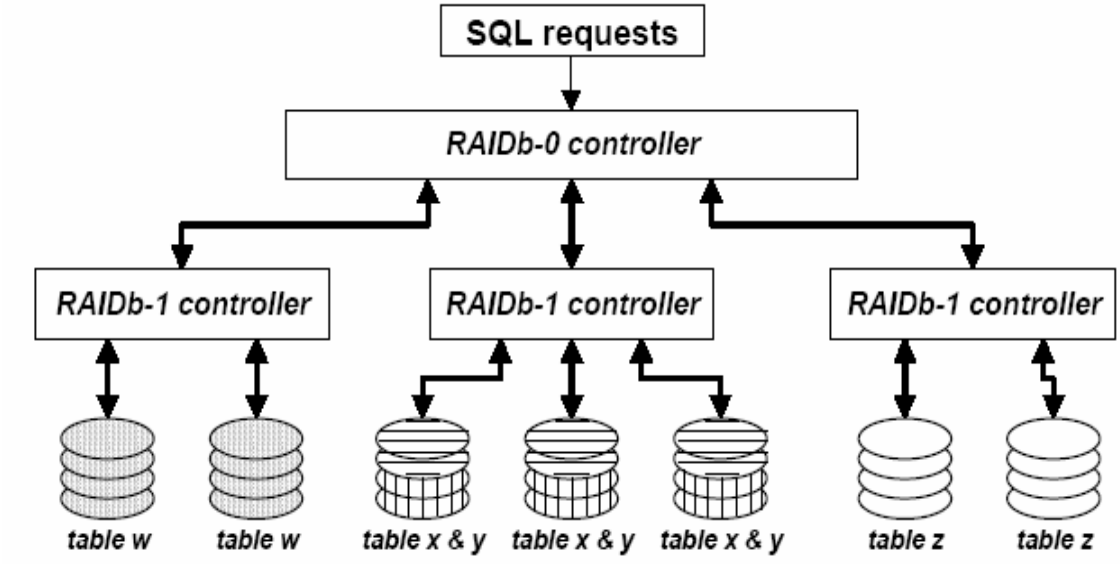

Figure 2.9. Partitioned RAIDb-1 clusters showing vertical scalability (Cecchet, 2003)

The RAIDb standard also provides for horizontal scalability. As defined by Cecchet (2003) for RAIDb cluster, horizontal scalability is the ability, "to have two or more controllers that synchronize the incoming requests to agree on a common serializable order." This is performed by creating a multi-master cluster. Multi-master RAIDb clusters allow for failover, client-side load-balancing, and controller redundancy. Low latency communication channels handle synchronization among multiple controllers in a multi-master RAIDb cluster (Figure 2.10).

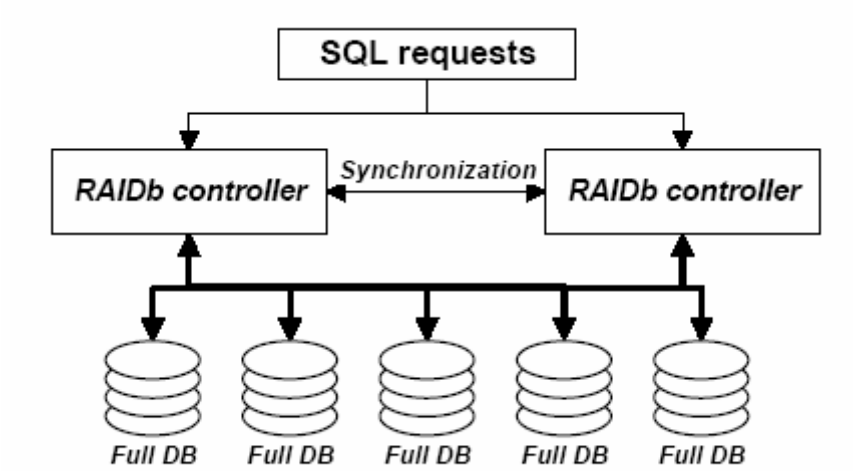

Figure 2.10. Multi-master RAIDb cluster showing horizontal scalability.

2.3 Clustered Java Database Connectivity

Clustered Java Database Connectivity, or C-JDBC, is a Java-based middle-ware implementation of RAIDb that is maintained and developed at INRIA. C-JDBC is heavily based on Java Database Connectivity (JDBC). JDBC drivers are used to connect and communicate with database servers from the C-JDBC controller. C-JDBC drivers are used to connect from a client to the C-JDBC controller. C-JDBC supports most common database functions, including stored procedures, SQL (Structured Query Language), and security rules.

While C-JDBC supports most of the features specified in the RAIDb specification, there are many RAIDb requirements that remain unsupported. These requirements include true multi-master configurations, multi-node joins, and complete communication libraries (Cecchet, 2004).

True multi-master configurations allow all controllers to manage, through communication and coordination, the processes of the database cluster. The failure of C-JDBC to meet this requirement is noticed most readily in horizontal configurations. While the RAIDb standard allows multiple masters to control and access all nodes in a cluster, C-JDBC partitions node control among masters. If a client requests information from a node over which the connected controller does not have authority, the request is passed to the authoritative controller. Only upon failure can a controller relinquish authority over a node to another controller.

During the setup of a C-JDBC cluster, database nodes must be assigned to a controller. To create a hot-standby cluster, every node needs to be assigned to a primary and a secondary controller. The secondary controller would then be deactivated. Upon failure of the primary controller, authority over the nodes would be relinquished to the secondary controller. C-JDBC, however, does not allow assigning nodes to multiple masters. The inability of C-JDBC clusters to effectively handle multi-master environments makes hot standby solutions impossible.

For RAIDb-0 and RAIDb-2 configurations, multi-node joins are essential. As tables are distributed across nodes, a method to perform queries on related tables, regardless of physical separation, is important. C-JDBC, however, does not provide this capability. To perform any query among multiple tables, all tables involved in the query

must be present on a single node. This reduces the usefulness of RAIDb-0 and RAIDb-2 clustering configurations.

C-JDBC currently uses JGroups, a Java-based communications library, as a messaging layer for controllers and database nodes. JGroups is not optimized for low-latency communication that is essential for controller and node synchronization. To enforce ACID compliance, the controller must allow for the inefficient communication library.

2.4 Additional Research on Middleware Database Clusters

RAIDb is not the only middleware database clustering solution available. Other solutions provide similar advantages to RAIDb, such as load balancing, fault tolerance, data partitioning and replication, and ACID-compliance. The two most prominent alternatives to RAIDb clusters are HA-JDBC and Emic Application Clusters for MySQL.

High-Availability JDBC, or HA-JDBC, "is a JDBC driver proxy that provides light-weight, transparent clustering capability to any underlying JDBC driver." Built on Java, the HA-JDBC tools take advantage of the portability and compatibility of middleware systems. While providing full support for JDBC 3.0, HA-JDBC offers similar functionality to RAIDb. HA-JDBC is largely incomplete, however, providing only rudimentary replication and clustering support (HA-JDBC, 2004).

Emic Networks has also created a middleware database clustering engine called Emic Application Clusters for MySQL (EAC for MySQL). EAC for MySQL provides a single virtual database from multiple nodes running the MySQL database platform. Like RAIDb, EAC for MySQL also provides "software-based load balancing and

comprehensive fault management." Unlike RAIDb, EAC for MySQL is limited to the MySQL platform. No other database platforms can be used with EAC (Emic Networks, 2005).

2.5 Benchmarks

Benchmarks are often used to measure various aspects of system performance. Many measurements make database benchmarks useful, including throughput performance, speed, cost, and ACID-compliance. Throughput performance can be described as the ability of the database system to effective process many concurrent query streams. This characteristic is often commonly referred to as a stress test. A throughput test is generally created by sending multiple tasks to the system under test (SUT) and measuring the performance as a result. The speed of a system describes how quickly a serial workload can be executed. The result is generally measured in execution time. Cost is also an important factor for all system evaluations. While it may be possible to have a system that responds favorably to speed and throughput tests, the cost to implement such a system may be prohibitive. ACID-compliance testing is another essential piece of the database puzzle. While a database may be cheap, quick, and able to handle heavy load, if ACID attributes are not met, retrieving data from the system becomes useless and the solution unusable.

The industry-accepted benchmarking tool used for profiling databases is a suite of tests provided by the Transaction Processing Performance Council (TPC). The TPC benchmarks each employ the four characteristic measurements of a good database benchmark. The TPC collection includes four benchmarks: TPC-C, TPC-H, TPC-R, and

TPC-W. Each of these tools has relevance to a distinct aspect of performance measurement.

2.5.1 TPC-C Benchmark

The first of these benchmarking suites, TPC-C, is an online transaction processing benchmark. TPC-C includes many features of performance analysis. Among these features are tests to determine access and transaction speed, total cost of ownership, checkpoint durations, and the number of lost connections during a given measurement interval. Each of these measurement statistics help to determine the viability of a database solution.

The TPC-C benchmark consists of nine tables. Using these tables, the TPC-C benchmark attempts to profile a database system's ability to handle online transaction processing (OLTP). The TPC has continued to evolve its benchmarks to remain as representative of current practice as possible. The TPC-C benchmark continues to be a popular yardstick for comparing OLTP performance on various hardware and software configurations. The TPC-C benchmark is also much more involved than any of the other TPC benchmarks and comes without tools other than a detailed specification of the process. Using the TPC benchmark suite, a complete profile of any database system can be made which is accepted and understandable by the industry.

2.5.2 TPC-H Benchmark

The TPC-H benchmark is a decision support benchmark that focuses on three distinct types of tests: the load test, the power test, and the throughput test. The load test

is intended to measure how quickly and efficiently the database system can load raw data and prepare it for use. The method of loading data is not specified and the time measured includes table and index creation. The power test is used to determine the speed of the database in doing serial operations. A single query stream is sent to the database system and the time to complete the single query stream is measured. Finally, the throughput test is intended to determine the ability of the database in handling parallel query streams. Multiple query streams are sent to the database system and the time to execute all the query streams is measured.

The performance metric reported by TPC-H is called the composite query-per-hour performance metric (QphH@Size). This performance metric reflects multiple aspects of the capability of the system to process queries, including database size, "the query processing power when queries are submitted by a single stream and the query throughput when queries are submitted by multiple concurrent users." The TPC-H Price/Performance metric is expressed as cost per composite query-per-hour at a specified database size ($/QphH@Size).

2.5.3 TPC-R Benchmark

The TPC-R benchmark is a decision support benchmark similar to TPC-H, but which allows for query optimization based on an advanced knowledge of the queries. It consists of a suite of business oriented queries and concurrent data modifications.

The performance metric reported by TPC-R is called the composite query-per-hour performance metric (QphR@Size). This performance metric reflects multiple aspects of the capability of the system to process queries, including selected database

size, single stream tests aimed to quantify system power and the query throughput tests made of multiple concurrent query streams. The TPC-R Price/Performance metric is expressed as cost per composite query-per-hour at a specified database size ($/QphR@Size).

2.5.4 TPC-W Benchmark

The TPC-W benchmark is a transactional web benchmark that simulates an online bookstore. “The workload is performed in a controlled internet commerce environment that simulates the activities of a business oriented transactional web server.” The benchmark workload measures many factors generally associated with online database systems, such as:

- Multiple browser sessions
- Database-driven page generation (dynamic web pages)
- Online transactions
- Several database access profiles, including shopping, browsing, and ordering mixes.
- Databases consisting of many tables with a wide variety of sizes, attributes, and relationships
- Transaction integrity (ACID properties)
- Contention on data access and update

The performance metrics reported by TPC-W is the number of web interactions processed per second (WIPS) and associated price per WIPS.

2.6 Review of Literature Conclusions

The demand for distributed data networks is growing at a phenomenal rate. Many research projects and prototypes have been launched to fill the rising demand for these computing clusters. Among these prototypes are middleware database clustering solutions that are designed to take advantage of the increased processing power of parallel computing environments.

Before implementing any of these database clusters in a commercial environment, system performance and characteristics must be evaluated. The cluster's abilities to remain ACID compliant and provide transparency are important characteristics to model. Other important factors that contribute to a proper system model include reaction to system strain, throughput enhancements available from load balancing and redundancy, and resiliency to node failure.

An accurate measure of these characteristics can be obtained through widely accepted benchmarks. The industry standard for database benchmarking is the TPC suite of benchmarks. These four benchmarks provide tools to measure throughput performance, speed, ACID compliance, and cost. An accurate depiction of the performance of a database solution can be inferred through the analysis of at least one of the TPC benchmarking tools.

Chapter 3

3 RESEARCH PROCEDURES

This chapter explains the design of the statistical experiment and the physical and logical cluster configurations. The experimental design describes an ANOVA 2-level factorial experiment. The requirements for the ANOVA two-level factorial experiment are also explained. The Minitab 14 statistical software was used to generate the statistical designs. The physical cluster configuration, hardware requirements, and network configurations necessary to complete the statistical experiment design are also explained. The logical cluster configuration of the RAIDb controller and backend nodes are described. The logical plan consists of benchmark implementation specifics, software network configurations, and operating system configurations that allow the statistical tests to be completed.

3.1 Statistical Experiment Design

The statistical Analysis of Variance (ANOVA) test is a powerful tool for multivariate testing and regression modeling. Most modern regression modeling is accomplished through ANOVA or specialized tests based on the ANOVA procedures. A common ANOVA design includes two important pieces: blocks and treatments. Blocks consist of designed variations in the experimental process that attempt to monitor the effects of uncontrollable, or hidden, variables in the test. Some common examples of

blocks include location, time of the day, and experiment order. Varying the location, for example, affects the outcome due to uncontrollable environmental conditions. Treatments consist of designed variations in the experimental process that attempt to monitor the effects of controllable variables in the test. These treatments are the independent variables that the experiment was designed to investigate. Effects caused by treatments are measured and used to build regression models of important data (Keller, 2001). Table 3.1 shows an ANOVA experiment design with blocks and treatments.

Table 3.1. ANOVA experiment design with three blocks and three treatments.

	Treatments		
Blocks	*Chemical A*	*Chemical B*	*Chemical C*
Location 1	Outcome A1	Outcome B1	Outcome C1
Location 2	Outcome A2	Outcome B2	Outcome C2
Location 3	Outcome A3	Outcome B3	Outcome C3

Several important statistics are generated to evaluate the outcome of an ANOVA test. One of the most important statistics generated in an ANOVA test is the F-statistic. The F-statistic describes where the most significant differences lie. If significant variation exists among the blocks, or locations, in Table 3.1, the F-statistic for rows will be uncharacteristically high. If significant variation exists between the chemicals in Table 3.1, the F-statistic for all columns will be unusually high. In other words, the probability that the observed differences between treatments, or blocks, occurred entirely by chance is relatively low (Keller, 2001).

The probability of the F-statistic occurring by random variation is often called the p-value of the statistic. This p-value is compared against a chosen level of significance, denoted by the Greek letter α. The value of α is any number between 0 and 1 that is

chosen to represent the significance of the results of the test. For example, to be sure that the results of an experiment would be accurate 95% of the time, α would be set equal to 0.05. To be sure that the results of an experiment would be accurate 99% of the time, α would be equal to 0.01. Only a p-value that is less than α would indicate a significant event (Montgomery, 2004).

Closely associated with the significance level of an experiment is the confidence level. Confidence is equal to one minus the chosen α. By choosing an α-level of 0.05, we have a confidence level of 0.95, or 95%. For the purpose of this research, our significance level is chosen to be 0.01 ($\alpha = 0.01$) and, by definition, a confidence level equal to 0.99 (Keller, 2001).

When the results of an ANOVA test indicate that significant variation exists between treatments, a regression model can be created. A regression model is a mathematical equation that is used to model, or to predict, the outcome of the system under test. A regression model is created by using the treatments of an experiment to create an equation of the form $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n + \varepsilon$, where β_0 symbolizes the y-intercept of the equation, x_i indicates a variable included in the treatments, β_i indicates the coefficient of the related treatment variable, and ε symbolizes the error inherent in modeling estimations. Any terms for insignificant treatments are removed from the equation for simplification (Keller, 2001).

When a regression model is created it should be tested for goodness-of-fit. This is accomplished through the coefficient of determination, denoted by the symbol R^2. The coefficient of determination measures the amount of variation explainable by the significant treatment variables. If the experiment design in Table 3.1 were to have a

coefficient of determination equal to 0.6501 ($R^2 = 0.6501$), the statistic would explain that 65.01% of the variation or differences found in the dependent variable are explained by the differences in the independent variables. In short, the regression model can explain roughly 65% of the differences occurring in the overall system. The remaining 35% of variation in the above model would remain unexplained (Montgomery, 2004).

Although a regression model may be created with an appealing coefficient of determination, the model may be greatly flawed without testing the coefficient of correlation, denoted by ρ. The coefficient of correlation can be used to determine if a linear relationship between independent variables actually exists. This is important to determine the validity of the regression model and the goodness-of-fit tests. The coefficient of correlation is a number such that $-1 <= \rho <= 1$. The coefficient of correlation is measured against the Student's t-statistic to conclude the significance of the correlation between the treatments (Keller, 2001).

For ANOVA experiments to be valid, four requirements must be met. These requirements are:

1) Homoscedasticity – The equality of variance throughout the residual distribution which represents the data to be studied.

2) Normality – The residuals obtained from the experiment must be normally distributed.

3) Independence – The residuals from the experiment should not show trends, but should be randomly scattered on the plot.

4) Outliers – Outliers should be accounted for as data is analyzed.

Each of these requirements can be verified by testing the residuals from the experiment. The residuals of the data in ANOVA analysis are "the difference[s] between the observed value of y and the value predicted by a model, sometimes referred to as the error of prediction." A plot of the residuals against the predicted values of the model will help to determine that the distribution has uniform variation. A histogram of the residuals can determine the normality of the data. By graphing the residuals over time or experiment order and looking for trends, independence can be verified. A scatter plot of the residuals can help to locate outliers (Keller, 2001).

A specialized version of the ANOVA experiment design is the 2-level factorial experiment design. The two-level factorial design experiment is generally used when trying to determine the significant effect of multiple factors in a system. The 2-level factorial experiment systematically varies multiple factors or treatments between high (+) and a low (-) operational levels. The effects of each of the factors on the overall outcome are calculated and analyzed. Additionally, effects of the interactions between multiple factors are also evaluated. These effects can be tested against standard test statistics to determine their significance. These statistics can then be used to create a regression model that can be used to predict an outcome (Montgomery, 2004).

As with most experiments, a measure of success must be set to document the accomplishment of the experiment. This measure of success, or hypothesis, can be stated in two ways. The first is most commonly referred to the null hypothesis. The null hypothesis for this research states that variation does not exist between RAIDb cluster configurations and with that variation an effective regression model can not be generated. The second statement of the hypothesis is closely related to the first. This statement is

commonly called the alternative hypothesis. The alternative hypothesis for this research states that variation does exist between RAIDb cluster configurations and with that variation an effective regression model can be generated (Montgomery, 2004).

To test these hypotheses and create a valid statistical experiment design, three factors, that are judged most influential, must be chosen for each RAIDb level. As C-JDBC cannot perform multi-node data joining, benchmarks which define relations between multiple tables cannot be used (Cecchet, 2004). All TPC benchmarks define relations for each table included in the schema (Transaction Processing Performance Council, 2003). Due to the incongruence of the C-JDBC implementation with the RAIDb specification, RAIDb-0 clusters cannot be tested.

Other than RAIDb-0, three factors for each additional RAIDb level were chosen as significant contributors to the performance of the cluster. For RAIDb-1, these three factors include:

1) Cluster Size (x_1) – The number of dedicated database nodes in the cluster. Cluster size will be varied between a low level of 2 machines and a high level of 6 machines.

2) TPC-H Thread Count (x_2) – The number of throughput threads executed simultaneously during the TPC-H throughput test. This is equivalent to system load. Thread count will by varied between a low level of 1 thread and a high level of 12 threads.

3) Number of Controllers (x_3) – The number of active RAIDb controllers included in the system. Controller number will be varied from a low level of one controller

to a high level of two controllers. When multiple controllers are employed in the cluster, controllers are horizontally scaled.

Using these factors and the levels associated with them, a two-level factorial experiment design can be generated. A random number table was used to randomize run order to eliminate bias. Table 3.2 shows the two-level, two replicate factorial experiment design for a RAIDb-1 system.

Table 3.2. Three factor, two-level factorial experiment design for RAIDb-1 clusters.

Sequence	Run Order	Replicate	Cluster State		
			Controllers	Threads	Nodes
1	11	1	+	+	+
2	6	1	+	+	-
3	15	1	+	-	+
4	8	1	+	-	-
5	14	1	-	+	+
6	13	1	-	+	-
7	12	1	-	-	+
8	9	1	-	-	-
9	4	2	+	+	+
10	16	2	+	+	-
11	3	2	+	-	+
12	1	2	+	-	-
13	10	2	-	+	+
14	5	2	-	+	-
15	2	2	-	-	+
16	7	2	-	-	-

An experiment design was also created for RAIDb-2. The three significant factors in RAIDb-2 performance were judged to be:

1) Cluster Ability – Cluster ability describes the power of each node to execute the entire query set. Cluster ability is a simple average of the number of tables available per node. Cluster ability will be varied from a low level of 2 to a high level of 7. A high level of 7 was chosen because a level of 8 would have

transformed the cluster to RAIDb-1. The tables used on each node during high and low cycles can be found in Figure 3.1.

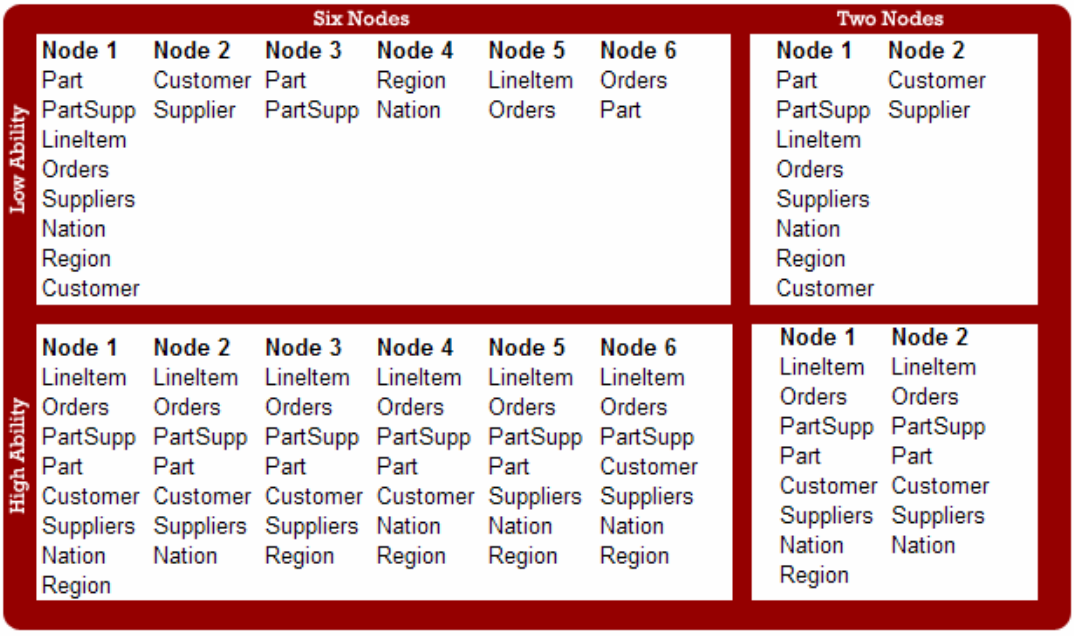

Figure 3.1. Table allocations for each node for high and low values of cluster ability and size.

2) Cluster Size – The number of dedicated database nodes in the cluster. Cluster size will be varied between a low level of 2 machines and a high level of 6 machines.

3) TPC-H Thread Count – The number of throughput threads executed simultaneously during the TPC-H throughput test. This is equivalent to system load. Thread count will by varied between a low level of 1 threads and a high level of 12 threads.

Using these factors and the levels associated with them, a two-level factorial experiment design can be generated. A random number table was used to randomize run order to eliminate bias. Table 3.3 shows the two-level, two run factorial experiment design for a RAIDb-2 system.

Table 3.3. Three factor, two-level factorial experiment design for RAIDb-2 clusters

Sequence	Run Order	Replicate	Cluster State		
			Ability	Threads	Nodes
1	12	1	+	+	+
2	9	1	+	+	-
3	4	1	+	-	+
4	15	1	+	-	-
5	13	1	-	+	+
6	16	1	-	+	-
7	7	1	-	-	+
8	5	1	-	-	-
9	11	2	+	+	+
10	8	2	+	+	-
11	2	2	+	-	+
12	10	2	+	-	-
13	3	2	-	+	+
14	14	2	-	+	-
15	6	2	-	-	+
16	1	2	-	-	-

The results for these experiments are calculated and compared to appropriate statistics for the ANOVA test. A regression model is then generated and tested for goodness-of-fit and applicability to the collected data. All statistics are calculated using the Minitab 14 statistical software package.

3.2 Cluster Configuration

By definition, a RAIDb cluster consists of at least two nodes and a controller. The outline of our statistical experiment dictates that there will be two cluster sizes: two- and six-nodes clusters. It is necessary to configure each of these nodes prior to configuring the RAIDb controller.

3.2.1 Cluster Node Hardware

As a middleware distributed database, RAIDb does not require uniform hardware throughout the cluster. Any device capable of connecting to a network and running a database platform is a candidate for a RAIDb cluster node. However, it can be difficult to assure that changes in performance are not caused by the hardware differences that exist in heterogeneous environments. To reduce variability between cluster nodes due to hardware differences, a uniform hardware profile was chosen. This allows for more accurate measures of the relative performance change when cluster characteristics are altered. The hardware profile used on all nodes in the RAIDb clusters that are used in this research can be found in Table 3.4.

Table 3.4. Selected hardware resources for the controller and six nodes.

Machine Property	Selected Hardware
Disk	20 GB Internal 3.5 inch hard drive
RAM	256 MB RAM
CPU	350 MHz Pentium II Processor
Network Interface(s)	100 Megabit 3Com FastEthernet XL Interface Card

3.2.2 Cluster Node Software

Once a hardware profile has been chosen, an operating system must be installed on each node. The operating system, or OS, is responsible for interfacing with all hardware devices in the system. A mature OS will provide simple methods to configure hardware for normal operation, such as setting a network address, communicating with the disks and network, and running the database platform. The operating system chosen for the cluster nodes is the Fedora Core 1 Linux distribution. The Linux operating system, or kernel, version 2.4 is included in this distribution and was installed without modification. A "Minimal" operating system installation was performed with the Fedora

Core 1 distribution installer. As part of the installation, all hard drives in the cluster were partitioned to 20 GB using an ext3 journaling file system.

Networks require a well-defined protocol to allow for communication between the nodes on the network. The network protocol suite TCP/IP is a requirement for C-JDBC operation. For a correct TCP/IP configuration, a network address, conforming to the TCP/IP standard, is required. This was accomplished by setting the correct parameters in the network configuration files. The "hosts" file, a file used to resolve names to TCP/IP addresses without the need for a Domain Name Server (DNS), was also modified to reflect the chosen address scheme. These files can be found in Appendix A. The TCP/IP address assignment for each node in the RAIDb cluster can be found in Table 3.5.

Table 3.5. TCP/IP address and names for all machines in the test environment.

Node Name	Node TCP/IP Address
Node1	172.16.1.1
Node2	172.16.1.2
Node3	172.16.1.3
Node4	172.16.1.4
Node5	172.16.1.5
Node6	172.16.1.6

When the networking subsystem is configured for each database node, a database platform can be installed. A database platform that is simple, due to the small capacities of the nodes, yet complete, is necessary for the database nodes. The TPC-H benchmark requires a database system that supports nested queries, data views, and transaction processing. The PostgreSQL database server, version 7.4.6, complies with these requirements. PostgreSQL was installed on each of the backend nodes using the RPM package for Fedora Core 1 available on the PostgreSQL website.

PostgreSQL server's default installation requires that each database username has a corresponding user account on the host system. This configuration was modified to force PostgreSQL to keep a user account database separate from the host system. Users were then added to the PostgreSQL server database using the user account tools packaged with the database platform. A database by the name of "tpch" was also created using the accompanying tools. Changes to the database server configuration files and scripts can be found in Appendix B. Finally, a compressed archive of the PostgreSQL server installation was made. After each run in the statistical experiment, the stored server image would be restored, resetting the database server to its initial state after installation. These procedures were accomplished on each of the six backend database nodes.

3.2.3 Cluster Topology

To reduce unexplained or uncontrollable variation, a clear cluster topology must be chosen. The nodes in the cluster should be isolated from network devices not associated with the RAIDb cluster. This will reduce unanticipated network congestion not related to the cluster traffic. To accomplish this, all database nodes in the cluster will be networked through a Cisco Catalyst 2900, 100 megabit per second switch. The RAIDb controller will serve as a network gateway between the outside network and the cluster network and will receive all database requests. The internal network interface card of the controller will be networked using the Cisco 2900 switch. The external interface card installed on the controller will be networked with the test client by way of a crossover cable. In commercial environments the external connection of the controller would be networked to another switch, which allows the external network access to the

cluster. This configuration is common to most storage area networks or computer clustering topologies. The topology chosen for the RAIDb cluster used in this research can be seen in Figure 3.1.

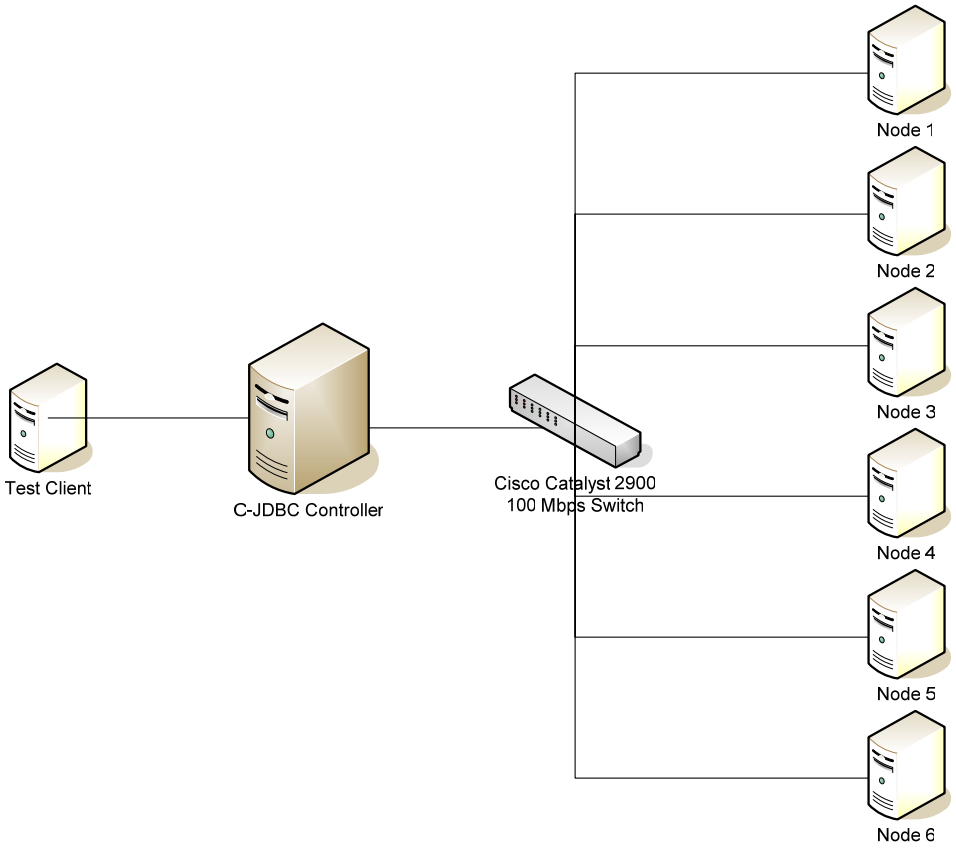

Figure 3.2. RAIDb cluster topology including six backend database nodes and one controller.

This topology allows for easy removal of nodes from the cluster for two-node tests and includes all nodes for six-node tests. The two central nodes, Node3 and Node4, were used for all two-node tests (Figure 3.2). For RAIDb-1 tests that utilized two controllers, multiple controllers are networked through an additional Cisco Catalyst 2900 100 Mbps switch. This topological change can be found in Figure 3.3.

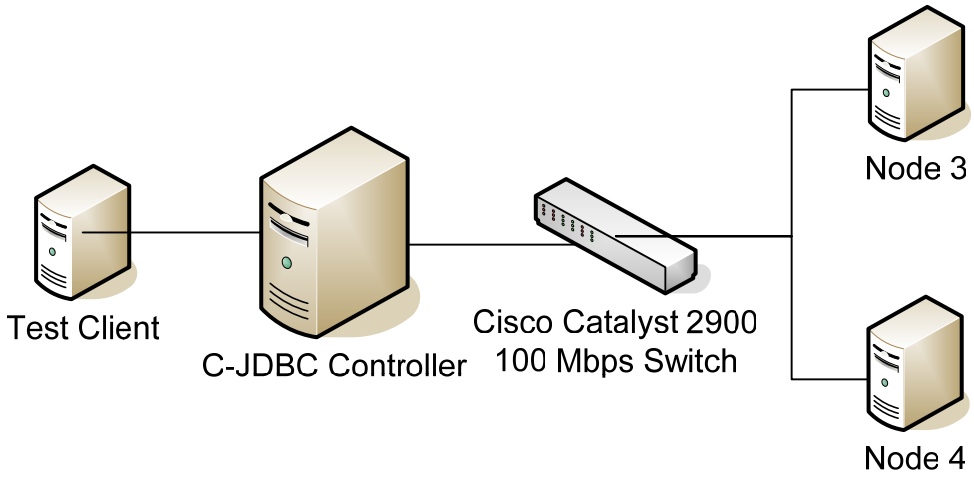

Figure 3.3. Two-node, single controller RAIDb cluster configuration.

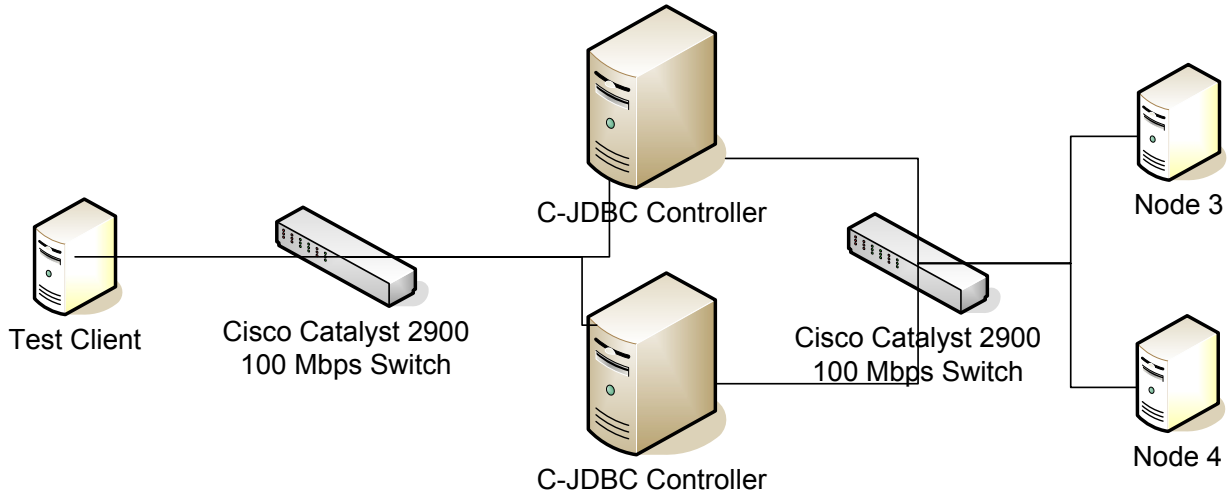

Figure 3.4. Two node, two controller RAIDb cluster configuration.

3.2.4 RAIDb Controller Configuration

The RAIDb controller employs similar software to that of the cluster nodes. A 350 MHz Pentium II processor was used for each of the controllers. Hard drives with 20 GB partitions were used. Each controller contained two 100 Mbps network interface cards for communication. One of these network interface cards is designated the internal interface and directly communicates with the nodes of the RAIDb cluster. The other

network interface card is designated the external interface and receives requests from the database clients. Service requests are received on the external interface, evaluated to determine where they should be routed, and finally sent to the appropriate node on the internal interface.

The RAIDb controller also requires a standard operating system. For consistency, Fedora Core 1 (Linux kernel 2.4) was also installed using the "Minimal" installation method. As part of the installation, a TCP/IP address scheme was chosen for the external and internal interfaces. The internal interfaces are configured with IP addresses to that can be used to communicate with the backend nodes. The TCP/IP addresses for each interface on the controllers can be found in Table 3.6.

Table 3.6. RAIDb controller interfaces and TCP/IP address assignments

Network Interface	TCP/IP Address
Primary Controller – Internal Interface	172.16.0.10
Primary Controller – External Interface	10.0.1.10
Secondary Controller – Internal Interface	172.16.0.11
Secondary Controller – External Interface	10.0.1.11

The C-JDBC software is based on the JDBC database connectivity suite provided by Sun MicroSystems. The C-JDBC controller implementation is also written in Java and requires a Java Runtime Environment (JRE) for operation. Java code is not compiled into native machine code. Instead, Java programs are compiled into a byte code that can be executed on a Java interpreter. This interpreter is commonly called a virtual machine and is bundled with several libraries in the JRE. The Java Runtime Environment version 1.4.2 was installed on the RAIDb controllers using the default method provided by Sun MicroSystems. The JDBC 3 drivers for the PostgreSQL database platform were also installed on each controller. These JDBC drivers were obtained from PostgreSQL's

website and provide a method to connect to the database system on each node in the cluster.

To complete the configuration of the controllers, version 1.0.4 of the C-JDBC software was installed from the compressed archive available from ObjectWeb. Configuration files for the C-JDBC clustering solution are written in XML and specify many important parameters to a C-JDBC-implemented RAIDb controller. Two main configuration file types exist for the controller: the controller configuration file and the virtual database configuration file. The controller configuration file is primarily responsible for the general configuration of the controller and includes options for data backup, TCP configuration, additional controller detection, the location of the virtual database configuration file, and controller message logging.

The virtual database configuration file consists of the definition of one or more virtual databases. A virtual database, as defined by the C-JDBC documentation, is a database published from the C-JDBC controller that consists of the summation of capabilities of each of the database platforms residing on the nodes in the cluster. A virtual database presents a continuous, congruent picture of a database that may be partitioned across several backend database nodes. A virtual database consists of many properties including:

- A name, or a identifier, which allows the database on the controller to be referenced;
- A set of nodes and their JDBC connection strings, usernames, and passwords that are to be included in the virtual database;

- C-JDBC authentication, or the username/password keys used to connect to a virtual database on the controller;
- A request scheduler, or query parser that evaluates where pending queries should be routed;
- A load balancer, or a method to decide how queries should be balanced among replicated nodes;
- And a request cache, which stores common query execution results and data.

For all cluster configurations, the name "tpch" will be used for the virtual database identifier. For two-node configurations, JDBC connection strings for nodes three and four will be used as the source for the virtual database. For six node configurations, all six available nodes will be used as a virtual database source. The username/password combination tpch/tpch will be used as the virtual database authentication key. The query-based request scheduler will be used for all virtual database configurations. The "least pending requests first" load balancing technique will be used for all virtual databases as suggested by Emmanuel Cecchet. This load balancing technique sends queries to the least busy database node first. A query-based request cache will also be used. All configuration files for the C-JDBC controllers can be found in Appendix C.

3.2.5 Test Client Configuration

While possibly the least important machine to the performance of the cluster, the test client is the most complex to configure and setup. The test client may not influence

the performance of the cluster, but does act as the repository for the TPC-H benchmark suite, test results, and message logs.

Once more, for consistency, the Fedora Core 1 Linux Distribution (Linux kernel 2.4) was chosen as the operating system for the test client. A 350 MHz Pentium II processor was used as the test client's CPU. Hard drives with 20 GB ext3 partitions were used. A network interface card capable of 100 Mbps was assigned an address that could be used to submit requests to the external interfaces of the controllers. The networking address of 10.0.1.1 was chosen for the test client.

3.2.5.1 Test Client Software Configuration

The test client contains all the necessary programs to connect and submit queries to the cluster controller, the programs and files necessary to run the TPC-H benchmark, and the directory structure to log experimental results. To connect to the cluster controller, the C-JDBC database driver is needed. This driver is packaged with the controller software that was installed on each controller machine. The driver file was copied from one of the controller machines and placed in the Java classpath. A directory structure to store the results from the experiment tests was also created and can be seen in Figure 3.4.

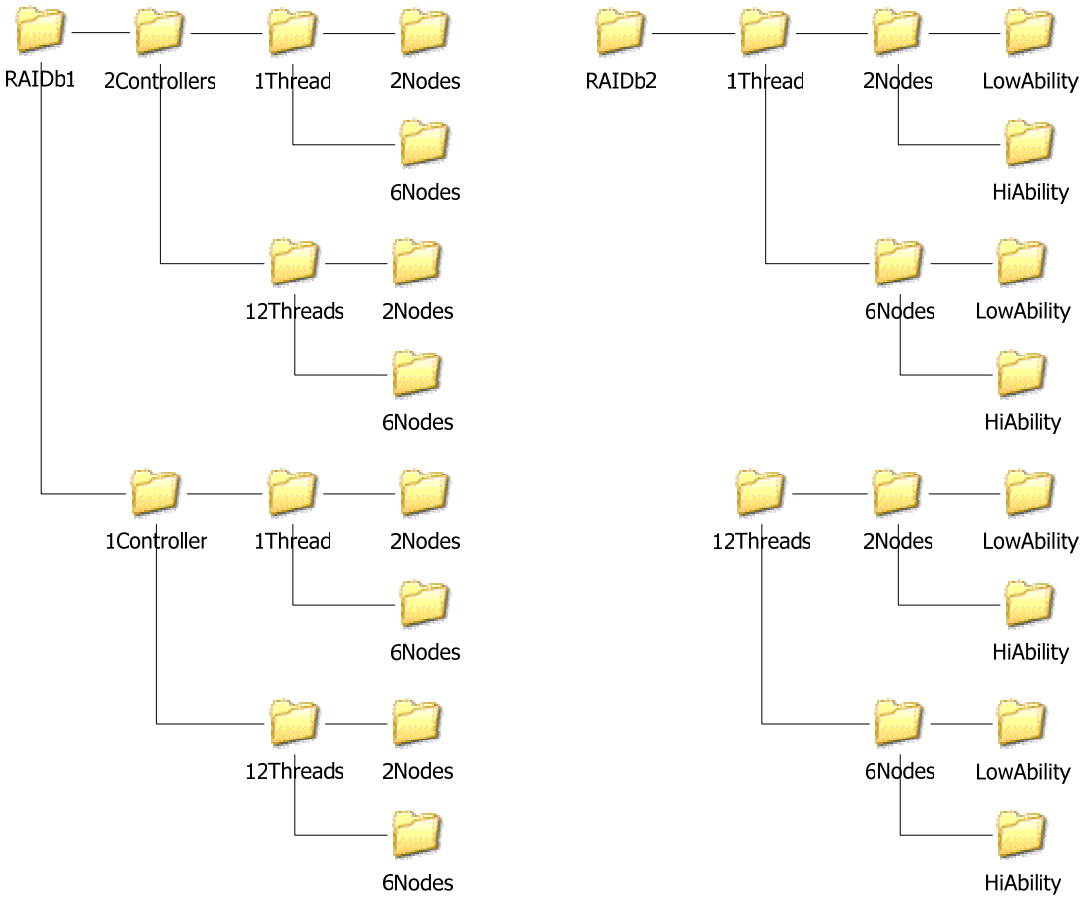

Figure 3.5. Directory structure for RAIDb1 and RAIDb2 experiment results.

To properly execute the TPC-H benchmark, several additional software tools were installed on the test client. As with the controller, version 1.4.2 of Sun's JRE was installed on the test client for JDBC support and benchmark program execution. The Gnu C Compiler (gcc) was later installed on each machine to allow for the compilation of the benchmark suite and database platforms.

3.2.5.2 TPC-H Benchmark

The TPC-H benchmark is a decision support benchmark. TPC-H consists of a stream of unoptimized, business oriented queries. The Transaction Processing Performance Council has developed the queries and data contained in the TPC-H

benchmark "to have broad industry-wide relevance while maintaining a sufficient degree of ease of implementation." This benchmark tests systems for "ability to examine large amounts of data and execute complex queries. Query sets included in the TPC-H benchmark simulate generated ad-hoc queries, are far more complex than most OLTP transactions, and generate intensive activity on the part of the database server."

The TPC-H benchmark reports performance as a function of database size. Database size is reported in raw data size, or size of the data when generated in comma-delimited text files. Several standard database sizes are permitted in the TPC-H benchmark, which are reported by a corresponding Scale Factor (SF). Scale Factor values and related raw database sizes can be seen in Table 3.7. A Scale Factor of 1 was chosen for this research.

Table 3.7. Scale Factors (SF) and raw database size.

Scale Factor	Raw Database Size
1	1 GB
10	10 GB
30	30 GB
100	100 GB
300	300 GB
1000	1000 GB
3000	3000 GB
10000	10000 GB
30000	30000 GB
100000	100000 GB

The query sets used for the TPC-H benchmark are performed against a database schema consisting of eight tables (Figure 3.5). The size of some of these tables increases proportionally to the Scale Factor, including tables Part, Partsupp, Lineitem, Orders, Supplier, and Customer. The tables Nation and Region are of constant size regardless of

the Scale Factor chosen. The cardinality, in rows per table, can also be found in Figure 3.5.

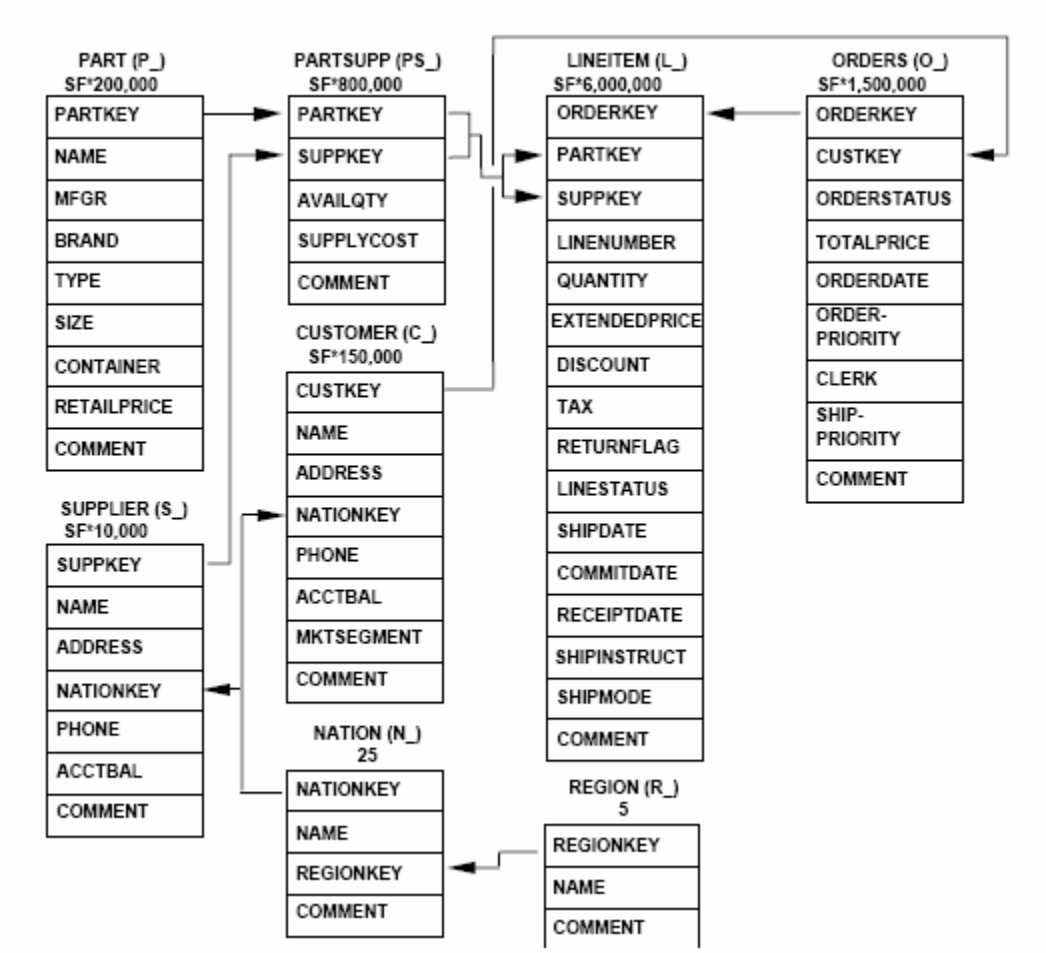

Figure 3.6. TPC-H eight table database schema and cardinality.

As part of the TPC-H benchmark, two distinct tests, the load test and the performance test, must be performed in order. Each type of test is designed to measure a distinct aspect of the performance of the database system. One execution of load test is followed by two sets of the performance test. The performance test consists of two component tests: the load test and the power test. The sequence of database performance tests executed by the TPC-H benchmark can be seen in Figure 3.6.

Figure 3.7. The flow of a TPC-H benchmark run.

The load test measures the time a database system spends to create tables and indexes and to load data into the database engine. The load test allows the use of proprietary loading engines and tools to load data from pre-existing comma-delimited text files. These comma-delimited text files can be created using the DBGEN program provided with the TPC-H specification. The DBGEN program generates data that is compliant to the TPC-H specification. A custom Java program was created to automate the generation and loading of the data. The source code to the Java implementation of the TPC-H benchmark used in this research can be found in Appendix D.

The power test, one component of the performance test, consists of three distinct parts: the first refresh function, one serial executing of the TPC-H query set, and the second refresh function. The first refresh function, or New Sales Refresh Function,

"inserts new rows into the ORDERS and LINEITEM tables in the database following the scaling and data generation methods used to populate the database." The additional data that will be used to execute the first refresh function is created using the DBGEN program which is bundled with the TPC-H specification. A Java program was written to automate the execution of the New Sales Refresh Function (Appendix D).

The second refresh function, or Old Sales Refresh Function, "removes rows from the ORDERS and LINEITEM tables in the database to emulate the removal of stale or obsolete information." Data must be removed from the ORDERS and LINEITEM table in a way that will maintain ACID properties of the data set. In other words, the removal of the items from the database must be performed with a transaction. The item identifiers of the records to be removed are generated using the DBGEN program. A Java program was also written to automate the execution of the Old Sales Refresh Function.

The TPC-H query stream consists of 22 queries. Order of execution of the queries is randomized using a random number table. Execution order and query values are generated from SQL templates and the QGEN program available with the TPC-H specification. Once generated with QGEN, the randomized queries are sent to the RAIDb controller for execution. Java programs have been written to automate the power test. The source code for these programs can be found in Appendix D.

The throughput test, the second component of the performance test, consists of multiple query streams performed in parallel. The TPC-H specification details the lowest number of parallel query streams required for the database size being used. A table showing the allowable number of query streams can be found in Table 3.8.

Table 3.8. The number of required query streams per TPC-H scale factor (SF).

SF	Number of Required Streams
1	2
10	3
30	4
100	5
300	6
1000	7
3000	8
10000	9
30000	10
100000	11

The RAIDb cluster used in this research will be using a Scale Factor of 1, or a raw data size of 1 GB. It is therefore required to use at least two query streams in parallel. There is not a ceiling for query streams in the TPC-H specification.

The TPC-H benchmark also contains tests to ensure that a database system is ACID-compliant. These tests include several queries run against the data loaded into the database system. To test atomicity, the TPC-H specification requires executing a transaction, allowing the transaction to commit, and testing the database system for correct values. To further test atomicity, the TPC-H specifications require that the same transaction be executed, substituting a rollback for the commit, and verifying that the data remains unchanged. To test consistency, the total price of an order is compared to the sum of the line items for that order. The system is consistent if these numbers match. Each of the isolation failures discussed in chapter two must also be tested. This is accomplished with a set of six transactions designed to test these failures. Comparing the results of these six tests with expected values allows the isolation requirement to be verified. To test durability, a transaction is executed and the permanent storage of the data changed, added or removed from the database system is verified. The queries that

drive each of these tests will be generated using the tools provided with the TPC-H specification.

3.3 Research Procedure Conclusions

In this chapter, a method for creating a statistical performance model has been established. A three factor, two-level factorial design was chosen to isolate the factors that contribute the most to the performance of a RAIDb cluster. Two replicates will be performed creating 16 experimental runs. The results from this experiment were used to produce a regression model. This regression model can be used to predict the performance of a system. The strength of the regression model can also be evaluated to determine the reliability of the regression. A confidence interval of 99% was used to determine the confidence level for the results.

The TPC-H benchmark was used to capture data about the cluster. The benchmark was executed against each cluster configuration specified in the statistical experiment design. One benchmark run consists of one load test execution and two executions of the performance test. Execution of the TPC-H benchmark were automated and facilitated through the use of scripts and Java programs. The results from each run were stored in a pre-structured repository. The cluster was then reconfigured for the next run in the statistical experiment design.

Chapter 4

4 RESULTS AND DATA ANALYSIS

In this chapter the data collected from the procedures described in chapter three will be discussed and analyzed. The results of this analysis were used to identify significant performance factors in a RAIDb cluster. The results will also be used to generate a regression model to predict the performance of the database clusters.

4.1 RAIDb-1 Cluster Performance

The most common purpose of establishing a RAIDb database cluster is to create the highest performance advantage with available resources. Analyzing each factor in the various tests included in the TPC-H benchmark can be useful to determine the method for highest performance gain.

4.1.1 Overall Results

To determine the significance of each variable used in this research, the statistical experiment design discussed in chapter three was executed. Table 4.1 shows the statistical experiment design, including the standard order, the run order, and the replicate number. Table 4.1 also shows the high (+) and low (-) configuration levels for the experimental variables with the corresponding run time (Overall), measured in hours, of the TPC-H benchmark.

Table 4.1. Statistical Experiment Design with overall results from the TPC-H benchmark in hours

Sequence	Run Order	Replicate	Cluster State			Overall
			Controllers	Threads	Nodes	
1	11	1	+	+	+	16.4678
2	6	1	+	+	-	19.8719
3	15	1	+	-	+	19.3572
4	8	1	+	-	-	19.3207
5	14	1	-	+	+	16.3978
6	13	1	-	+	-	19.4544
7	12	1	-	-	+	18.8002
8	9	1	-	-	-	19.2476
9	4	2	+	+	+	15.9542
10	16	2	+	+	-	19.8116
11	3	2	+	-	+	19.1165
12	1	2	+	-	-	18.6261
13	10	2	-	+	+	16.9484
14	5	2	-	+	-	19.3797
15	2	2	-	-	+	18.8977
16	7	2	-	-	-	18.871

4.1.1.1 F-Statistic

The F-statistic describes which factors and interactions in the experiment results significantly affect the regression model. The F-statistic is calculated from the ANOVA values generated with the results. When a value for the F-statistic is calculated, a corresponding p-value can be generated. If the p-value of the statistic is greater than the chosen significance level ($\alpha = 0.01$), the ANOVA source carries no significant effect on the regression model. If the p-value of the statistic is less than the significance level, significant effects do exist. The ANOVA tests for the RAIDb-1 factorial experiment can be found in Table 4.2.

Table 4.2. Analysis of Variance for Main Effects and Interactions of the RAIDb-1 tests.

Analysis of Variance for RAIDb-1						
Source	*DF*	*Seq SS*	*Adj SS*	*Adj MS*	*F*	*P*
Main Effects	3	13.9599	13.9599	4.65329	58.72	0
2-Way Interactions	3	10.4006	10.4006	3.46687	43.75	0
3-Way Interactions	1	0.4628	0.4628	0.46283	5.84	0.042
Residual Error	8	0.634	0.634	0.07925		
Pure Error	8	0.634	0.634	0.07925		
Total	15	25.4573				

Table 4.2 shows a p-value for main effects to be less than 0.0001. The p-value for two-way interactions is also less than 0.0001. The very low p-values for each of these ANOVA sources indicates that a study of the main effects and two-way interactions is necessary to determine which of the factors, and which interactions between multiple factors, are causing such significant variation in the results of the tests.

4.1.1.2 Main Variable Effects

Many of the statistics used to evaluate the significance of the factors in this experiment are best calculated from the averages of the execution time for identical cluster configurations. The average time for each configuration allows the statistical analysis to provide more meaningful and accurate inferences. These averaged results were used as a basis for all statistical analysis and regression and are found, with the accompanying configuration of high (+) and low (-) variable levels, in Table 4.3.

Table 4.3. Statistical Experiment Design with averaged overall results from the TPC-H benchmark in hours

Cluster State			Overall
Controllers	Threads	Nodes	
+	+	+	16.21099986
+	+	-	19.84176153
+	-	+	19.23686903
+	-	-	18.97340722
-	+	+	16.67306639
-	+	-	19.41705319
-	-	+	18.84893097
-	-	-	19.05932569

The effect of a factor on the experiment results is the difference between the average outcome at the high level of operation and the average outcome at the low level of operation. These effects explain the amount of influence a factor has on the overall experiment results. If a factor has little influence on the outcome of the test overall, the average execution time of the high level for all sixteen runs should be approximately equal to the average execution time of the low level for all sixteen runs. Should a factor have significant influence on the outcome of the test, a large difference in the averages of the high and low level tests should be apparent. The effects for the three factors used in the RAIDb-1 tests are graphed in Figure 4.1.

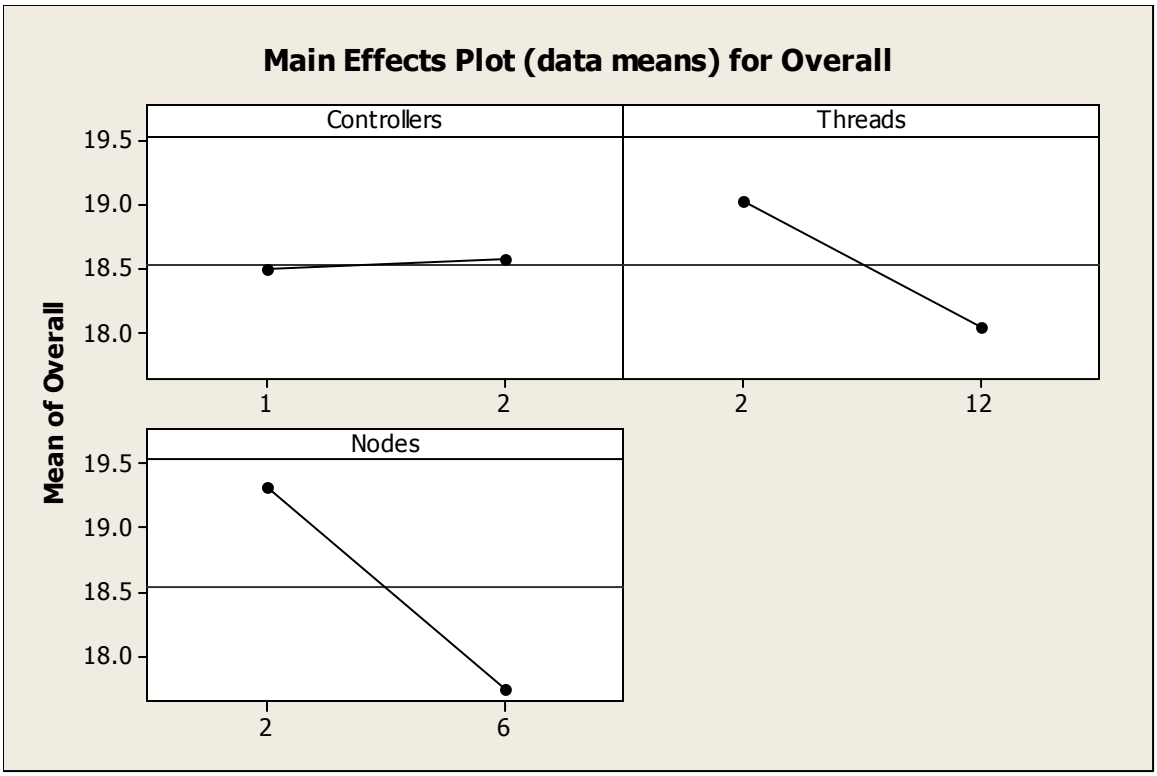

Figure 4.1. Graphs of the effects of three factors in the RAIDb-1 experiment.

As can be seen in Figure 4.1, a very weak effect can be observed for the change in the number of controllers. However, plots for thread and node factors indicate that there is significant effect between high and low levels. Figure 4.1 also shows that performance of RAIDb-1 clusters increases as the number of replicated nodes also increases. The inverse of this statement is also true. As the number of nodes in the cluster decrease, the performance of the cluster also decreases. Threads also affect the performance of the system. Like the number of nodes, an increase in the number of threads will also increase the performance of the system. Decreasing the number of threads will also decrease the performance of the system.

An interesting interpretation of these results is that RAIDb-1 clusters handle high load very efficiently, yet performance decreases as load decreases. This observation suggests that RAIDb-1 clusters are most effective in high load environments. These

results also suggest that the effectiveness of the cluster increases as the size of the cluster increases. As cluster size increases, performance also increases.

While the graphs in Figure 4.1 give us a general idea of the effect for each of the factors, a test statistic is needed to determine if these effects are truly significant. The effect values can be converted to a t-statistic, which can be used as a measure of the significance of these factors. The t-statistic for the effect for each of the three factors can be seen in Table 4.4.

Table 4.4. Sample effects, t-statistics and p-values for all main factors.

Estimated Effects and Coefficients			
Term	*Effect*	*t*	*p*
Controllers	0.0662	0.47	0.651
Threads	-0.9939	-7.06	0
Nodes	-1.5804	-11.23	0

The t-statistic for controllers is equal to 0.47 with a corresponding p-value of 0.651. As the significance level for this experiment is low ($\alpha = 0.01$), this p-value is extremely high. A high p-value indicates that controllers are not a significant factor in the performance of RAIDb-1 clusters and should not be included in the regression model. The t-statistic for threads is equal to -7.06 with a corresponding p-value of less than 0.0001. This extremely low p-value indicates that threads are a major contributor to the regression model. The t-statistic for nodes is equal to -11.23 with a corresponding p-value also less than 0.0001. Node number is also a significant contributor to the regression model. As a result of this statistical analysis, it is reasonable to conclude that the number of threads and the number of nodes in a RAIDb-1 cluster should be included in building a performance model. Furthermore, our observations about the performance

of the cluster, as a function of thread and node number, are accurate. It is also reasonable to conclude that the number of active controllers does not affect the performance of the clusters used in this thesis.

4.1.1.3 Two-way Interaction Effects

The synergistic effect of multiple factors can also be measured. These effects, created by a combination of factors, are called interactions. Interaction effects enhance the resolution of the model by describing influences that multiple factors jointly demonstrate on experiment results. While main effects can provide an overall picture of the influential main factors, interaction effects begin to describe when the significant change in the main factors occur in relation to the other factors.

The significance of factor interactions are calculated using the t-statistic. The t-statistics and their corresponding p-values can be found in Table 4.5.

Table 4.5. Interaction effects with calculated t-statistics and p-values.

Estimated Effects and Coefficients for Overall (coded units)			
Term	*Effect*	*t*	*p*
Controllers*Threads	-0.0848	-0.6	0.563
Controllers*Nodes	-0.1032	-0.73	0.484
Threads*Nodes	-1.607	-11.42	0

All two-factor interactions, except the Thread-Node interaction, are insignificant as the p-value for these effects are lower than the chosen significance level. The Thread-Node interaction has a t-statistic of -11.42 with a corresponding p-value less than 0.0001. Due to the small p-value, it is reasonable to conclude that the number of threads and the

number of nodes interact to create a cooperative effect on the performance of the cluster. A plot of each of these interactions can be found in Figure 4.2.

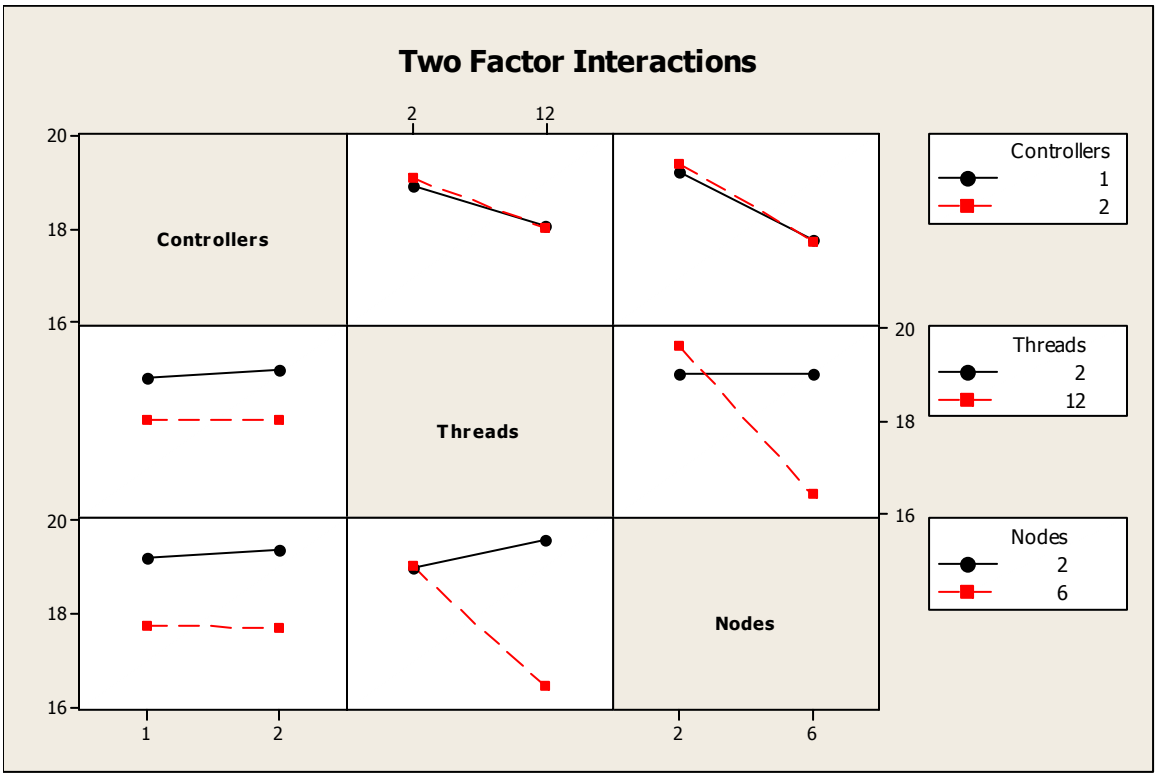

Figure 4.2. Two factor interactions for RAIDb-1 experiment results.

Several observations can be made from the statistics in Table 4.5 and the graphs in Figure 4.2. One of these observations is that altering cluster size has little or no effect while the workloads stay small. As workload remains small, the performance enhancements obtained by adding more nodes is neglible; a two-node system is just as able to handle the low workload as a six-node cluster. However, as workload increases, a two-node cluster becomes less able to handle the required processing. By increasing the number of nodes available to the system, the ability of the cluster to process the requests and enhance performance greatly increases. Through increased load balancing and caching abilities, a six node cluster has greater capacity to process requests during high

load runs than low load runs. Increasing cluster size is an effective method for increasing performance when the workload of the cluster is high.

4.1.1.4 Statistical Performance Model for RAIDb-1 Clusters

To create a regression model, the coefficients for each of the significant factors and interactions must be calculated. Only two main factors and one two-factor interaction were found to be significant. Each of the coefficients is relevant to the model up to a 99% confidence interval. Simply put, the terms in this model will significantly apply to this model 99% of the time. The coefficients for any insignificant factors have been excluded. These coefficients can be found in Table 4.6.

Table 4.6. Regression coefficients for the model constant and main factors.

Estimated Coefficients for RAIDb-1 Tests	
Term	*Coefficients*
Constant	18.5591
Threads	0.221999
Nodes	0.167329
Threads*Nodes	-0.0803477

As stated in chapter three, the regression model takes the form of $y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \ldots + \beta_n x_n + \varepsilon$, where β_0 symbolizes the y-intercept, or constant, of the equation, x_i indicates a variable included in the treatments, β_i indicates the coefficient of the related treatment variable, and ε symbolizes the error inherent in modeling estimations. Using the coefficients for each of the factors and interactions in Table 4.6, the following regression model can be created:

Execution Time = 18.5591 + 0.221999(Number Of Threads) + 0.167329*(Number of Nodes) – 0.0803477*(Number of Threads)*(Number of Nodes)*

4.1.1.5 Understanding the Model Constant

The constant in the statistical regression model serves to calibrate the model to the machine speed of the computing nodes. The constant used for the machines in the RAIDb-1 configurations is 18.5591. By changing the speed of the machine used to build the cluster, this constant will change. While some evidence shows that the relative performance changes between machine speeds is equal (Cecchet, 2003 and Rogers, 2004), the exact value of this constant for machines other than those tested in this thesis can not be determined from this model. An estimate of this constant for different machine speeds can be obtained by executing the TPC-H benchmark on a RAIDb-1 cluster. The results from the benchmark can be used to solve the equation for the missing constant. A more definite evaluation of the model based on machine speed should be obtained through further research.

4.1.1.6 Coefficient of Determination

While a regression model that can be used to predict the behavior of RAIDb clusters can be built with the factor and interaction coefficients, the accuracy of the model applied to the data in this research must be verified. One method to accomplish this is to use the coefficient of determination, or R^2. The R^2 value for RAIDb-1 cluster was calculated to be 0.9751. This means that 97.51% of all the variation in the data is

explained by the regression model stated above. Less than three percent of the data found in this research is unexplained. This high value for the coefficient of determination indicates an excellent goodness-of-fit, or accuracy, of the regression model.

An adjusted value for R^2 was also calculated. The coefficient of determination can often be overestimated during ANOVA or factorial experiments. This is caused by the relatively small sample needed to obtain large amounts of information. To allow for this inflated value of R^2, an coefficient of determination adjusted for the small sample size and overestimation is often calculated. The adjusted R^2 value for RAIDb-1 cluster is 0.9533, or 95.33%. This adjusted value explains that, in reality, only 95.33% of the variation in the performance of the RAIDb-1 cluster is accounted for by the regression model. Less than five percent of the variation remains unexplained. While slightly lower value, the adjusted coefficient of determination is an excellent indicator that the regression model can predict, to a high accuracy, the performance of a RAIDb-1 cluster.

4.1.2 ACID Compliance

The performance model of any database system would be useless without an assurance that the data is kept pure through ACID-compliance. All tests specified in the TPC-H specification were executed. All nodes in the cluster, as well as the system as a whole, are ACID-compliant. The results for each of these tests can be seen in Table 4.7.

Table 4.7. Results of TPC-H ACID tests for individual nodes and entire RAIDb-1 system.

System Tested	ACID Test	Pass/Fail	Conclusion
Node 1-6	Atomicity	Pass	PostgreSQL transaction implementation
	Consistency	Pass	PostgreSQL-implemented consistency-compliance
	Isolation	Pass	PostgreSQL-implemented isolation-compliance
	Durability	Pass	Postgresql hard disk storage
RAIDb-1 System	Atomicity	Pass	All PostgreSQL transactions correctly passed to all relevant backends for handling by database platform
	Consistency	Pass	PostgreSQL -implemented consistency-compliance correctly supported by RAIDb controller
	Isolation	Pass	PostgreSQL -implemented isolation-compliance correctly supported by RAIDb controller
	Durability	Pass	RAIDb controller correctly reports finished transactions made compliant through database platform storage mechanism

4.1.3 Statistical Model Validation

The statistical performance model presented in this thesis for RAIDb-1 clusters is mathematically valid for two- and six-node clusters. However, to extend the validity of the model to larger clusters, some minimal validation of the model must exist. To verify the accuracy of this model, results obtained from experiments with larger clusters must be statistically equal to the predicted values obtained from the model. This is accomplished by employing the prediction interval of the model. A prediction interval specifies a lower and upper bounds for the cluster configuration. If the results obtained from experiments designed to validate the model fall within these bounds, the predicted and actual results are accepted as equal. Any variation from the predicted values that is within this

prediction interval is credited to the random variation that exists in any system (Keller, 2001).

Using Minitab 14, prediction intervals were generated for eight- and ten-node clusters. Experiments for the larger cluster size were then executed. These prediction intervals, with the results from the corresponding experiments, can be found in Table 4.8.

Table 4.8. Prediction intervals for 8- and 10-node RAIDb-1 clusters.

Cluster Configuration			**Prediction Interval**		**Actual Results**
Controllers	*Threads*	*Nodes*	*Lower Bound*	*Upper Bound*	
1	12	8	13.8842	16.7180	15.5905
1	12	10	12.1619	15.6963	15.4732

The prediction interval for an eight-node RAIDb-1 cluster is from 13.8842 hours to 16.7180 hours. The validation experiment resulted in an execution time of 15.5905 hours. The validation results for an eight-node cluster are within the bounds of the prediction interval and are, therefore, statistically equal. The predication interval for a ten-node cluster ranges from 12.1619 hours to 15.6963 hours. The validation experiment resulted in an execution time of 15.4732 hours. The validation results for a ten-node cluster are also within the bounds of the prediction interval and are statistically equal. A graph of the predicted execution times, the prediction interval and actual values can be found in Figure 4.3.

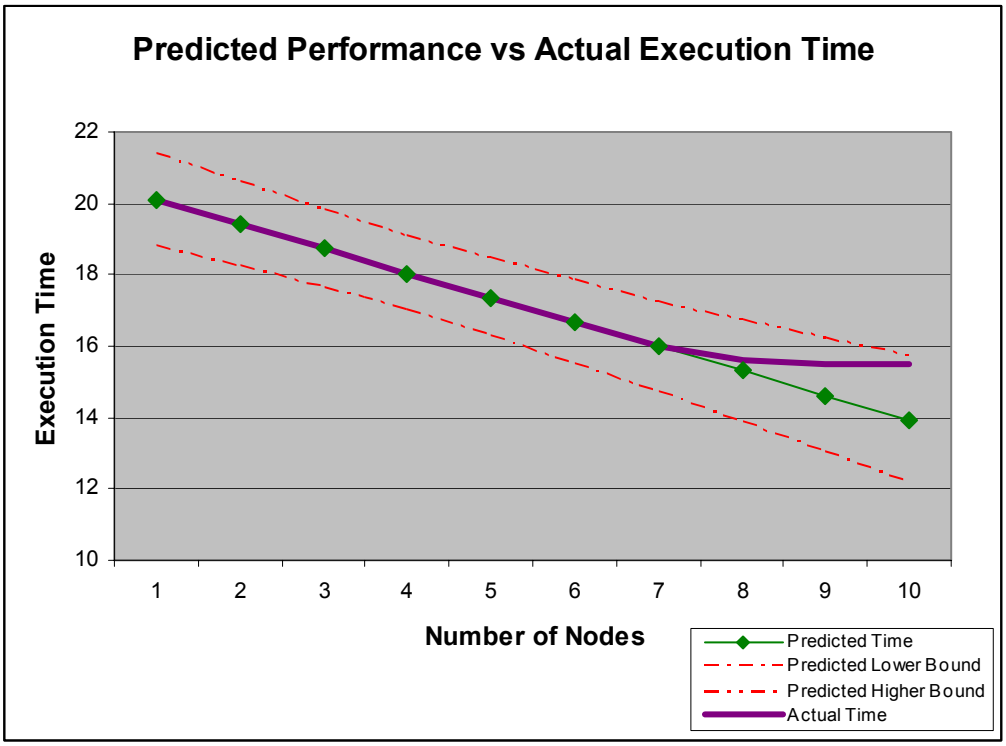

Figure 4.3. Predicted and actual execution times within the prediction interval for RAIDb-1 clusters.

Although statistically equal, execution times for ten-node systems begin to diverge from the predicted model. This behavior is representative of RAIDb-1 clusters. As shown in section 4.1.1.3, the performance enhancements obtained by increasing the cluster size decreases as the work presented to each node decreases. By adding more nodes to the cluster in the validation experiments, the workload per node has decreased. The resulting decrease in performance for eight- and ten-node clusters is shown in Figure 4.3

Other possible reasons for this departure are many: the test client is beginning to bottleneck with server requests; the ability of the RAIDb-1 controller to enhance performance through load balancing, caching or other techniques is beginning to reach a limit; or, that the nature of the TPC-H benchmark is such that execution time can not be

compressed any further. While the performance model is statistically valid for clusters with up to ten nodes, further research should be conducted to explain the divergent behavior of a ten-node cluster.

4.1.4 RAIDb-1 Conclusions

An effective performance model for RAIDb-1 clusters has been created. This performance model was built using a three factor, two-level factorial design. The coefficients generated from each factor apply to the model with a 99% confidence interval. The correlation of determination shows that a 99% certainty exists that 95% of the variation in the data is explained by this performance model. This performance model can be used to predict the behavior of RAIDb-1 clusters containing two to six nodes. Although the model is statistically valid for clusters with more than six nodes, validation experiment results indicate that clusters with eight and ten nodes begin to diverge from the regression model. Possible causes of this divergence include the ability of the test client to sufficiently load the cluster or performance degradation of the RAIDb-1 controller.

4.2 RAIDb-2 Cluster Performance

4.2.1 Overall Results

To determine the significance of each variable used in this research, the statistical experiment design for RAIDb-2 discussed in chapter three was executed. Table 4.9 shows the statistical experiment design, including the standard order, the run order, and

the replicate number. Table 4.9 also shows the high (+) and low (-) configuration levels for the experimental variables with the corresponding run time (Overall), measured in hours, of the TPC-H benchmark.

Table 4.9. Statistical Experiment Design with overall results from the TPC-H benchmark in hours

Sequence	Run Order	Replicate	Cluster State			Overall
			Ability	Threads	Nodes	
1	12	1	+	+	+	16.173
2	9	1	+	+	-	18.9718
3	4	1	+	-	+	18.9138
4	15	1	+	-	-	19.0348
5	13	1	-	+	+	17.5845
6	16	1	-	+	-	19.3516
7	7	1	-	-	+	18.9784
8	5	1	-	-	-	19.0204
9	11	2	+	+	+	16.3852
10	8	2	+	+	-	19.3463
11	2	2	+	-	+	18.5505
12	10	2	+	-	-	18.9062
13	3	2	-	+	+	17.5281
14	14	2	-	+	-	19.1857
15	6	2	-	-	+	18.5724
16	1	2	-	-	-	19.0348

4.2.1.1 F-Statistic

As with the RAIDb-1 cluster results, the F-statistics for RAIDb-2 cluster results can be calculated to determine the factors that influence the regression model. The calculated F-statistics, for the RAIDb-2 cluster results, with the corresponding p-values can be found in Table 4.10.

Table 4.10. Analysis of Variance for Main Effects and Interactions of the RAIDb-2 tests.

Analysis of Variance for RAIDb-2						
Source	*DF*	*Seq SS*	*Adj SS*	*Adj MS*	*F*	*P*
Main Effects	3	9.6403	9.6403	3.21342	97.10	0
2-Way Interactions	3	4.9528	4.9528	1.65094	49.89	0
3-Way Interactions	1	0.3490	0.3490	0.34898	10.55	0.012
Residual Error	8	0.2648	0.2648	0.03309		
Pure Error	8	0.2648	0.2648	0.03309		
Total	15	15.2068				

Table 4.10 shows a p-value for main effects equal to approximately zero. The p-value for two-way interactions is also approximately zero. The very low p-values for each of these ANOVA sources indicates that a study of the main effects and two-way interactions is necessary to determine which of the factors, and which interactions between multiple factors, are causing such significant variation in the results of the tests. The p-value for three-way interactions is equal to 0.012. While very low, this value is not less than the chosen significance level and cannot be considered significant to the regression model.

4.2.1.2 Main Variable Effects

Like the analysis for RAIDb-1 cluster results, the average time for each cluster configuration allows the statistical analysis to provide more meaningful and accurate inferences. These averaged results, with the accompanying configuration of high (+) and low (-) variable levels, can be found in Table 4.11.

Table 4.11. Statistical Experiment Design with overall results from the TPC-H benchmark in hours

Cluster Configuration			Overall
Ability	Threads	Nodes	
+	+	+	16.27906194
+	+	-	19.15907083
+	-	+	18.73212819
+	-	-	18.97048181
-	+	+	17.55631708
-	+	-	19.26866
-	-	+	18.77538792
-	-	-	19.02757028

As with the analysis of the RAIDb-1 cluster results, effects for each factor can be calculated to identify significant factors. The effects for the three factors can be graphed to depict their level of significance. The effects for the three factors used in the RAIDb-2 tests are graphed in Figure 4.4.

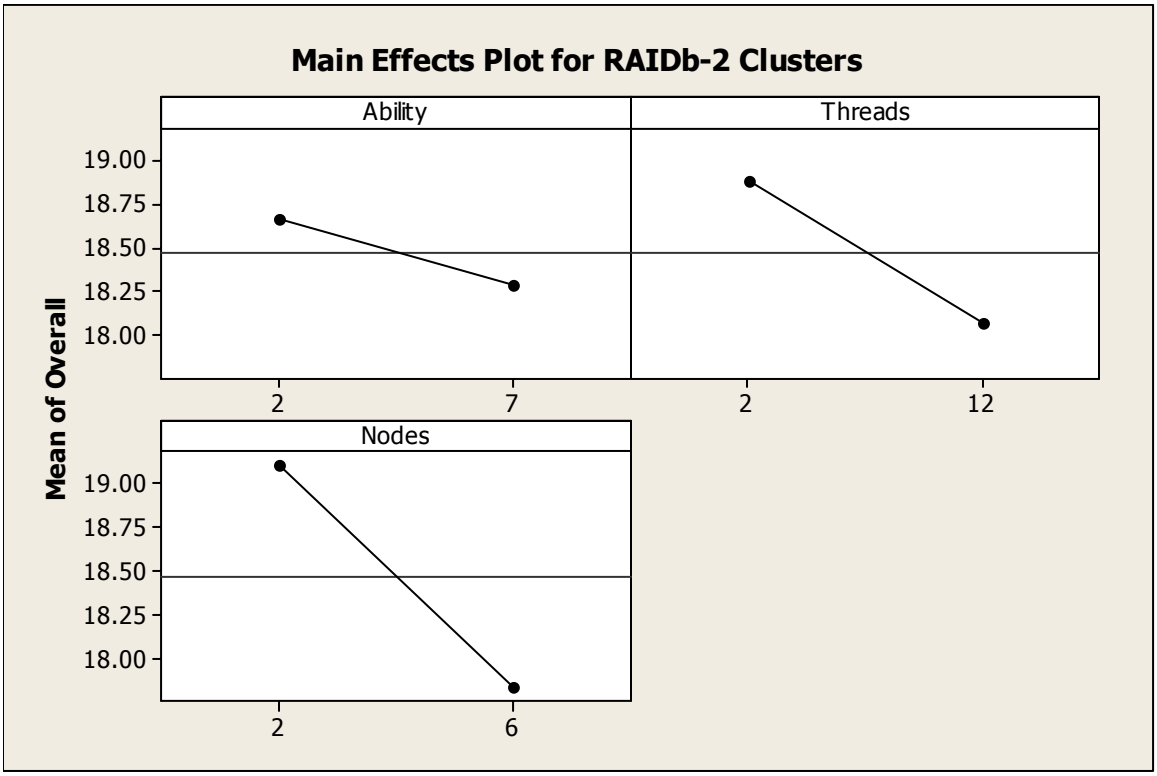

Figure 4.4. Graphs of the effects of three factors in the RAIDb-2 experiment.

Relatively strong effects exist for all three factors of the experiment. As the ability of the cluster (the number of tables available on each node), the number of threads, or the number of nodes increases in a cluster, the performance of the cluster also increases. It can also be observed that the number of available nodes has the greatest effect on cluster performance. Secondly, the number of threads, or system load, causes the next most visible effect on cluster performance. Finally, modifications to the various levels of the cluster ability provide the least impact on performance.

The graphs in Figure 4.3 only provide a general idea of the effects for each of the factors. A test statistic is needed to determine if these effects are truly significant. The effect values can be converted to a t-statistic, which can be used as a measure of the significance of these factors. The t-statistic for the effect for each of the three factors can be seen in Table 4.12.

Table 4.12. The effects, t-statistics, and p-values for all main factors.

Estimated Effects			
Term	*Effect*	*t*	*p*
Ability	-0.3718	-4.09	0.003
Threads	-0.8106	-8.91	0
Nodes	-1.2707	-13.97	0

The t-statistic for ability is equal to -4.09 with a corresponding p-value of 0.003. This low p-value indicates that the level of ability is significant to the regression model. The t-statistic for threads is equal to -8.91 with a corresponding p-value of approximately zero. This p-value indicates that threads are a major contributor to the regression model. The t-statistic for nodes is equal to -13.97 with a corresponding p-value also approximately zero. Node number is also a significant contributor to the regression

model. As a result of this statistical analysis, it is reasonable to conclude that the node ability, the number of threads and the number of nodes in a RAIDb-2 cluster should be included in building a performance model.

4.2.1.3 Two-way Interaction Effects

The interaction effects of multiple factors in RAIDb-2 clusters can also be measured. The significance of multi-factor interactions are also calculated using the t-statistic. The t-statistics and their corresponding p-values for interaction effects can be found in Table 4.13.

Table 4.13. Interaction effects with calculated t-statistics and p-values.

Estimated Effects for Two-way Interactions			
Term	*Effect*	*t*	*p*
Ability*Threads	-0.3216	-3.54	0.008
Ability*Nodes	-0.2885	-3.17	0.013
Threads*Nodes	-1.0255	-11.27	0.000

As noted by the F-statistic, all two-factor interactions, except the Ability-Node interaction, are significant as the p-value for these effects are lower than the chosen significance level. The Ability-Node interaction has a t-statistic of -3.17 with a corresponding p-value equal to 0.013. This p-value is larger than the significance level and, therefore, it can be inferred that the Ability-Node interaction is not significant. The t-statistic for the Ability-Thread interaction is equal to -3.54. This value for the t-statistic has a corresponding p-value of 0.008. It is reasonable to conclude that the Ability-Thread interaction is significant. The t-statistic for the Thread-Node interaction has a value of -11.27. This value corresponds to a p-value of approximately zero. It is also reasonable

to infer that the Thread-Node interaction is significant for the RAIDb-2 cluster regression model. A plot of each of these interactions can be found in Figure 4.4.

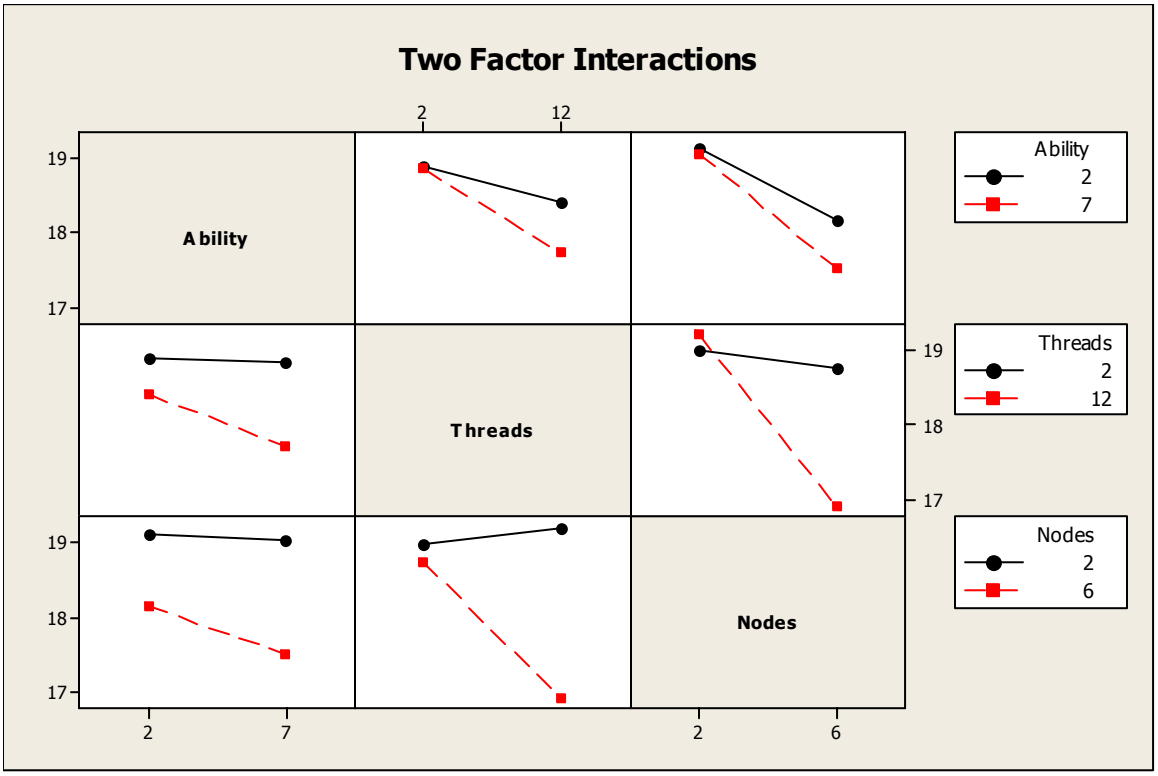

Figure 4.5. Two factor interactions for RAIDb-2 experiment results.

As can be seen in Table 4.13 and Figure 4.4, the number of threads and number of available nodes interact. As with RAIDb-1 clusters, increasing cluster size during high load tests greatly increases performance. Jointly modifying the number of threads and cluster ability also result in cluster performance enhancements. During high load, increasing the ability of the cluster produces a large improvement in execution time.

4.2.1.4 Statistical Performance Model for RAIDb-2 Clusters

The coefficients for each of the significant factors and interactions were calculated. All three main factor effects were found to be significant. Both the Ability-

Thread interaction effect and the Thread-Node interaction effect were also found to be significant. The coefficients for each of these significant effects can be used to create a regression model. Each of the coefficients is relevant to the model up to a 99% confidence interval. Simply put, the terms in this model will significantly apply to this model 99% of the time. The coefficients for any insignificant factors have been excluded. These coefficients can be found in Table 4.14.

Table 4.14. Regression model coefficients for RAIDb-2 clusters.

Estimated Coefficients for RAIDb-2 Tests	
Term	*Coefficients*
Constant	18.8030
Ability	0.0156951
Threads	0.181922
Nodes	0.0412284
Ability*Threads	-0.0128650
Threads*Nodes	-0.0512727

Using the coefficients for each of the factors and interactions in Table 4.14, the following regression model can be created for RAIDb-2 clusters:

Execution Time = 18.8030 + 0.0156951(Cluster Ability) + 0.181922*(Number Of Threads) + 0.0412284*(Number of Nodes) – 0.0128650*(Cluster Ability)*(Number of Threads) – 0.0512727*(Number of Threads)*(Number of Nodes)*

4.2.1.5 Understanding the Model Constant

As with the RAIDb-1 model, the constant in the RAIDb-2 statistical regression model serves to calibrate the model to the machine speed of the computing nodes. The

constant used for the machines in this thesis is 18.8030. By changing the speed of the machine used to build the cluster, this constant will change. While some evidence shows that the relative performance changes between machine speeds is equal (Cecchet, 2003 and Rogers, 2004), the exact value for this constant, dependent on machine speed, can not be determined from this model. An estimate of this constant for different machine speeds can be obtained by running the TPC-H benchmark on a RAIDb-2 cluster and using the model to solve the equation for the missing constant. A more definite evaluation of the model based on machine speed should be obtained through further research.

4.2.1.6 Coefficient of Determination

The coefficient of determination can be used to describe the goodness-of-fit, or accuracy of the model applied to the data. The R^2 value for RAIDb-2 cluster was calculated to be 0.9826. This means that 98.26% of all the variation in the data is explained by the regression model stated above. Less than two percent of the data found in this research is unexplained. This high value for the coefficient of determination indicates an excellent goodness-of-fit, or accuracy, of the regression model.

An adjusted value for R^2 was also calculated for RAIDb-2 cluster results. The adjusted R^2 value for RAIDb-1 cluster is 0.9674, or 96.74%. This adjusted value explains that, more accurately, only 96.74% of the variation in the performance of the RAIDb-2 cluster is accounted for by the regression model. Less than four percent of the variation remains unexplained. While a slightly lower value, the adjusted coefficient of determination is an excellent indicator that the regression model can predict, to a high accuracy, the performance of a RAIDb-2 cluster.

4.2.2 ACID Compliance

As with RAIDb-1, the performance model of the RAIDb-2 clusters are only valid of ACID-compliance can be met. Each of the ACID tests specified in the TPC-H benchmark specifications were executed against the RAIDb-2 clusters. The results of these tests can be seen in Table 4.15.

Table 4.15. Results of TPC-H ACID tests for individual nodes and entire RAIDb-2 system.

System Tested	ACID Test	Pass/Fail	Conclusion
Node 1-6	Atomicity	Pass	PostgreSQL transaction implementation
	Consistency	Pass	PostgreSQL -implemented consistency-compliance
	Isolation	Pass	PostgreSQL -implemented isolation-compliance
	Durability	Pass	PostgreSQL hard disk storage
RAIDb-2 System	Atomicity	Pass	All PostgreSQL transactions correctly passed to all relevant backends for handling by database platform
	Consistency	Pass	PostgreSQL -implemented consistency-compliance correctly supported by RAIDb controller
	Isolation	Pass	PostgreSQL -implemented isolation-compliance correctly supported by RAIDb controller
	Durability	Pass	RAIDb controller correctly reports finished transactions made compliant through database platform storage mechanism

4.2.3 Statistical Model Validation

The statistical performance model presented in this thesis for RAIDb-2 clusters is mathematically valid for two- and six-node clusters. However, as with the RAIDb-1 model, in order to extend the validity of the model to larger clusters, some minimal

validation of the model must exist. To verify the accuracy of this model, results obtained from experiments with larger clusters must be statistically equal to the predicted values obtained from the model. Again, this is accomplished by employing the prediction interval of the model. If the results obtained from validation experiments fall within these bounds, the predicted and actual results are accepted as equal. Any variation from the predicted values that is within this prediction interval is credited to the random variation that exists in any system (Keller, 2001).

Using Minitab 14, prediction intervals were generated for eight- and ten-node RAIDb-2 clusters. Benchmark tests for the larger cluster sizes were then executed. These prediction intervals, with the results from the corresponding experiments, can be found in Table 4.16.

Table 4.16. Prediction intervals for 8- and 10-node RAIDb-2 clusters.

Cluster Configuration			**Prediction Interval**		**Actual Results**
Ability	*Threads*	*Nodes*	*Lower Bound*	*Upper Bound*	
2	12	8	14.8264	17.4062	16.5137
2	12	10	13.4655	16.4709	16.2793

For the eight-node RAIDb-2 clusters shown in Table 4.16, the prediction interval spans from 14.8264 to 17.4062. The experiment results produced an execution time of 16.5137, a value well within the prediction interval for the cluster. For the ten-node RAIDb-2 clusters shown in Table 4.16, the prediction interval spans from 13.4655 to 16.4709. The experiment results produced an execution time of 16.2793, also a value that is withing the prediction interval. The experiment results and the predicted values of the regression model are statistically equal. The model was successful at predicting the execution time of eight- and ten-node RAIDb-2 clusters. The predicted and actual execution times are graphed, between the prediction interval, in Figure 4.6.

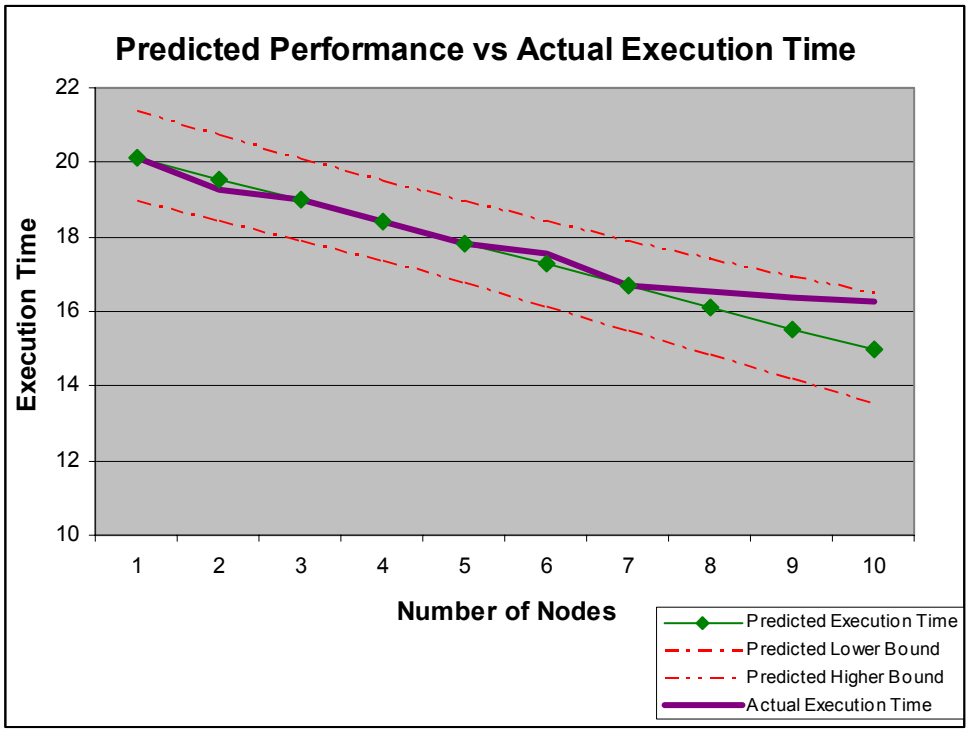

Figure 4.6. Predicted and actual execution times within the prediction interval for RAIDb-2 clusters.

Although statistically equal, execution times for ten-node systems begin to diverge from the predicted model. This behavior is representative of RAIDb-2 clusters. As shown in section 4.2.1.3, the performance enhancements obtained by increasing the cluster size decreases as the work presented to each node decreases. By adding more nodes to the cluster in the validation experiments, the workload per node has decreased. The resulting decrease in performance for eight- and ten-node clusters is shown in Figure 4.6.

Other possible reasons for this departure are many: the test client is beginning to bottleneck with server requests; the ability of the RAIDb-1 controller to enhance performance through load balancing, caching or other techniques is beginning to reach a limit; or, that the nature of the TPC-H benchmark is such that execution time can not be compressed any further. While the performance model is statistically valid for clusters

with up to ten nodes, further research should be conducted to explain the divergent behavior of a ten-node cluster.

4.2.4 RAIDb-2 Conclusions

An effective performance model for RAIDb-2 clusters has been created. This performance model was built using a three factor, two-level factorial design. The coefficients generated from each factor apply to the model with a 99% confidence interval. The correlation of determination shows that a 99% certainty exists that 95% of the variation in the data is explained by this performance model. This performance model can be used to predict the behavior of RAIDb-2 clusters containing two to six nodes. Although the model is statistically valid for clusters with more than six nodes, validation experiment results indicate that clusters with eight and ten nodes begin to diverge from the regression model. Possible causes of this divergence include the ability of the test client to sufficiently load the cluster or performance degradation of the RAIDb-1 controller.

Chapter 5

5 CONCLUSIONS AND RECOMMENDATIONS

5.1 Research Summary

By combining powerful, commercially-available Relational Database Management Systems (RDBMS) with the Clustered Java Database Connectivity (C-JDBC) software suite, significant increases in database performance can be obtained through the construction of Redundant Arrays of Inexpensive Databases (RAIDb). An RAIDb cluster consists of one master node, or controller, and one or more database nodes. The RAIDb controller coordinates the operations of each of these database nodes to create a single, powerful database system. Load balancing, query caching, and heterogeneity are just a few of the benefits provided by a RAIDb cluster.

The RAIDb clusters used in this research were homogenous. Cluster homogeneity allows for less unexplained variation and more specific regression analysis. Two levels of RAIDb clusters were evaluated. RAIDb level 1, or RAIDb-1, replicates full databases across multiple nodes. This type of replication is often referred to as full replication. RAIDb level 2, or RAIDb-2, is a partial replication solution that allows for partitioning and replicating large database across several nodes. This type of replication is often referred to as database striping or partial replication. RAIDb level 0, or RAIDb-0, partitions one database across multiple nodes and does not allow replication. Due to

the limitations of the C-JDBC implementation of the RAIDb standard, RAIDb-0 was not evaluated in this research.

A statistical experiment design was chosen to create an effective regression model of database clusters. A three factor, two-level factorial design was chosen to identify significant factors in the performance of a RAIDb cluster. Significant factors were then used to create a regression model to predict system performance. Other statistics, such as the coefficient of determination and the F-statistic, were used to verify the validity of the model.

There are many tools available to benchmark database systems. The most widely accepted database benchmarks are supported by the Transaction Processing Performance Council (TPC). Four TPC benchmarking suites were available for this research: TPC-C, TPC-H, TPC-R, and TPC-W. Of these four benchmarks, the TPC-H benchmark was chosen for benchmarking the RAIDb clusters. The broad range of evaluations available from TPC-H benchmark includes tests for data loading, serial query streams, parallel query streams, and data updates.

Using the overall execution time of TPC-H benchmark for varying RAIDb cluster configurations, a statistically meaningful performance model was created for each RAIDb level. Component tests of the TPC-H benchmark were also evaluated.

5.2 Conclusions

An analysis of the results produced from the statistical experiment design led to the following conclusions and performance models for RAIDb clusters:

- For RAIDb-1 clusters of two to six nodes, a performance model for predicting cluster performance can be expressed as *Execution Time = 18.5591 + 0.221999*(Number Of Threads) + 0.167329*(Number of Nodes) - 0.0803477*(Number of Threads)*(Number of Nodes).*
- For RAIDb-2 clusters of two to six nodes, a performance model for predicting cluster performance can be expressed as *Execution Time = 18.8030 + 0.0156951*(Cluster Ability) + 0.181922*(Number Of Threads) + 0.0412284*(Number of Nodes) - 0.0128650*(Cluster Ability)*(Number of Threads) - 0.0512727*(Number of Threads)*(Number of Nodes).*
- The number of nodes available to handle workload in all levels of RAIDb contributes the most performance enhancement. As more database nodes are added to the cluster, higher performance is possible.
- The number of threads, or workload, executed against all levels of RAIDb contributes the second-most performance enhancement. As more load given to the cluster, more efficient procedures can be used to enhance performance and faster execution time is available.
- The interaction between the number of threads and nodes for all levels of RAIDb contribute to the performance of the cluster. As the number of threads and nodes are jointly increased, higher system performance is obtainable.
- For RAIDb-2 clusters of two to six nodes, the ability, or amount of replicated information available per database node, also contributes to the performance of the cluster. As the ability of a cluster increases, execution time decreases.

- For RAIDb-2 clusters of two to six nodes, the interaction between the ability and workload of the cluster positively affects performance. As workload and ability simultaneously increase, greater performance enhancements are possible.
- For both RAIDb-1 and RAIDb-2 clusters, eight- and ten-node configurations begin to deviate from the model. It is estimated that this departure from the model results from the inabililty of the test client to sufficiently load the system or that a limit exists for performance enhancement.

5.3 Recommendations for Future Research

A number of factors known to affect the performance of C-JDBC clusters were not evaluated in this research. These factors include processor speed, processor type, storage medium and speed, database platform, and database size. No known research involving the effect of these factors on the performance of RAIDb clusters has been conducted and should be pursued.

While some validation for the models presented in this thesis has been established, a more complete evaluation of the model with higher node counts is required. There is some evidence that a ten-node cluster begins to approach a performance limit. As cluster size increases to ten nodes, a deviating trend can be noticed. Possible causes of this bottleneck include the inability of the test client to load the system or a limit of performance enhancement. The exact cause of this deviation, however, is unclear and additional research is required in this area.

Statistical analysis in this research also identified various factors that were only slightly insignificant to the regression model. These factors were found in the RAIDb-2

cluster configurations and include the interaction between cluster ability and available nodes and the three way interaction between cluster ability, thread number, and node count. The very low p-value for these interaction effects warrant further research to more fully understand how these interactions effect the performance of the cluster.

6. BIBLIOGRAPHY

Ault, M., Tummam M., (2003). *Oracle9i RAC: Oracle Real Application Clusters Configuration and Internals*, Rampant Press.

Austin, D., Thakkar M. (1999). *Using Oracle 8,* Indianapolis, IN: Macmillan Computer Publishing.

Balkanski, D., Trams, M., Rehm, W. (2003). Communication middleware systems for heterogeneous clusters: a comparative study. *Proceedings of the 2003 IEEE International Conference on Cluster Computing, 1*, 504- 507.

Brewer, E. (2004). *System R & DBMS Overview*. Retrieved November 22, 2004 from http://www.cs.berkeley.edu/~brewer/cs262/SystemR.html

Brown, S., Vallas, R., Ibrahim, M., Al-Zobaidie, A. (1998). Object design of a distributed client/server system. *Proceedings of the Ninth International Workshop on Database and Expert Systems Applications, 1*, 957-966.

Cecchet, E., Marguerite, J., Zwaenepoel, W. (2003). RAIDb: Redundant Array of Inexpensive Databases. *Institut National De Recherche En Informatique Et En Automatique: Rapport de recherché, 4921,* 1-25. Retrieved January 13, 2004 from http://c-jdbc.objectweb.org/current/doc/RR-C-JDBC.pdf

Cecchet, E., (2004). C-JDBC: General Mailing List. Message posted to ObjectWeb electronic mailing list, archived at http://www.objectweb.org/wws/arc/c-jdbc

Chen, C., Sun, W., Rishe, N. (1998). Performance Comparison of Three Alternatives of Distributed Multidatabase Systems: A Global Query Perspective. *IEEE International, 1,* 53-59.

Chen, C. H., Ting, W., Lu, W. B., Wang, G. K. (2003) Recovery Mechanism Design for Hot Standby Computer System. *IEEE International Conference on Systems, Man and Cybernetics, 3,* 3027-3031.

Codd, E. F. (1970). A Relational Model of Data for Large Shared Data Banks. *Communications of the ACM*, 13(6), 377-387.

Depue, A., (2004). C-JDBC: General Mailing List. Message posted to ObjectWeb electronic mailing list, archived at http://www.objectweb.org/wws/arc/c-jdbc

Deutsch, J.M. (1994). The evolution of customer middleware requirements. *Proceedings of the Third International Conference on Parallel and Distributed Information Systems, 1*, 262-263.

Emic Networks (2005). *Emic Application Clusters for MySQL*. Retrieved January 01, 2005 from http://www.emicnetworks.com/products/mysql.html

High Availability JDBC Group, (2004). *An Overview of HA-JDBC*. Retrieved December 30, 2004 from http://ha-jdbc.sourceforge.net/

Hemibigner, D., McLeod, D., (1985). A Federated Architecture for Information Management. *ACM Transactions on Office Information Systems, 3 (3)*, 253-278.

Hsiao, D. (1992). Federated databases and systems: part I - A tutorial on their data sharing. *The VLDB Journal: The International Journal on Very Large Data Bases, 1 (1)*, 127-180.

Keller, G. (2001). *Applied Statistics with Microsoft Excel,* Pacific Grove, CA: Duxbury – Thomson Learning.

Matisse (2004). *The Emergence of the Object-SQL Database,* Retrieved November 01, 2004 from http://matisse.com/pdf/product_information/collateral/the_emergence_of_the_object-sql_database.pdf.

Microsoft Corporation, (1999). *Glossary*. Retrieved November 01, 2004 from http://iishelp.web.cern.ch/IISHelp/iis/htm/core/iigloss.htm.

Montgomery, D. C., (2004). *Engineering Statistics*, New York, NY: John Wiley and Sons.

Nielsen, C. (2003). *A Descriptive Performance Model Of Small, Low Cost, Diskless Beowulf Clusters*. Provo, UT: Brigham Young University Press

Rogers, B. (2004). Performance Testing and Analysis of Redundant Arrays of Inexpensive Databases, *2004 Annual Conference of the American Society of Engineering Education, 1,* 1-9.

Sadoski, D. (2004). *Database Two Phase Commit*. Retrieved September 09, 2004 from http://www.sei.cmu.edu/str/descriptions/dtpc.html

Schilit, B.N., Sengupta, U., (2004). Device Ensembles. *IEEE Computer, 37(12),* 56-64.

Shen, C., Wei, J. Y. (1998). The Network as a Distributed Object Database. *Network Operations and Management Symposium, 2 (15-22),* 540-548.

Shiroshita, T., Takahashi, O., Yamashita, M., Nakamura, Y. (1996). Reliable data distribution middleware for large-scale massive data replication. *Fourth International Conference on Parallel and Distributed Information Systems, 1, 196-205.*

Takegaki, M. (1995). The middleware services for industrial computer systems. *Proceedings of the Second International Workshop on Real-Time Computing Systems and Applications,* 180-181.

Tanenbaum, A. (2002). *Distributed Systems: principles and paradigms*. Upper Saddle River, NJ: Prentice Hall.

Transaction Processing Performance Council. (2004). TPC Benchmark C Standard Specification. Retrieved September 25, 2004 from http://www.tpc.org/

Transaction Processing Performance Council. (2003). TPC Benchmark H Standard Specification. Retrieved September 25, 2004 from http://www.tpc.org/

Transaction Processing Performance Council. (2003). TPC Benchmark R Standard Specification. Retrieved September 25, 2004 from http://www.tpc.org/

Transaction Processing Performance Council. (). TPC Benchmark W Standard Specification. Retrieved September 25, 2004 from http://www.tpc.org/

Vaughn, J. (2003). *A Short Database History.* Retrieved October 12, 2004 from http://math.hws.edu/vaughn/cpsc/343/2003/history.html

APPENDIX

APPENDIX A – CLUSTER NETWORK CONFIGURATION

Node 1

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       node1
172.16.1.2      node2
172.16.1.3      node3
172.16.1.4      node4
172.16.1.5      node5
172.16.1.6      node6
172.16.0.10     controller1
172.16.0.11     controller2
```

Network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.1
NETWORK=172.16.0.0
BROADCAST=172.16.255.255
```

Node 2

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       node2
172.16.1.1      node1
172.16.1.3      node3
172.16.1.4      node4
172.16.1.5      node5
172.16.1.6      node6
172.16.0.10     controller1
172.16.0.11     controller2
```

Network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.2
NETWORK=172.16.0.0
BROADCAST=172.16.255.255
```

Node 3

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       node3
172.16.1.1      node1
172.16.1.2      node2
172.16.1.4      node4
172.16.1.5      node5
172.16.1.6      node6
```

172.16.0.10 controller1
172.16.0.11 controller2

Network configuration file

Please read /usr/share/doc/initscripts-*/sysconfig.txt
for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.3
NETWORK=172.16.0.0
BROADCAST=172.16.255.255

Node 4

Hosts file

Do not remove the following line, or various programs
that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost node4
172.16.1.1 node1
172.16.1.2 node2
172.16.1.3 node3
172.16.1.5 node5
172.16.1.6 node6
172.16.0.10 controller1
172.16.0.11 controller2

Network configuration file

Please read /usr/share/doc/initscripts-*/sysconfig.txt
for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.4
NETWORK=172.16.0.0
BROADCAST=172.16.255.255

Node 5

Hosts file

Do not remove the following line, or various programs
that require network functionality will fail.
127.0.0.1 localhost.localdomain localhost node5
172.16.1.1 node1
172.16.1.2 node2
172.16.1.3 node3
172.16.1.4 node4
172.16.1.6 node6
172.16.0.10 controller1
172.16.0.11 controller2

Network configuration file

Please read /usr/share/doc/initscripts-*/sysconfig.txt
for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.5
NETWORK=172.16.0.0
BROADCAST=172.16.255.255

Node 6

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       node6
172.16.1.1      node1
172.16.1.2      node2
172.16.1.3      node3
172.16.1.4      node4
172.16.1.5      node5
172.16.0.10     controller1
172.16.0.11     controller2
```

Network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

GATEWAY=172.168.0.10
TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.1.6
NETWORK=172.16.0.0
BROADCAST=172.16.255.255
```

Controller 1

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       controller1
172.16.1.1      node1
172.16.1.2      node2
```

172.16.1.3 node3
172.16.1.4 node4
172.16.1.5 node5
172.16.1.6 node6
172.16.0.11 controller2_internal
10.0.1.11 controller2_external
10.0.1.1 testclient

Internal network configuration file

Please read /usr/share/doc/initscripts-*/sysconfig.txt
for the documentation of these parameters.

TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.0.10
NETWORK=172.16.0.0
BROADCAST=172.16.255.255

External network configuration file

Please read /usr/share/doc/initscripts-*/sysconfig.txt
for the documentation of these parameters.

TYPE=Ethernet
DEVICE=eth1
BOOTPROTO=static
NETMASK=255.0.0.0
ONBOOT=yes
IPADDR=10.0.1.10
NETWORK=10.0.0.0
BROADCAST=10.255.255.255

Controller 2

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       controller2
172.16.1.1      node1
172.16.1.2      node2
172.16.1.3      node3
172.16.1.4      node4
172.16.1.5      node5
172.16.1.6      node6
172.16.0.10     controller1_internal
10.0.1.10       controller1_external
10.0.1.1        testclient
```

Internal network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.255.0.0
ONBOOT=yes
IPADDR=172.16.0.11
NETWORK=172.16.0.0
BROADCAST=172.16.255.255
```

External network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

TYPE=Ethernet
DEVICE=eth1
BOOTPROTO=static
NETMASK=255.0.0.0
ONBOOT=yes
IPADDR=10.0.1.11
NETWORK=10.0.0.0
BROADCAST=10.255.255.255
```

Test Client

Hosts file

```
# Do not remove the following line, or various programs
# that require network functionality will fail.
127.0.0.1       localhost.localdomain       localhost       testclient
10.0.1.10       controller1
10.0.1.11       controller2
```

Network configuration file

```
# Please read /usr/share/doc/initscripts-*/sysconfig.txt
# for the documentation of these parameters.

TYPE=Ethernet
DEVICE=eth0
BOOTPROTO=static
NETMASK=255.0.0.0
ONBOOT=yes
IPADDR=10.0.1.1
NETWORK=10.0.0.0
BROADCAST=10.255.255.255
```

APPENDIX B – POSTGRESQL CONFIGURATION FILES

Postgresql.conf

```
# -----------------------------
# PostgreSQL configuration file
# -----------------------------
#
# This file consists of lines of the form:
#
#   name = value
#
# (The '=' is optional.) White space may be used. Comments are introduced
# with '#' anywhere on a line. The complete list of option names and
# allowed values can be found in the PostgreSQL documentation. The
# commented-out settings shown in this file represent the default values.
#
# Any option can also be given as a command line switch to the
# postmaster, e.g. 'postmaster -c log_connections=on'. Some options
# can be changed at run-time with the 'SET' SQL command.
#
# This file is read on postmaster startup and when the postmaster
# receives a SIGHUP. If you edit the file on a running system, you have
# to SIGHUP the postmaster for the changes to take effect, or use
# "pg_ctl reload".


#---------------------------------------------------------------------------
# CONNECTIONS AND AUTHENTICATION
#---------------------------------------------------------------------------

# - Connection Settings -

tcpip_socket = true
max_connections = 100
          # note: increasing max_connections costs about 500 bytes of shared
          # memory per connection slot, in addition to costs from shared_buffers
       # and max_locks_per_transaction.
#superuser_reserved_connections = 2
#port = 5432
#unix_socket_directory = ''
```

```
#unix_socket_group = ''
#unix_socket_permissions = 0777     # octal
#virtual_host = ''                  # what interface to listen on; defaults to any
#rendezvous_name = ''                       # defaults to the computer name

# - Security & Authentication -

#authentication_timeout = 60 # 1-600, in seconds
#ssl = false
#password_encryption = true
#krb_server_keyfile = ''
#db_user_namespace = false


#---------------------------------------------------------------------------
# RESOURCE USAGE (except WAL)
#---------------------------------------------------------------------------

# - Memory -

shared_buffers = 1000                       # min 16, at least max_connections*2, 8KB each
#sort_mem = 4096                # min 64, size in KB
#vacuum_mem = 8192                  # min 1024, size in KB

# - Free Space Map -

#max_fsm_pages = 20000              # min max_fsm_relations*16, 6 bytes each
#max_fsm_relations = 1000   # min 100, ~50 bytes each

# - Kernel Resource Usage -

#max_files_per_process = 1000       # min 25
#preload_libraries = ''


#---------------------------------------------------------------------------
# WRITE AHEAD LOG
#---------------------------------------------------------------------------

# - Settings -

fsync = false                       # turns forced synchronization on or off
#wal_sync_method = fsync    # the default varies across platforms:
                                    # fsync, fdatasync, open_sync, or open_datasync
#wal_buffers = 50                   # min 4, 8KB each
```

```
# - Checkpoints -

checkpoint_segments = 15     # in logfile segments, min 1, 16MB each
checkpoint_timeout = 300     # range 30-3600, in seconds
#checkpoint_warning = 30     # 0 is off, in seconds
#commit_delay = 0            # range 0-100000, in microseconds
#commit_siblings = 5         # range 1-1000


#---------------------------------------------------------------------------
# QUERY TUNING
#---------------------------------------------------------------------------

# - Planner Method Enabling -

#enable_hashagg = false
#enable_hashjoin = true
#enable_indexscan = true
#enable_mergejoin = true
#enable_nestloop = true
#enable_seqscan = true
#enable_sort = true
#enable_tidscan = true

# - Planner Cost Constants -

#effective_cache_size = 1000# typically 8KB each
#random_page_cost = 4                # units are one sequential page fetch cost
#cpu_tuple_cost = 0.01               # (same)
#cpu_index_tuple_cost = 0.001        # (same)
#cpu_operator_cost = 0.0025 # (same)

# - Genetic Query Optimizer -

geqo = true
#geqo_threshold = 11
#geqo_effort = 1
#geqo_generations = 0
#geqo_pool_size = 0          # default based on tables in statement,
                             # range 128-1024
#geqo_selection_bias = 2.0   # range 1.5-2.0

# - Other Planner Options -

#default_statistics_target = 10         # range 1-1000
#from_collapse_limit = 8
```

```
#join_collapse_limit = 8        # 1 disables collapsing of explicit JOINs


#---------------------------------------------------------------------------
# ERROR REPORTING AND LOGGING
#---------------------------------------------------------------------------

# - Syslog -

syslog = 2                      # range 0-2; 0=stdout; 1=both; 2=syslog
#syslog_facility = 'LOCAL0'
#syslog_ident = 'postgres'

# - When to Log -

#client_min_messages = notice        # Values, in order of decreasing detail:
                                    #   debug5, debug4, debug3, debug2, debug1,
                                    #   log, info, notice, warning, error

#log_min_messages = notice # Values, in order of decreasing detail:
                                    #   debug5, debug4, debug3, debug2, debug1,
                                    #   info, notice, warning, error, log, fatal,
                                    #   panic

#log_error_verbosity = default   # terse, default, or verbose messages

#log_min_error_statement = panic # Values in order of increasing severity:
                                    #   debug5, debug4, debug3, debug2, debug1,
                                    #   info, notice, warning, error, panic(off)

#log_min_duration_statement = -1 # Log all statements whose
                                    # execution time exceeds the value, in
                                    # milliseconds.  Zero prints all queries.
                                    # Minus-one disables.

#silent_mode = false             # DO NOT USE without Syslog!

# - What to Log -

#debug_print_parse = false
#debug_print_rewritten = false
#debug_print_plan = false
#debug_pretty_print = false
#log_connections = false
#log_duration = false
#log_pid = false
```

```
#log_statement = false
#log_timestamp = false
#log_hostname = false
#log_source_port = false


#------------------------------------------------------------------------
# RUNTIME STATISTICS
#------------------------------------------------------------------------

# - Statistics Monitoring -

#log_parser_stats = false
#log_planner_stats = false
#log_executor_stats = false
#log_statement_stats = false

# - Query/Index Statistics Collector -

#stats_start_collector = true
#stats_command_string = false
#stats_block_level = false
#stats_row_level = false
#stats_reset_on_server_start = true


#------------------------------------------------------------------------
# CLIENT CONNECTION DEFAULTS
#------------------------------------------------------------------------

# - Statement Behavior -

#search_path = '$user,public'  # schema names
#check_function_bodies = true
#default_transaction_isolation = 'read committed'
#default_transaction_read_only = false
#statement_timeout = 0                  # 0 is disabled, in milliseconds

# - Locale and Formatting -

#datestyle = 'iso, mdy'
#timezone = unknown             # actually, defaults to TZ environment setting
#australian_timezones = false
#extra_float_digits = 0                 # min -15, max 2
#client_encoding = sql_ascii  # actually, defaults to database encoding
```

```
# These settings are initialized by initdb -- they may be changed
lc_messages = 'en_US.UTF-8'                     # locale for system error message strings
lc_monetary = 'en_US.UTF-8'                     # locale for monetary formatting
lc_numeric = 'en_US.UTF-8'              # locale for number formatting
lc_time = 'en_US.UTF-8'                         # locale for time formatting

# - Other Defaults -

#explain_pretty_print = true
#dynamic_library_path = '$libdir'
#max_expr_depth = 10000                 # min 10


#---------------------------------------------------------------------------
# LOCK MANAGEMENT
#---------------------------------------------------------------------------

#deadlock_timeout = 1000     # in milliseconds
#max_locks_per_transaction = 64      # min 10, ~260*max_connections bytes each


#---------------------------------------------------------------------------
# VERSION/PLATFORM COMPATIBILITY
#---------------------------------------------------------------------------

# - Previous Postgres Versions -

#add_missing_from = true
#regex_flavor = advanced      # advanced, extended, or basic
#sql_inheritance = true

# - Other Platforms & Clients -

#transform_null_equals = false
```

pg_hba.conf

```
# PostgreSQL Client Authentication Configuration File
# ===================================================
#
# Refer to the PostgreSQL Administrator's Guide, chapter "Client
# Authentication" for a complete description.  A short synopsis
# follows.
#
# This file controls: which hosts are allowed to connect, how clients
```

```
# are authenticated, which PostgreSQL user names they can use, which
# databases they can access.  Records take one of seven forms:
#
# local      DATABASE  USER  METHOD  [OPTION]
# host       DATABASE  USER  IP-ADDRESS  IP-MASK   METHOD  [OPTION]
# hostssl    DATABASE  USER  IP-ADDRESS  IP-MASK   METHOD  [OPTION]
# hostnossl  DATABASE  USER  IP-ADDRESS  IP-MASK   METHOD  [OPTION]
# host       DATABASE  USER  IP-ADDRESS/CIDR-MASK  METHOD  [OPTION]
# hostssl    DATABASE  USER  IP-ADDRESS/CIDR-MASK  METHOD  [OPTION]
# hostnossl  DATABASE  USER  IP-ADDRESS/CIDR-MASK  METHOD  [OPTION]
#
# (The uppercase quantities should be replaced by actual values.)
# The first field is the connection type: "local" is a Unix-domain socket,
# "host" is either a plain or SSL-encrypted TCP/IP socket, "hostssl" is an
# SSL-encrypted TCP/IP socket, and "hostnossl" is a plain TCP/IP socket.
# DATABASE can be "all", "sameuser", "samegroup", a database name (or
# a comma-separated list thereof), or a file name prefixed with "@".
# USER can be "all", an actual user name or a group name prefixed with
# "+" or a list containing either.  IP-ADDRESS and IP-MASK specify the
# set of hosts the record matches.  CIDR-MASK is an integer between 0
# and 32 (IPv6) or 128(IPv6) inclusive, that specifies the number of
# significant bits in the mask, so an IPv4 CIDR-MASK of 8 is equivalent
# to an IP-MASK of 255.0.0.0, and an IPv6 CIDR-MASK of 64 is equivalent
# to an IP-MASK of ffff:ffff:ffff:ffff::. METHOD can be "trust", "reject",
# "md5", "crypt", "password", "krb4", "krb5", "ident", or "pam".  Note
# that "password" uses clear-text passwords; "md5" is preferred for
# encrypted passwords.  OPTION is the ident map or the name of the PAM
# service.
#
# This file is read on server startup and when the postmaster receives
# a SIGHUP signal.  If you edit the file on a running system, you have
# to SIGHUP the postmaster for the changes to take effect, or use
# "pg_ctl reload".

# Put your actual configuration here
# ----------------------------------
#
# CAUTION: The default configuration allows any local user to connect
# using any PostgreSQL user name, including the superuser, over either
# Unix-domain sockets or TCP/IP.  If you are on a multiple-user
# machine, the default configuration is probably too liberal for you.
# Change it to use something other than "trust" authentication.
#
# If you want to allow non-local connections, you need to add more
# "host" records.  Also, remember TCP/IP connections are only enabled
# if you enable "tcpip_socket" in postgresql.conf.
```

```
# TYPE  DATABASE    USER        IP-ADDRESS        IP-MASK           METHOD

# IPv4-style local connections:
host    all         all         127.0.0.1         255.255.255.255   trust
host    all         all         172.16.0.1        255.255.0.0       trust
# IPv6-style local connections:
#host    all         all         ::1               ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff       trust

# Using sockets credentials for improved security. Not available everywhere,
# but works on Linux, *BSD (and probably some others)

#local  all         all         trust
local  all    all           password
```

refreshPG.bash

```
#!/bin/bash

# (c) 2004 Brandon Rogers

# This file automates the restoring of the database image after a TPC-H run
# 1. Stop the database service
# 2. Clean up old files
# 3. Extract the archived database image
# 4. Start the database service
# 5. Create the tables needed for the TPC-H Benchmark

# The tables in step 5 will be removed and recreated as part of the load test.

service postgresql stop
rm -Rf pgsql/
tar -xvzf pgsql.tar.gz
service postgresql start
psql -u < createTables.sql
echo Done!
```

createTables.sql

```
create table region ( r_regionkey INTEGER NOT NULL, r_name CHAR(25),
r_comment VARCHAR(152), PRIMARY KEY(r_regionkey)) ;
create table nation ( n_nationkey INTEGER NOT NULL, n_name CHAR(25),
n_regionkey INTEGER , n_comment VARCHAR(152), PRIMARY KEY(n_nationkey))
;
```

```
create table part ( p_partkey INTEGER NOT NULL, p_name VARCHAR(55), p_mfgr
CHAR(25), p_brand CHAR(10), p_type VARCHAR(25), p_size INTEGER, p_container
char(10), p_retailprice DECIMAL(15,2), p_comment VARCHAR(23), PRIMARY
KEY(p_partkey)) ;

create table supplier ( s_suppkey INTEGER NOT NULL, s_name CHAR(25), s_address
VARCHAR(40), s_nationkey INTEGER , s_phone CHAR(15), s_acctbal
DECIMAL(15,2), s_comment VARCHAR(101), PRIMARY KEY(s_suppkey)) ;

create table partsupp ( ps_partkey INTEGER NOT NULL, ps_suppkey INTEGER NOT
NULL, ps_availqty INTEGER, ps_supplycost DECIMAL(15,2), ps_comment
VARCHAR(199), PRIMARY KEY (ps_partkey, ps_suppkey)) ;

create table customer ( c_custkey INTEGER NOT NULL, c_name VARCHAR(25),
c_address VARCHAR(40), c_nationkey INTEGER , c_phone CHAR(15), c_acctbal
DECIMAL(15,2), c_mktsegment CHAR(10), c_comment VARCHAR(117), PRIMARY
KEY(c_custkey)) ;

create table orders ( o_orderkey INTEGER NOT NULL, o_custkey INTEGER ,
o_orderstatus char(1), o_totalprice DECIMAL(15,2), o_orderdate DATE, o_orderpriority
CHAR(15), o_clerk CHAR(15), o_shippriority INTEGER, o_comment VARCHAR(79),
PRIMARY KEY(o_orderkey)) ;

create table lineitem ( l_orderkey INTEGER NOT NULL, l_partkey INTEGER ,
l_suppkey INTEGER , l_linenumber INTEGER NOT NULL, l_quantity
DECIMAL(15,2), l_extendedprice DECIMAL(15,2), l_discount DECIMAL(15,2), l_tax
DECIMAL(15,2), l_returnflag CHAR(1), l_linestatus CHAR(1), l_shipdate DATE,
l_commitdate DATE, l_receiptdate DATE, l_shipinstruct CHAR(25), l_shipmode
CHAR(10), l_comment VARCHAR(44), PRIMARY KEY (l_orderkey, l_linenumber)) ;

create table orders_temp ( o_orderkey INTEGER NOT NULL, o_custkey INTEGER ,
o_orderstatus char(1), o_totalprice DECIMAL(15,2), o_orderdate DATE, o_orderpriority
CHAR(15), o_clerk CHAR(15), o_shippriority INTEGER, o_comment VARCHAR(79),
PRIMARY KEY(o_orderkey)) ;

create table lineitem_temp ( l_orderkey INTEGER NOT NULL, l_partkey INTEGER ,
l_suppkey INTEGER , l_linenumber INTEGER NOT NULL, l_quantity
DECIMAL(15,2), l_extendedprice DECIMAL(15,2), l_discount DECIMAL(15,2), l_tax
DECIMAL(15,2), l_returnflag CHAR(1), l_linestatus CHAR(1), l_shipdate DATE,
l_commitdate DATE, l_receiptdate DATE, l_shipinstruct CHAR(25), l_shipmode
CHAR(10), l_comment VARCHAR(44), PRIMARY KEY (l_orderkey, l_linenumber)) ;
```

APPENDIX C – C-JDBC CONTROLLER CONFIGURATION FILES

RAIDb-1 Configuration Files

One Controller-One Node Controller Configuration File

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
  <Controller port="25322">

    <JmxSettings>
      <RmiJmxAdaptor/>
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc/setup/raidb1.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

One Controller-One Node Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC 1.0rc6//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-1.0rc6.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
```

```
    <VirtualUsers>
        <VirtualLogin vLogin="tpch" vPassword="tpch" />
    </VirtualUsers>
  </AuthenticationManager>


  <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <RequestManager>
    <RequestScheduler>
      <RAIDb-1Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-1>
        <WaitForCompletion policy="first"/>
        <RAIDb-1-LeastPendingRequestsFirst />
      </RAIDb-1>
    </LoadBalancer>

  </RequestManager>


 </VirtualDatabase>

</C-JDBC>
```

One Controller-Six Node Controller Configuration File

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
```

```
  <Controller port="25322">

    <JmxSettings>
      <RmiJmxAdaptor/>
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc/setup/raidb1-6.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

One Controller-Six Node Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC 1.0rc6//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-1.0rc6.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>


    <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">
      <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
        <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
      </ConnectionManager>
    </DatabaseBackend>

    <DatabaseBackend name="node1" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.1/tpch" connectionTestStatement="select 1">
      <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
        <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
```

```
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node2" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.2/tpch" connectionTestStatement="select 1">
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node4" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.4/tpch" connectionTestStatement="select 1">
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node5" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.5/tpch" connectionTestStatement="select 1">
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node6" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.6/tpch" connectionTestStatement="select 1">
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>


  <RequestManager>
    <RequestScheduler>
      <RAIDb-1Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>
```

```
    <LoadBalancer>
      <RAIDb-1>
        <WaitForCompletion policy="first"/>
        <RAIDb-1-LeastPendingRequestsFirst />
      </RAIDb-1>
    </LoadBalancer>

  </RequestManager>


 </VirtualDatabase>

</C-JDBC>
```

Two Controller-One Node Controller Configuration File (Controller 1)

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
 <Controller port="25322">

   <JmxSettings>
    <RmiJmxAdaptor port="1090" />
   </JmxSettings>

   <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/setup2/raidb1-1.xml"/>
 </Controller>
</C-JDBC-CONTROLLER>
```

Two Controller-Six Node Virtual Database Configuration File (Controller 1)

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC 1.0rc6//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-1.0rc6.dtd">

<C-JDBC>

 <VirtualDatabase name="tpch">
     <Distribution>
      <BackendRecoveryPolicy backendName="node4" recoveryPolicy="on"/>
     </Distribution>
```

```
    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>


    <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">
      <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
        <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
      </ConnectionManager>
    </DatabaseBackend>


    <RequestManager>
      <RequestScheduler>
        <RAIDb-1Scheduler level="pessimisticTransaction"/>
      </RequestScheduler>

      <RequestCache>
        <MetadataCache/>
        <ResultCache granularity="table" />
      </RequestCache>

      <LoadBalancer>
        <RAIDb-1>
          <WaitForCompletion policy="first"/>
          <RAIDb-1-LeastPendingRequestsFirst />
        </RAIDb-1>
      </LoadBalancer>

    </RequestManager>


  </VirtualDatabase>

</C-JDBC>
```

Two Controller-One Node Controller Configuration File (Controller 2)

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
  <Controller port="25323">

    <JmxSettings>
      <RmiJmxAdaptor port="1091" />
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/setup2/raidb1-2.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

Two Controller-Six Node Virtual Database Configuration File (Controller 2)

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC 1.0rc6//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-1.0rc6.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">

       <Distribution>
         <BackendRecoveryPolicy backendName="node3" recoveryPolicy="on"/>
          </Distribution>


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>

    <DatabaseBackend name="node4" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.4/tpch" connectionTestStatement="select 1">
      <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
```

```
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <RequestManager>
    <RequestScheduler>
      <RAIDb-1Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-1>
        <WaitForCompletion policy="first"/>
        <RAIDb-1-LeastPendingRequestsFirst />
      </RAIDb-1>
    </LoadBalancer>

  </RequestManager>


 </VirtualDatabase>

</C-JDBC>
```

RAIDb-2 Configuration Files

One Node-High Ability Controller Configuration File

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
 <Controller port="25322">

  <JmxSettings>
    <RmiJmxAdaptor/>
  </JmxSettings>
```

```
   <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/level2/hiAbility/raidb2.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

One Node-High Ability Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC latest//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-latest.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>


    <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">

            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
```

```
                    <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="customer" nbOfColumns="8">
                    <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                    <DatabaseColumn columnName="c_name" isUnique="false"/>
                    <DatabaseColumn columnName="c_address"
isUnique="false"/>
                    <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="c_phone" isUnique="false"/>
                    <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                    <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                </DatabaseTable>

                    <DatabaseTable tableName="part" nbOfColumns="9">
                    <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="p_name" isUnique="false"/>
                  <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                    <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                  <DatabaseColumn columnName="p_type" isUnique="false"/>
                    <DatabaseColumn columnName="p_size" isUnique="false"/>
                  <DatabaseColumn columnName="p_container"
isUnique="false"/>
                 <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

              <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                  <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                  <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                  </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
```

```
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                <DatabaseColumn columnName="l_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="revenue0" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue1" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue2" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue3" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue4" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue5" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue6" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue7" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue8" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue9" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue10" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue11" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue12" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue13" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue14" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue15" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue16" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue17" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue18" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue19" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue20" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue21" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue22" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue23" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue24" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue25" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>


            </DatabaseStaticSchema>
        </DatabaseSchema>
```

```
    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node2" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.2/tpch" connectionTestStatement="select 1">
            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer" nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name" isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone" isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>
```

```
                    <DatabaseTable tableName="part" nbOfColumns="9">
                        <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="p_name" isUnique="false"/>
                        <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                        <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                        <DatabaseColumn columnName="p_type" isUnique="false"/>
                        <DatabaseColumn columnName="p_size" isUnique="false"/>
                        <DatabaseColumn columnName="p_container"
isUnique="false"/>
                        <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                        <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="partsupp" nbOfColumns="5">
                        <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                        <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                        <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                        <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                    <DatabaseTable tableName="nation" nbOfColumns="4">
                        <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                        <DatabaseColumn columnName="n_name" isUnique="false"/>
                        <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                        <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="orders" nbOfColumns="9">
                        <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                        <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                        <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>
```

```
                    <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                        <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                        <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                        <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                        <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                        <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                        <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                        <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                        <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                        <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                        <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                        <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                        <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                        <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                        <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                        <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                        <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                        <DatabaseColumn columnName="l_tax" isUnique="false"/>
                        <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                        <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
```

```
                        <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                        <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                    </DatabaseTable>

                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>


  <RequestManager>
    <RequestScheduler>
      <RAIDb-2Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-2>
        <WaitForCompletion policy="first"/>
        <RAIDb-2-LeastPendingRequestsFirst />
      </RAIDb-2>
    </LoadBalancer>

  </RequestManager>


 </VirtualDatabase>

</C-JDBC>
```

Six Node-High Ability Controller Configuration File

```
<?xml version="1.0" encoding="ISO-8859-1" ?>
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER 1.0rc6//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
1.0rc6.dtd">
<C-JDBC-CONTROLLER>
  <Controller port="25322">

    <JmxSettings>
      <RmiJmxAdaptor/>
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/level2/hiAbility/raidb2-6.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

Six Node-High Ability Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC latest//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-latest.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>


    <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">

            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
```

```
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer" nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name" isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone" isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>

                        <DatabaseTable tableName="part" nbOfColumns="9">
                      <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="p_name" isUnique="false"/>
                    <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                        <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                    <DatabaseColumn columnName="p_type" isUnique="false"/>
                        <DatabaseColumn columnName="p_size" isUnique="false"/>
                        <DatabaseColumn columnName="p_container"
isUnique="false"/>
                     <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
```

```
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
```

```
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="lineitem" nbOfColumns="16">
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                <DatabaseColumn columnName="l_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
```

```
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="revenue0" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue1" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue2" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue3" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue4" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue5" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue6" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue7" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue8" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue9" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue10" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue11" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue12" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue13" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue14" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue15" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue16" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue17" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue18" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue19" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue20" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue21" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue22" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue23" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue24" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue25" nbOfColumns="2">
```

```
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>



                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node2" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.2/tpch" connectionTestStatement="select 1">
            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer" nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name" isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="c_phone" isUnique="false"/>
                    <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                    <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                </DatabaseTable>

                    <DatabaseTable tableName="part" nbOfColumns="9">
                    <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="p_name" isUnique="false"/>
                    <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                    <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                    <DatabaseColumn columnName="p_type" isUnique="false"/>
                    <DatabaseColumn columnName="p_size" isUnique="false"/>
                    <DatabaseColumn columnName="p_container"
isUnique="false"/>
                    <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
```

```
                        <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                        <DatabaseColumn columnName="l_tax" isUnique="false"/>
                        <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                        <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                        <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                    </DatabaseTable>

                      </DatabaseStaticSchema>
                </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node1" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.1/tpch" connectionTestStatement="select 1">

          <DatabaseSchema dynamicPrecision="static">
              <DatabaseStaticSchema>
                  <DatabaseProcedure name="loadtables" />
                  <DatabaseProcedure name="cleanup" />

                  <DatabaseTable tableName="supplier" nbOfColumns="7">
                      <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                      <DatabaseColumn columnName="s_name" isUnique="false"/>
                      <DatabaseColumn columnName="s_address"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="s_phone" isUnique="false"/>
                    <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="customer" nbOfColumns="8">
                    <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                    <DatabaseColumn columnName="c_name" isUnique="false"/>
                    <DatabaseColumn columnName="c_address"
isUnique="false"/>
                    <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="c_phone" isUnique="false"/>
                    <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                    <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                </DatabaseTable>

                    <DatabaseTable tableName="part" nbOfColumns="9">
                    <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="p_name" isUnique="false"/>
                  <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                    <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                  <DatabaseColumn columnName="p_type" isUnique="false"/>
                    <DatabaseColumn columnName="p_size" isUnique="false"/>
                    <DatabaseColumn columnName="p_container"
isUnique="false"/>
                 <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

              <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                  <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
```

```
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
```

```
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
```

```
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

            </DatabaseStaticSchema>
        </DatabaseSchema>

    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node4" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.4/tpch" connectionTestStatement="select 1">

        <DatabaseSchema dynamicPrecision="static">
```

```
<DatabaseStaticSchema>
    <DatabaseProcedure name="loadtables" />
    <DatabaseProcedure name="cleanup" />

    <DatabaseTable tableName="customer" nbOfColumns="8">
        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
        <DatabaseColumn columnName="c_name" isUnique="false"/>
        <DatabaseColumn columnName="c_address"
isUnique="false"/>
        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
        <DatabaseColumn columnName="c_phone" isUnique="false"/>
        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
    </DatabaseTable>

        <DatabaseTable tableName="part" nbOfColumns="9">
      <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
      <DatabaseColumn columnName="p_name" isUnique="false"/>
    <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
      <DatabaseColumn columnName="p_brand"
isUnique="false"/>
    <DatabaseColumn columnName="p_type" isUnique="false"/>
      <DatabaseColumn columnName="p_size" isUnique="false"/>
     <DatabaseColumn columnName="p_container"
isUnique="false"/>
   <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
      <DatabaseColumn columnName="p_comment"
isUnique="false"/>
    </DatabaseTable>

  <DatabaseTable tableName="partsupp" nbOfColumns="5">
      <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
    <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
      <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
     <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
```

```
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                <DatabaseColumn columnName="l_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>


                </DatabaseStaticSchema>
            </DatabaseSchema>

    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
```

```
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node5" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.5/tpch" connectionTestStatement="select 1">

          <DatabaseSchema dynamicPrecision="static">
              <DatabaseStaticSchema>
                  <DatabaseProcedure name="loadtables" />
                  <DatabaseProcedure name="cleanup" />

                  <DatabaseTable tableName="supplier" nbOfColumns="7">
                      <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                      <DatabaseColumn columnName="s_name" isUnique="false"/>
                      <DatabaseColumn columnName="s_address"
isUnique="false"/>
                      <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                      <DatabaseColumn columnName="s_phone" isUnique="false"/>
                      <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                      <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                  </DatabaseTable>

                      <DatabaseTable tableName="part" nbOfColumns="9">
                    <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                      <DatabaseColumn columnName="p_name" isUnique="false"/>
                   <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                      <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                   <DatabaseColumn columnName="p_type" isUnique="false"/>
                      <DatabaseColumn columnName="p_size" isUnique="false"/>
                      <DatabaseColumn columnName="p_container"
isUnique="false"/>
                    <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                      <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                  </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
```

```
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                </DatabaseTable>


            <DatabaseTable tableName="region" nbOfColumns="3">
                <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                <DatabaseColumn columnName="r_name" isUnique="false"/>
                <DatabaseColumn columnName="r_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="nation" nbOfColumns="4">
                <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                <DatabaseColumn columnName="n_name" isUnique="false"/>
                <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                <DatabaseColumn columnName="n_comment"
isUnique="false"/>
            </DatabaseTable>
            <DatabaseTable tableName="orders" nbOfColumns="9">
                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>
```

```
                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node6" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.6/tpch" connectionTestStatement="select 1">

        <DatabaseSchema dynamicPrecision="static">
            <DatabaseStaticSchema>
                <DatabaseProcedure name="loadtables" />
                <DatabaseProcedure name="cleanup" />

                <DatabaseTable tableName="supplier" nbOfColumns="7">
                    <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                    <DatabaseColumn columnName="s_name" isUnique="false"/>
                    <DatabaseColumn columnName="s_address"
isUnique="false"/>
                    <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="s_phone" isUnique="false"/>
                    <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                    <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="customer" nbOfColumns="8">
                    <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                    <DatabaseColumn columnName="c_name" isUnique="false"/>
                    <DatabaseColumn columnName="c_address"
isUnique="false"/>
                    <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                    <DatabaseColumn columnName="c_phone" isUnique="false"/>
                    <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                    <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
```

```
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="lineitem" nbOfColumns="16">
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                <DatabaseColumn columnName="l_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="orders_temp" nbOfColumns="9">
```

```
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
```

```
                        <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                        <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                    </DatabaseTable>

                  </DatabaseStaticSchema>
              </DatabaseSchema>

    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
   </DatabaseBackend>

   <RequestManager>
    <RequestScheduler>
      <RAIDb-2Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-2>
        <WaitForCompletion policy="first"/>
        <RAIDb-2-LeastPendingRequestsFirst />
      </RAIDb-2>
    </LoadBalancer>

   </RequestManager>


  </VirtualDatabase>

</C-JDBC>
```

One Node-LowAbility Controller Configuration File

```
<?xml version="1.0" encoding="UTF8" ?>
```

```
<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER latest//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
latest.dtd">

<C-JDBC-CONTROLLER>
  <Controller port="25322">

    <Report />

    <JmxSettings>
      <RmiJmxAdaptor/>
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/level2/lowAbility/raidb2.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

One Node-Low Ability Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC latest//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-latest.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
          <VirtualLogin vLogin="tpch" vPassword="tpch" />
      </VirtualUsers>
    </AuthenticationManager>


    <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">

            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
```

```
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer" nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name" isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone" isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>

                        <DatabaseTable tableName="part" nbOfColumns="9">
                      <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="p_name" isUnique="false"/>
                    <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                        <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                    <DatabaseColumn columnName="p_type" isUnique="false"/>
                        <DatabaseColumn columnName="p_size" isUnique="false"/>
                        <DatabaseColumn columnName="p_container"
isUnique="false"/>
                     <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
```

```
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
```

```
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
            </DatabaseTable>

            <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                <DatabaseColumn columnName="l_tax" isUnique="false"/>
                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="revenue0" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue1" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue2" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue3" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue4" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue5" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue6" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue7" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue8" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue9" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue10" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue11" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue12" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue13" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue14" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue15" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue16" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue17" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue18" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue19" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue20" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue21" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue22" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue23" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue24" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue25" nbOfColumns="2">
```

```
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>



                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node2" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.2/tpch" connectionTestStatement="select 1">
                <DatabaseSchema dynamicPrecision="static">
                    <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                        <DatabaseTable tableName="supplier"
nbOfColumns="7">
                            <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                            <DatabaseColumn columnName="s_name"
isUnique="false"/>
                            <DatabaseColumn columnName="s_address"
isUnique="false"/>
                            <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                            <DatabaseColumn columnName="s_phone"
isUnique="false"/>
                            <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                            <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                        </DatabaseTable>

                        <DatabaseTable tableName="customer"
nbOfColumns="8">
```

```
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name"
isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone"
isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>

                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <RequestManager>
    <RequestScheduler>
      <RAIDb-2Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-2>
        <WaitForCompletion policy="first"/>
        <RAIDb-2-LeastPendingRequestsFirst />
      </RAIDb-2>
    </LoadBalancer>

  </RequestManager>
```

```
  </VirtualDatabase>

</C-JDBC>
```

Six Node-Low Ability Controller Configuration File

```
<?xml version="1.0" encoding="UTF8" ?>

<!DOCTYPE C-JDBC-CONTROLLER PUBLIC "-//ObjectWeb//DTD C-JDBC-
CONTROLLER latest//EN"  "http://c-jdbc.objectweb.org/dtds/c-jdbc-controller-
latest.dtd">

<C-JDBC-CONTROLLER>
  <Controller port="25322">

    <Report />

    <JmxSettings>
      <RmiJmxAdaptor/>
    </JmxSettings>

    <VirtualDatabase virtualName="tpch" autoEnableBackends="true"
configFile="/usr/local/cjdbc1.0/level2/lowAbility/raidb2-6.xml"/>
  </Controller>
</C-JDBC-CONTROLLER>
```

Six Node-Low Ability Virtual Database Configuration File

```
<?xml version="1.0" encoding="UTF8"?>
<!DOCTYPE C-JDBC PUBLIC "-//ObjectWeb//DTD C-JDBC latest//EN" "http://c-
jdbc.objectweb.org/dtds/c-jdbc-latest.dtd">

<C-JDBC>

  <VirtualDatabase name="tpch">


    <AuthenticationManager>
      <Admin>
          <User username="admin" password="" />
      </Admin>
      <VirtualUsers>
```

```
        <VirtualLogin vLogin="tpch" vPassword="tpch" />
    </VirtualUsers>
   </AuthenticationManager>


   <DatabaseBackend name="node3" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.3/tpch" connectionTestStatement="select 1">

            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="supplier" nbOfColumns="7">
                        <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="s_name" isUnique="false"/>
                        <DatabaseColumn columnName="s_address"
isUnique="false"/>
                        <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="s_phone" isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer" nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name" isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone" isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>

                        <DatabaseTable tableName="part" nbOfColumns="9">
```

```
                    <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="p_name" isUnique="false"/>
                    <DatabaseColumn columnName="p_mfgr" isUnique="false"/>
                    <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                    <DatabaseColumn columnName="p_type" isUnique="false"/>
                    <DatabaseColumn columnName="p_size" isUnique="false"/>
                    <DatabaseColumn columnName="p_container"
isUnique="false"/>
                    <DatabaseColumn columnName="p_retailprice" isUnique="false"/>
                    <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="partsupp" nbOfColumns="5">
                    <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                    <DatabaseColumn columnName="ps_suppkey" isUnique="true"/>
                    <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                    <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                <DatabaseTable tableName="region" nbOfColumns="3">
                    <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                    <DatabaseColumn columnName="r_name" isUnique="false"/>
                    <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="nation" nbOfColumns="4">
                    <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                    <DatabaseColumn columnName="n_name" isUnique="false"/>
                    <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                    <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                </DatabaseTable>
                <DatabaseTable tableName="orders" nbOfColumns="9">
```

```
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="orders_temp" nbOfColumns="9">
                    <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk" isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="lineitem_temp" nbOfColumns="16">
                    <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                    <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                    <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                    <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                    <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                    <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                    <DatabaseColumn columnName="l_tax" isUnique="false"/>
                    <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
```

```
                    <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                    <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                    <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                </DatabaseTable>

                <DatabaseTable tableName="revenue0" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue1" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue2" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue3" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue4" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue5" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue6" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue7" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue8" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue9" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue10" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue11" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue12" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue13" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue14" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue15" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue16" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue17" nbOfColumns="2">
```

```
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue18" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue19" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue20" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue21" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue22" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                </DatabaseTable>

                <DatabaseTable tableName="revenue23" nbOfColumns="2">
                    <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                    <DatabaseColumn columnName="total_revenue"
isUnique="false" />
```

```
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue24" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>

                    <DatabaseTable tableName="revenue25" nbOfColumns="2">
                        <DatabaseColumn columnName="supplier_no"
isUnique="false" />
                        <DatabaseColumn columnName="total_revenue"
isUnique="false" />
                    </DatabaseTable>



                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node2" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.2/tpch" connectionTestStatement="select 1">
                <DatabaseSchema dynamicPrecision="static">
                    <DatabaseStaticSchema>
                    <DatabaseProcedure name="loadtables" />
                    <DatabaseProcedure name="cleanup" />

                        <DatabaseTable tableName="supplier"
nbOfColumns="7">
                            <DatabaseColumn columnName="s_suppkey"
isUnique="true"/>
                            <DatabaseColumn columnName="s_name"
isUnique="false"/>
                            <DatabaseColumn columnName="s_address"
isUnique="false"/>
                            <DatabaseColumn columnName="s_nationkey"
isUnique="false"/>
```

```
                        <DatabaseColumn columnName="s_phone"
isUnique="false"/>
                        <DatabaseColumn columnName="s_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="s_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="customer"
nbOfColumns="8">
                        <DatabaseColumn columnName="c_custkey"
isUnique="true"/>
                        <DatabaseColumn columnName="c_name"
isUnique="false"/>
                        <DatabaseColumn columnName="c_address"
isUnique="false"/>
                        <DatabaseColumn columnName="c_nationkey"
isUnique="false"/>
                        <DatabaseColumn columnName="c_phone"
isUnique="false"/>
                        <DatabaseColumn columnName="c_acctbal"
isUnique="false"/>
                        <DatabaseColumn columnName="c_mktsegment"
isUnique="false"/>
                        <DatabaseColumn columnName="c_comment"
isUnique="false"/>
                    </DatabaseTable>

                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node1" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.1/tpch" connectionTestStatement="select 1">

        <DatabaseSchema dynamicPrecision="static">
            <DatabaseStaticSchema>

            <DatabaseProcedure name="loadtables"/>
                <DatabaseProcedure name="cleanup" />
```

```
                    <DatabaseTable tableName="part" nbOfColumns="9">
                        <DatabaseColumn columnName="p_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="p_name"
isUnique="false"/>
                        <DatabaseColumn columnName="p_mfgr"
isUnique="false"/>
                        <DatabaseColumn columnName="p_brand"
isUnique="false"/>
                        <DatabaseColumn columnName="p_type"
isUnique="false"/>
                        <DatabaseColumn columnName="p_size"
isUnique="false"/>
                        <DatabaseColumn columnName="p_container"
isUnique="false"/>
                        <DatabaseColumn columnName="p_retailprice"
isUnique="false"/>
                        <DatabaseColumn columnName="p_comment"
isUnique="false"/>
                    </DatabaseTable>

                    <DatabaseTable tableName="partsupp" nbOfColumns="5">
                        <DatabaseColumn columnName="ps_partkey"
isUnique="true"/>
                        <DatabaseColumn columnName="ps_suppkey"
isUnique="true"/>
                        <DatabaseColumn columnName="ps_availqty"
isUnique="false"/>
                        <DatabaseColumn columnName="ps_supplycost"
isUnique="false"/>
                        <DatabaseColumn columnName="ps_comment"
isUnique="false"/>
                    </DatabaseTable>


                </DatabaseStaticSchema>
            </DatabaseSchema>

        <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
          <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
        </ConnectionManager>
      </DatabaseBackend>
```

```
   <DatabaseBackend name="node4" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.4/tpch" connectionTestStatement="select 1">

                <DatabaseSchema dynamicPrecision="static">
                        <DatabaseStaticSchema>

                    <DatabaseProcedure name="loadtables"/>
                    <DatabaseProcedure name="cleanup" />

                                <DatabaseTable tableName="region" nbOfColumns="3">
                                        <DatabaseColumn columnName="r_regionkey"
isUnique="true"/>
                                        <DatabaseColumn columnName="r_name"
isUnique="false"/>
                                        <DatabaseColumn columnName="r_comment"
isUnique="false"/>
                                </DatabaseTable>

                                <DatabaseTable tableName="nation" nbOfColumns="4">
                                        <DatabaseColumn columnName="n_nationkey"
isUnique="true"/>
                                        <DatabaseColumn columnName="n_name"
isUnique="false"/>
                                        <DatabaseColumn columnName="n_regionkey"
isUnique="false"/>
                                        <DatabaseColumn columnName="n_comment"
isUnique="false"/>
                                </DatabaseTable>


                        </DatabaseStaticSchema>
                </DatabaseSchema>

    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
   </DatabaseBackend>

   <DatabaseBackend name="node5" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.5/tpch" connectionTestStatement="select 1">

                <DatabaseSchema dynamicPrecision="static">
                        <DatabaseStaticSchema>

                    <DatabaseProcedure name="loadtables"/>
```

```
                <DatabaseProcedure name="cleanup" />

                        <DatabaseTable tableName="orders" nbOfColumns="9">
                                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                                <DatabaseColumn columnName="o_clerk"
isUnique="false"/>
                                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                        </DatabaseTable>

                        <DatabaseTable tableName="lineitem"
nbOfColumns="16">
                                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
                                <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                                <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                                <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                                <DatabaseColumn columnName="l_tax"
isUnique="false"/>
                                <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                                <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
```

```
                <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                <DatabaseColumn columnName="l_comment"
isUnique="false"/>
        </DatabaseTable>

        <DatabaseTable tableName="orders_temp"
nbOfColumns="9">
                <DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_clerk"
isUnique="false"/>
                <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                <DatabaseColumn columnName="o_comment"
isUnique="false"/>
        </DatabaseTable>

        <DatabaseTable tableName="lineitem_temp"
nbOfColumns="16">
                <DatabaseColumn columnName="l_orderkey"
isUnique="true"/>
                <DatabaseColumn columnName="l_partkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_suppkey"
isUnique="false"/>
                <DatabaseColumn columnName="l_linenumber"
isUnique="true"/>
```

```
                        <DatabaseColumn columnName="l_quantity"
isUnique="false"/>
                        <DatabaseColumn columnName="l_extendedprice"
isUnique="false"/>
                        <DatabaseColumn columnName="l_discount"
isUnique="false"/>
                        <DatabaseColumn columnName="l_tax"
isUnique="false"/>
                        <DatabaseColumn columnName="l_returnflag"
isUnique="false"/>
                        <DatabaseColumn columnName="l_linestatus"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_commitdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_receiptdate"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipinstruct"
isUnique="false"/>
                        <DatabaseColumn columnName="l_shipmode"
isUnique="false"/>
                        <DatabaseColumn columnName="l_comment"
isUnique="false"/>
                    </DatabaseTable>

                </DatabaseStaticSchema>
            </DatabaseSchema>


    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
     <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <DatabaseBackend name="node6" driver="org.postgresql.Driver"
url="jdbc:postgresql://172.16.1.6/tpch" connectionTestStatement="select 1">

            <DatabaseSchema dynamicPrecision="static">
                <DatabaseStaticSchema>

                <DatabaseProcedure name="loadtables"/>
                <DatabaseProcedure name="cleanup" />

                    <DatabaseTable tableName="part" nbOfColumns="9">
```

<DatabaseColumn columnName="p_partkey"
isUnique="true"/>
<DatabaseColumn columnName="p_name"
isUnique="false"/>
<DatabaseColumn columnName="p_mfgr"
isUnique="false"/>
<DatabaseColumn columnName="p_brand"
isUnique="false"/>
<DatabaseColumn columnName="p_type"
isUnique="false"/>
<DatabaseColumn columnName="p_size"
isUnique="false"/>
<DatabaseColumn columnName="p_container"
isUnique="false"/>
<DatabaseColumn columnName="p_retailprice"
isUnique="false"/>
<DatabaseColumn columnName="p_comment"
isUnique="false"/>
</DatabaseTable>

<DatabaseTable tableName="orders" nbOfColumns="9">
<DatabaseColumn columnName="o_orderkey"
isUnique="true"/>
<DatabaseColumn columnName="o_custkey"
isUnique="false"/>
<DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
<DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
<DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
<DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
<DatabaseColumn columnName="o_clerk"
isUnique="false"/>
<DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
<DatabaseColumn columnName="o_comment"
isUnique="false"/>
</DatabaseTable>

<DatabaseTable tableName="orders_temp"
nbOfColumns="9">
<DatabaseColumn columnName="o_orderkey"
isUnique="true"/>

```
                    <DatabaseColumn columnName="o_custkey"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderstatus"
isUnique="false"/>
                    <DatabaseColumn columnName="o_totalprice"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderdate"
isUnique="false"/>
                    <DatabaseColumn columnName="o_orderpriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_clerk"
isUnique="false"/>
                    <DatabaseColumn columnName="o_shippriority"
isUnique="false"/>
                    <DatabaseColumn columnName="o_comment"
isUnique="false"/>
                </DatabaseTable>

            </DatabaseStaticSchema>
        </DatabaseSchema>

    <ConnectionManager vLogin="tpch" rLogin="tpch" rPassword="tpch">
      <VariablePoolConnectionManager initPoolSize="10" minPoolSize="5"
maxPoolSize="50" idleTimeout="300" waitTimeout="300"/>
    </ConnectionManager>
  </DatabaseBackend>

  <RequestManager>
    <RequestScheduler>
      <RAIDb-2Scheduler level="pessimisticTransaction"/>
    </RequestScheduler>

    <RequestCache>
      <MetadataCache/>
      <ResultCache granularity="table" />
    </RequestCache>

    <LoadBalancer>
      <RAIDb-2>
        <WaitForCompletion policy="first"/>
        <RAIDb-2-LeastPendingRequestsFirst />
      </RAIDb-2>
    </LoadBalancer>

  </RequestManager>
```

```
  </VirtualDatabase>
</C-JDBC>
```

APPENDIX D – JAVA TPC-H IMPLEMENTATION SOURCE CODE

tpchFullTest.bash

```
#(c) 2004 Brandon Rogers
#tpchFullTest.bash
#
# This file is used to instigate an entire TPC-H benchmark run.
# A ten minute sleep is executed to allow database servers to finish any
# leftover processing.

./loadTest.bash
echo Sleeping for 10 minutes...
sleep 10m
echo Woken up!
./perfTest.bash
```

loadTest.bash

```
#!/bin/bash

#(c) 2004 Brandon Rogers
#loadTest.bash
#
# This file is used to instigate a load test specified in the TPC-H specification

# Set the variables...

BASE_DIR=/home/brogers/thesis
SCRIPT_DIR=${BASE_DIR}/scripts
JAVA_EXE_LOCATION=${BASE_DIR}/java
INI_FILE=${BASE_DIR}/tpch.ini
LOG_FILE_LOCATION=${BASE_DIR}/logs
GTIME="date +%m/%d/%Y_%H:%M:%S.%N"
MKTIME="date +%s"

CLASS_PATH=${CLASS_PATH}:./c-jdbc-driver.jar

cd ${BASE_DIR}
```

```
        echo "`${GTIME}` (`${MKTIME}`) - Starting Database Load Test"


      cd ${JAVA_EXE_LOCATION}

      java -classpath $CLASS_PATH org.tpc.h.LT ${LOG_FILE_LOCATION}/LT.log
${INI_FILE}

        echo "`${GTIME}` (`${MKTIME}`) - Finished Database Load Test"
```

perfTest.bash

```
#!/bin/bash

#(c) 2004 Brandon Rogers
#perfTest.bash
#
# This file is used to instigate a performance test, containing both throughout and power
# tests, specified in the TPC-H specification

# Set the variables...

BASE_DIR=/home/brogers/thesis
SCRIPT_DIR=${BASE_DIR}/scripts
JAVA_EXE_LOCATION=${BASE_DIR}/java
INI_FILE=${BASE_DIR}/tpch.ini
LOG_FILE_LOCATION=${BASE_DIR}/logs
NUMBER_OF_RUNS=2
THROUGHPUT_TEST_THREADS=12
GTIME="date +%m/%d/%Y_%H:%M:%S.%N"
MKTIME="date +%s"

CLASS_PATH=${CLASS_PATH}:./c-jdbc-driver.jar

cd ${BASE_DIR}

#Do the Performance Test X number of Runs....

STREAM_COUNT=0
I=1
while [ $I -le $NUMBER_OF_RUNS ]
do

# Run the Power Test
```

```
        echo "`${GTIME}` (`${MKTIME}`) - Starting Power Test..."

        # Run RF1
        cd ${JAVA_EXE_LOCATION}
      echo "`${GTIME}` (`${MKTIME}`) - Starting RF1"

        java -classpath $CLASS_PATH org.tpc.h.RF1
${LOG_FILE_LOCATION}/PT.RF1 ${I} ${INI_FILE}
      echo "`${GTIME}` (`${MKTIME}`) - Finished RF1"

        # Run QS1
        cd ${JAVA_EXE_LOCATION}
      echo "`${GTIME}` (`${MKTIME}`) - Starting QS - Stream
${STREAM_COUNT}..."
        java -classpath $CLASS_PATH org.tpc.h.QS ${STREAM_COUNT}
${LOG_FILE_LOCATION}/PT.QS ${INI_FILE}
      echo "`${GTIME}` (`${MKTIME}`) - Finished QS - Stream
${STREAM_COUNT}..."

        # Run RF2
      cd ${JAVA_EXE_LOCATION}
      echo "`${GTIME}` (`${MKTIME}`) - Starting RF2..."
      java -classpath $CLASS_PATH org.tpc.h.RF2
${LOG_FILE_LOCATION}/PT.RF2 ${I} ${INI_FILE}
      echo "`${GTIME}` (`${MKTIME}`) - Finished RF2..."

        STREAM_COUNT=$((STREAM_COUNT+1))

# Run the Throughput Test
      # Run QS

        echo "`${GTIME}` (`${MKTIME}`) - Starting Throughput Test"

        J=1
        while [ $J -le $THROUGHPUT_TEST_THREADS ]
        do
        cd ${JAVA_EXE_LOCATION}
        echo "`${GTIME}` (`${MKTIME}`) - Starting QS - Stream
${STREAM_COUNT}..."
        java -classpath $CLASS_PATH org.tpc.h.QS ${STREAM_COUNT}
${LOG_FILE_LOCATION}/TT.QS ${INI_FILE} &
                J=$((J+1))
                STREAM_COUNT=$((STREAM_COUNT + 1))
        done
```

```
        #Throughput Test is still going on
        #Check for Throughput test threads still running… Loop until done.
        QS_THREAD_COUNT="`ps ax | grep java.*org.tpc.h.QS | grep -cv grep`"
        while [ $QS_THREAD_COUNT -ge 1 ]
        do
                QS_THREAD_COUNT="`ps ax | grep java.*org.tpc.h.QS | grep -cv
grep`"

                echo "" > /dev/null
        done

        echo "`${GTIME}` (`${MKTIME}`) - Finished Throughput Test"

        echo "`${GTIME}` (`${MKTIME}`) - Finished Run $I"

        I=$((I+1))



done

echo "`${GTIME}` (`${MKTIME}`) - TPC-H Benchmark Test Finished."
```

tpch.ini

DatabaseClass.java (org.tpc.h.DatabaseClass)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This file provides the main database interface for the entire TPC-H benchmark
 */

package org.tpc.h;

import java.sql.*;

public class DatabaseClass {

private Connection theConnect;
private Statement statement;
```

```
private String servers;
private String database;
private ResultSet result;
private String User = "";
private String Password = "";
public boolean isConnected = false;

        public DatabaseClass(String servers, String database, String Driver) {
                try {
                        Class.forName(Driver);
                        this.servers = servers;
                        this.database = database;
                }
                catch (ClassNotFoundException c) {
                        System.out.println(c);
                }

        } //constructor

        public DatabaseClass(String servers, String database, String Driver, String User,
String Password) {
                try {
                        Class.forName(Driver);
                        this.servers = servers;
                        this.database = database;
                        this.User = User;
                        this.Password = Password;
                }
                catch (ClassNotFoundException c) {
                        System.out.println(c);
                }

        } //constructor


        public boolean connect() {
                try {
                        //connect to the database
                        this.theConnect = DriverManager.getConnection("jdbc:cjdbc://" +
servers + "/" + database, this.User, this.Password);
                        //prepare a statement object
                        this.statement = this.theConnect.createStatement();
                        this.isConnected = true;

                        return this.isConnected;
                }
```

```
        catch (SQLException s) {
        System.out.println("jdbc:cjdbc://" + servers + "/" + database);
                System.out.println("Could not connect! ");
                System.out.println("ERROR: " + s);
                return false;
        }
        catch (NullPointerException s) {
                System.out.println("Could not connect! ");
                System.out.println("ERROR: " + s);
                return false;
        }

}

public void disconnect() {
        try {
                if( this.theConnect != null) {
                        this.theConnect.close();
                }
                this.theConnect = null;
        }
        catch (SQLException s) {
                System.out.println(s);
        }
        this.isConnected = false;
}

public void executeDDL(String ddlStatement) {
        try {
                if(!this.isConnected) {
                        this.connect();
                }

                this.result = null;
                this.statement.executeUpdate(ddlStatement);
        }
        catch (SQLException s) {
                System.out.println(s);
        }

}

public void executeQuery(String selectStatement) {
        try {
                if(!this.isConnected) {
                        this.connect();
```

```
                }

                this.result = this.statement.executeQuery(selectStatement);
        }
        catch (SQLException s) {
                System.out.println(s);
        }

}

public void execute(String statement) {
        try {
                if(!this.isConnected) {
                        this.connect();
                }

                this.result = null;
                this.statement.execute(statement);
        }
        catch (SQLException s) {
                System.out.println(s);
        }
}


public void query(String selectStatement) {
        try {
                this.result = this.statement.executeQuery(selectStatement);
        }
        catch (SQLException s) {
                System.out.println(s);
        }

}

public boolean first() {
        try {
                return this.result.first();
        }
        catch (SQLException s) {
                System.out.println(s);
                return false;
        }
}

public boolean nextRecord() {
```

```
        try {
                return this.result.next();
        }
        catch (SQLException s) {
                System.out.println(s);
                return false;
        }
}

public ResultSet getResult() {

        return this.result;
}

public void setAutoCommit(boolean OnOff) {
        try {
                this.theConnect.setAutoCommit(OnOff);
        }
        catch (SQLException s) {
                s.printStackTrace();
        }
}


public void startTransaction() {
        try {
                this.theConnect.setAutoCommit(false);
        }
        catch (SQLException s) {
                s.printStackTrace();
        }
}

public void commitTransaction() {
        try {
                this.theConnect.commit();
        }
        catch (SQLException s) {
                s.printStackTrace();
        }
}

public CallableStatement prepareCS(String sql) {
        return this.prepareCS(sql, true);
}
```

```
	public CallableStatement prepareCS(String sql, boolean readonly) {
		try {
			if(readonly) {
				return this.theConnect.prepareCall("{call " +
sql.toLowerCase() + " }");
			}
			else {
				return this.theConnect.prepareCall("}call " +
sql.toLowerCase() + " }");
			}
		}
		catch (SQLException s) {
			System.out.println(s);
			return null;
		}
	}

}
```

SQLConstants.java (org.tpc.h.SQLConstants)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This file loads all queries to be executed for a query stream.
 * This is accomplished using the tools QGEN and DBGEN provided with
 * the TPC-H Specification.  The DBGEN and QGEN programs are executed
 * and their results captured and queued for execution.
 *
 */
package org.tpc.h;

import java.io.*;
import java.util.*;

public class SQLConstants {
	/*
	These constants are for defining the queries!
	*/
	private final String QGEN_EXECUTE;
	private final String QGEN_LOCATION;
	private int scaleFactor;
	private String[] queryList = new String[22];
	private Properties configOptions;
```

```
        private String Seed;
        public int Permutation = 0;
        private String ReportFile;

        public SQLConstants(Properties configOptions, String ReportFile) {
                //we need to
                this.ReportFile = ReportFile;
                this.configOptions = configOptions;
                this.Seed = this.configOptions.getProperty("QGEN_SEED");
                this.QGEN_EXECUTE =
this.configOptions.getProperty("QGEN_EXECUTE");
                this.QGEN_LOCATION =
this.configOptions.getProperty("QGEN_LOCATION");
                this.scaleFactor =
Integer.parseInt(this.configOptions.getProperty("SCALE_FACTOR"));
        }

        private String getQueryByNumber(int index) {
                StringBuffer retVal = new StringBuffer();

                try {
                        Runtime rt = Runtime.getRuntime();
            Process proc = rt.exec(QGEN_EXECUTE + " " + scaleFactor + " " +
this.Permutation + " " + this.ReportFile  + " " + this.Seed  + " " + index);
//              InputStream stdin = proc.getErrorStream();
                InputStream stdin = proc.getInputStream();
                InputStreamReader isr = new InputStreamReader(stdin);
                BufferedReader br = new BufferedReader(isr);
                String line = null;

                //now just need to capture the query!
                while ( (line = br.readLine()) != null) {

                   if((!line.equals("")) && (line.trim().charAt(0) != '-'))
                        retVal.append(line.trim() + " ");
                }
                   }
                   catch (IOException e) {
                                e.printStackTrace();
                                System.exit(-1);
                   }

                   return retVal.toString().replaceAll("\r", " ").replaceAll("\n", " ").trim();
        }

        public String getQueryByQueueNumber(int index) {
```

```
                return this.queryList[(index-1)].replaceAll(";", "\n").trim();
        }

        public void prepareQueryStream() {
                for(int i = 0; i < queryList.length; i++) {
                        queryList[i] = this.getQueryByNumber((i+1));
                }
        }
}
```

Log.java (org.tpc.h.Log)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * A common interface for created logs of the TPC-H benchmark.
 */
package org.tpc.h;

import java.io.*;
import java.util.*;
import java.text.*;

public class Log {

        private String filename;
        private BufferedWriter fileLogger;
        private boolean autoFlushToFile = true;
        private StringBuffer RAMLog = new StringBuffer();
        private SimpleDateFormat df;

        public Log(String Filename) {
                this.filename = Filename;

                this.df = new SimpleDateFormat("MM/dd/yyyy hh:mm:ss.SSS a");

                //open the file and get it ready for appending.
                try {
                        this.fileLogger = new BufferedWriter(new
FileWriter(this.filename));
                }
                catch (IOException e) {
                        e.printStackTrace();
                        System.exit(0);
                }
```

```
	}

	public void addEntry(String EntryLine) {
		this.addEntry(EntryLine, false);
	}

	public void addEntry(String EntryLine, boolean echoToSTDIN) {
		if(autoFlushToFile) {
			try {
				this.fileLogger.write(this.df.format(new Date()) + " | " +
EntryLine + "\r\n");
				this.fileLogger.flush();
			}
			catch (IOException e) {
				e.printStackTrace();
				System.exit(0);
			}

		}
		else {
			RAMLog.append(this.df.format(new Date()) + " | " + EntryLine +
"\r\n");
		}
		if(echoToSTDIN) System.out.println(this.df.format(new Date()) + " | " +
EntryLine);
	}

	public void setAutoFlush() {
		this.autoFlushToFile = true;
		this.flushRAMToFile();
		RAMLog = new StringBuffer();
	}

	public void unsetAutoFlush() {
		RAMLog = new StringBuffer();
		this.autoFlushToFile = false;
	}

	public void close() {
		this.autoFlushToFile = true;
		this.flushRAMToFile();
		try {
			this.fileLogger.close();
		}
		catch (IOException e) {
			e.printStackTrace();
```

```
                        System.exit(0);
                }
        }

        private void flushRAMToFile() {
                try {
                        this.fileLogger.write(this.RAMLog + "\r\n");
                        this.fileLogger.flush();
                }
                catch (IOException e) {
                        e.printStackTrace();
                        System.exit(0);
                }
        }
}
```

LT.java (org.tpc.h.LT)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This executes the Load Test using the DatabaseClass.java and
 * SQLConstants.java files.
 */
package org.tpc.h;

import java.util.*;
import java.io.*;
import java.sql.*;

public class LT {
        public static void main(String[] args) {
                if(args.length == 1) {
                        LTEngine engine = new LTEngine(args[0]);
                        engine.go();
                }
                else if(args.length == 2) {
                        LTEngine engine = new LTEngine(args[0], args[1]);
                        engine.go();
                }
                else {
                        System.out.println("USAGE: java LT ReportFile [INIFile]");
                }

                System.out.println("LT Done.");
```

```
	}
}

class LTEngine {

	private final boolean echoToScreen = true;
	private String INIFile;
	private String ReportFile;
	private int Permutation;
	private Properties configurationOptions = new Properties();
	private int scaleFactor;
	private SQLConstants queries;
	private DatabaseClass db;
	private Log logFile;
	private BufferedReader fileReader;

	public LTEngine(String ReportFile) {
		this.initVars(ReportFile, "tpch.ini");
	}

	public LTEngine(String ReportFile, String INIFile) {
		this.initVars(ReportFile, INIFile);
	}

	private void initVars(String ReportFile, String INIFile) {
		this.INIFile = INIFile;
		this.ReportFile = ReportFile;

		try {
			FileInputStream propFile = new FileInputStream(INIFile);
			configurationOptions.load(propFile);
		}
		catch (FileNotFoundException e) { System.out.println("FILE NOT
FOUND: " + INIFile); e.printStackTrace(); }
		catch (IOException e) { e.printStackTrace(); }

		this.logFile = new Log(ReportFile);
		this.logFile.setAutoFlush();
	}

	public void go() {
		//all the queries are prepared and now ready to run...
		//just do them...
		String query;
		CallableStatement cs;
```

```
		this.logFile.addEntry("Creating database object", echoToScreen);
		this.db = new
DatabaseClass(configurationOptions.getProperty("SERVER_IP"),
configurationOptions.getProperty("DATABASE_NAME"),
configurationOptions.getProperty("DB_DRIVER_CLASS_NAME"),
configurationOptions.getProperty("DB_USER"),
configurationOptions.getProperty("DB_PASSWORD"));

		this.logFile.addEntry("Connecting to the Database...",echoToScreen);
		this.db.connect();
		this.logFile.addEntry("Connected!",echoToScreen);

		this.logFile.addEntry("Creating Tables...",echoToScreen);

		String line;
		try {
			this.fileReader = new BufferedReader(new
FileReader(configurationOptions.getProperty("SCRIPT_DIR") + "/createTables.sql"));

			while((line = fileReader.readLine()) != null) {
				line.trim();
				if(!line.equals("")) {
					this.logFile.addEntry("Running query: " + line);
					db.execute(line);
				}
			}
		}
		catch (IOException e) {
			e.printStackTrace();
		}

		this.logFile.addEntry("Tables Created.",echoToScreen);

		this.logFile.addEntry("Loading Data...",echoToScreen);

		try {
			this.logFile.addEntry("Loading Data [PART]...", echoToScreen);
			cs = this.db.prepareCS("loadTables('part');");
			cs.execute();

			this.logFile.addEntry("Loading Data [SUPPLIER]...",
echoToScreen);
			cs = this.db.prepareCS("loadTables('supplier');");
			cs.execute();
```

```
                this.logFile.addEntry("Loading Data [PARTSUPP]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('partsupp');");
                cs.execute();

                this.logFile.addEntry("Loading Data [CUSTOMER]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('customer');");
                cs.execute();

                this.logFile.addEntry("Loading Data [ORDERS]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('orders');");
                cs.execute();

                this.logFile.addEntry("Loading Data [LINEITEM]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('lineitem');");
                cs.execute();

                this.logFile.addEntry("Loading Data [NATION]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('nation');");
                cs.execute();

                this.logFile.addEntry("Loading Data [REGION]...",
echoToScreen);
                cs = this.db.prepareCS("loadTables('region');");
                cs.execute();

            }
            catch(SQLException s) {
                System.out.println(s);
            }

            this.logFile.addEntry("Data Loaded.", echoToScreen);

            this.logFile.addEntry("Preparing Database (ANALYZE)...",
echoToScreen);

            try {
               cs = this.db.prepareCS("cleanUp();");
          cs.execute();
            }
            catch(SQLException s) {
                System.out.println(s);
```

```
                }

                this.logFile.addEntry("Database Ready.", echoToScreen);

                this.logFile.addEntry("Creating Indexes.",echoToScreen);

                try {
                        this.fileReader = new BufferedReader(new
FileReader(configurationOptions.getProperty("SCRIPT_DIR") + "/createIndexes.sql"));

                        while((line = fileReader.readLine()) != null) {
                                line.trim();
                      if(!line.equals("")) {
                          this.logFile.addEntry("Running query: " + line);
                          db.execute(line);
                      }
                        }
                }
                catch (IOException e) {
                        e.printStackTrace();
                }

                this.logFile.addEntry("Indexes Created.",echoToScreen);

            this.logFile.addEntry("Preparing Database (ANALYZE)...", echoToScreen);

                try {
                        cs = this.db.prepareCS("cleanUp();");
                        cs.execute();
                }
                catch(SQLException s) {
                        System.out.println(s);
                }

                //this is removed to use valid CS
            //this.db.execute("{call cleanUp(); }");

            this.logFile.addEntry("Database Ready.", echoToScreen);


                this.db.disconnect();
                this.logFile.close();

        }
}
```

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This executes the New Sales Refresh Function
 * using the DatabaseClass.java file.
 */

package org.tpc.h;

import java.sql.*;
import java.util.*;
import java.io.*;

public class RF1 {
	public static void main(String[] args) {

		if(args.length == 2) {
			RF1Engine engine = new RF1Engine(args[0], args[1]);
			engine.go();
		}
		else if(args.length == 3) {
			RF1Engine engine = new RF1Engine(args[0], args[1], args[2]);
			engine.go();
		}
		else {
			System.out.println("USAGE: java RF1 ReportFile RefreshNumber
[INIFile]");
		}
	}

}

class RF1Engine {

	private String INIFile;
	private String ReportFile;
	private Properties configurationOptions = new Properties();
	private DatabaseClass db;
	private DatabaseClass db2;
	private String refreshCycle;
	private BufferedReader fileReader;
	private Log logFile;

	public RF1Engine(String ReportFile, String refreshCycle, String INIFile) {
		init(ReportFile, refreshCycle, INIFile);
```

```
	} //constructor

	public RF1Engine(String ReportFile, String refreshCycle) {
		init(ReportFile, refreshCycle, "tpch.ini");
	} //constructor

	private void init(String ReportFile, String refreshCycle, String INIFile) {
		this.INIFile = INIFile;
		this.ReportFile = ReportFile;
		this.refreshCycle = refreshCycle;

		try {
			FileInputStream propFile = new FileInputStream(INIFile);
			configurationOptions.load(propFile);
		}
		catch (FileNotFoundException e) { System.out.println("FILE NOT
FOUND: " + INIFile); e.printStackTrace(); }
		catch (IOException e) { e.printStackTrace(); }

		this.logFile = new Log(ReportFile + "." + this.refreshCycle);
		this.logFile.setAutoFlush();

		this.db = new
DatabaseClass(configurationOptions.getProperty("SERVER_IP"),
configurationOptions.getProperty("DATABASE_NAME"),
configurationOptions.getProperty("DB_DRIVER_CLASS_NAME"),
configurationOptions.getProperty("DB_USER"),
configurationOptions.getProperty("DB_PASSWORD"));
		this.db2 = new
DatabaseClass(configurationOptions.getProperty("SERVER_IP"),
configurationOptions.getProperty("DATABASE_NAME"),
configurationOptions.getProperty("DB_DRIVER_CLASS_NAME"),
configurationOptions.getProperty("DB_USER"),
configurationOptions.getProperty("DB_PASSWORD"));

	}

	public void go() {
		String flatRecord;
		CallableStatement cs;

		this.logFile.addEntry("RF1 Started.", true);

		this.db.connect();

		this.logFile.addEntry("Adding Records.", true);
```

```
		try {

			//insert order records into temp table.

			this.fileReader = new BufferedReader(new
FileReader(configurationOptions.getProperty("DBGEN_LOCATION") + "/orders.tbl.u"
+ refreshCycle));

			while((flatRecord = this.fileReader.readLine()) != null) {

				flatRecord = flatRecord.substring(0, flatRecord.length() -
1);
				flatRecord = flatRecord.replaceAll("\\|", "', '").trim();
				this.db.execute("INSERT INTO orders_temp VALUES ('"
+ flatRecord + "');");

			}
			//insert lineitem records into temp table.

			this.fileReader = new BufferedReader(new
FileReader(configurationOptions.getProperty("DBGEN_LOCATION") +
"/lineitem.tbl.u" + refreshCycle));

			while((flatRecord = this.fileReader.readLine()) != null) {

				flatRecord = flatRecord.substring(0, flatRecord.length() -
1);
				flatRecord = flatRecord.replaceAll("\\|", "', '").trim();
				this.db.execute("INSERT INTO lineitem_temp VALUES
('" + flatRecord + "');");

			}

			//all Records are now in tables and it is easy to
			//insert records keeping valid relationships!
			this.db.executeQuery("select DISTINCT o_orderkey from
orders_temp");
			while(this.db.nextRecord()) {
				this.db2.execute("INSERT INTO orders (select * from
orders_temp where o_orderkey = " + this.db.getResult().getString("o_orderkey") + ");");
				this.db2.execute("INSERT INTO lineitem (select * from
lineitem_temp where l_orderkey = " + this.db.getResult().getString("o_orderkey") +
");");
			}
		}
		catch (SQLException s) {
```

```
                    s.printStackTrace();
                }
                catch (IOException e) {
                    e.printStackTrace();
                }

                //clean up temp tables;
                this.db.execute("delete from orders_temp");
                this.db.execute("delete from lineitem_temp");

            this.logFile.addEntry("Preparing Database (ANALYZE)...", true);
                try {
                    cs = this.db.prepareCS("cleanUp();");
                    cs.execute();
                }
                catch(SQLException s) {
                    System.out.println(s);
                }

            this.logFile.addEntry("Database Ready.", true);
                this.logFile.addEntry("RF1 Finished.", true);
                this.logFile.close();

                this.db.disconnect();
                this.db2.disconnect();
        }
}
```

RF2.java (org.tpc.h.RF2)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This executes the Old Sales Refresh Function
 * using the DatabaseClass.java file.
 */
package org.tpc.h;

import java.util.*;
import java.io.*;
import java.sql.*;

public class RF2 {
        public static void main(String[] args) {
```

```
		if(args.length == 2) {
			RF2Engine engine = new RF2Engine(args[0], args[1]);
			engine.go();
		}
		else if(args.length == 3) {
			RF2Engine engine = new RF2Engine(args[0], args[1], args[2]);
			engine.go();
		}
		else {
			System.out.println("USAGE: java RF2 ReportFile RefreshNumber
[INIFile]");
		}
	}

}

class RF2Engine {

	private String INIFile;
	private String ReportFile;
	private Properties configurationOptions = new Properties();
	private DatabaseClass db;
	private String refreshCycle;
	private BufferedReader deleteFile;
	private Log logFile;

	public RF2Engine(String ReportFile, String refreshCycle, String INIFile) {
		init(ReportFile, refreshCycle, INIFile);
	} //constructor

	public RF2Engine(String ReportFile, String refreshCycle) {
		init(ReportFile, refreshCycle, "tpch.ini");
	} //constructor

	private void init(String ReportFile, String refreshCycle, String INIFile) {
		this.INIFile = INIFile;
		this.ReportFile = ReportFile;
		this.refreshCycle = refreshCycle;

		try {
			FileInputStream propFile = new FileInputStream(INIFile);
			configurationOptions.load(propFile);

			this.deleteFile = new BufferedReader(new
FileReader(configurationOptions.getProperty("DBGEN_LOCATION") + "/delete." +
refreshCycle));
```

```
                }
                catch (FileNotFoundException e) { System.out.println("FILE NOT
FOUND: " + INIFile); e.printStackTrace(); }
                catch (IOException e) { e.printStackTrace(); }

                this.logFile = new Log(ReportFile + "." + this.refreshCycle);
                this.logFile.setAutoFlush();

                this.db = new
DatabaseClass(configurationOptions.getProperty("SERVER_IP"),
configurationOptions.getProperty("DATABASE_NAME"),
configurationOptions.getProperty("DB_DRIVER_CLASS_NAME"),
configurationOptions.getProperty("DB_USER"),
configurationOptions.getProperty("DB_PASSWORD"));

        }

        public void go() {
                String deleteID;
                CallableStatement cs;

                this.db.connect();

                try {

                        while((deleteID = this.deleteFile.readLine()) != null) {
                                deleteID.trim();

                                this.logFile.addEntry("Queueing deletion of ID: " +
deleteID);
                                this.db.execute("INSERT INTO orders_temp VALUES ("
+ deleteID + ");");
                                this.db.execute("INSERT INTO lineitem_temp VALUES
(" + deleteID + ", null, null, 1);");
                        }

                        this.logFile.addEntry("Deleting queued records...");
                        this.db.execute("delete from lineitem where l_orderkey IN (select
l_orderkey from lineitem_temp);");
                        this.db.execute("delete from orders where o_orderkey IN (select
o_orderkey from orders_temp);");

                }
                catch (IOException e) {
                        e.printStackTrace();
                }
```

```
            this.logFile.addEntry("Preparing Database (ANALYZE)...", true);

                try {
                        cs = this.db.prepareCS("cleanUp();");
                        cs.execute();
                }
                catch(SQLException s) {
                        System.out.println(s);
                }


            this.logFile.addEntry("Database Ready.", true);

                this.logFile.addEntry("RF2 Finished.", true);
                this.logFile.close();

                this.db.disconnect();

        }
}
```

QS.java (org.tpc.h.QS)

```
/*
 * (c) 2004 Brandon Rogers
 *
 * This executes the one complete query stream of 22
 * queries as defined by TPC-H
 */

package org.tpc.h;

import java.util.*;
import java.io.*;

public class QS {
        public static void main(String[] args) {
                if(args.length == 2) {
                        QSEngine engine = new QSEngine(args[0], args[1]);
                        engine.go();
                }
                else if(args.length == 3) {
                        QSEngine engine = new QSEngine(args[0], args[1], args[2]);
                        engine.go();
```

```
                }
                else {
                        System.out.println("USAGE: java QS Permutation ReportFile
[INIFile]");
                }

        }
}

class QSEngine {

        private String INIFile;
        private String ReportFile;
        private int Permutation;
        private Properties configurationOptions = new Properties();
        private DatabaseClass db;
        private int scaleFactor;
        private SQLConstants queries;
        private Log logFile;

        public QSEngine(String Permutation, String ReportFile) {
                this.initVars(Permutation, ReportFile, "tpch.ini");
        }

        public QSEngine(String Permutation, String ReportFile, String INIFile) {
                this.initVars(Permutation, ReportFile, INIFile);
        }

        private void initVars(String Permutation, String ReportFile, String INIFile) {
                this.INIFile = INIFile;
                this.ReportFile = ReportFile;

                try {
                        FileInputStream propFile = new FileInputStream(INIFile);
                        configurationOptions.load(propFile);
                }
                catch (FileNotFoundException e) { System.out.println("FILE NOT
FOUND: " + INIFile); e.printStackTrace(); }
                catch (IOException e) { e.printStackTrace(); }

                this.logFile = new Log(ReportFile + "." + Permutation);
                this.logFile.setAutoFlush();

                this.logFile.addEntry("Setting up variables for test.");
```

```
		this.scaleFactor =
Integer.parseInt(configurationOptions.getProperty("SCALE_FACTOR"));

		this.Permutation = Integer.parseInt(Permutation);

		this.queries = new SQLConstants(this.configurationOptions,
this.ReportFile);
		this.queries.Permutation = this.Permutation;

		this.logFile.addEntry("Generating query stream with QGEN.");

		this.queries.prepareQueryStream();

	}

	public void go() {
		//all the queries are prepared and now ready to run...
		//just do them...
		String query;

		this.logFile.addEntry("Running Queries for Permutation " +
this.Permutation);

		//connect to the db
		this.db = new
DatabaseClass(configurationOptions.getProperty("SERVER_IP"),
configurationOptions.getProperty("DATABASE_NAME"),
configurationOptions.getProperty("DB_DRIVER_CLASS_NAME"),
configurationOptions.getProperty("DB_USER"),
configurationOptions.getProperty("DB_PASSWORD"));
		this.logFile.addEntry("Connecting to the Database...");

		this.db.connect();

		this.logFile.addEntry("Connected!");

		for(int i = 1; i <= 22; i++) {
			try {
				query = this.queries.getQueryByQueueNumber(i);

				String[] QUERY_RESULT = query.split("~");
				for (int x=0; x < QUERY_RESULT.length; x++) {
					QUERY_RESULT[x] =
QUERY_RESULT[x].trim();

					if(QUERY_RESULT[x].startsWith("select")) {
```

```
                                        //this is a selection
                                        this.logFile.addEntry("Preparing the
following query...");

        this.logFile.addEntry(QUERY_RESULT[x]);

                                        this.logFile.addEntry("Query Started (" + i +
").");

        this.db.executeQuery(QUERY_RESULT[x]);

                                        this.logFile.addEntry("Query Finished.");
                                        this.logFile.addEntry("ResultSet Retrieval
Started.");
                                        int rowsReturned = 0;
                                        while(this.db.nextRecord()) {

                                                int curCol = 1;
                                                StringBuffer rowString = new
StringBuffer();
                                                while(curCol <=
this.db.getResult().getMetaData().getColumnCount()) {
                                                        rowString.append("|" +
this.db.getResult().getString(curCol));
                                                        curCol++;
                                                }
                                                rowString.append("|");


        this.logFile.addEntry(rowString.toString());

                                                rowsReturned++;
                                                this.logFile.addEntry("Row " +
rowsReturned + " Retrieved.");
                                        }

                                }
                                else {
                                        this.logFile.addEntry("Preparing the
following query...");

        this.logFile.addEntry(QUERY_RESULT[x]);
                                        this.logFile.addEntry("Query Started (" + i +
").");
                                        this.db.execute(QUERY_RESULT[x]);
                                        this.logFile.addEntry("Query Finished.");
```

```
                                        }

                                }

                                this.logFile.addEntry("ResultSet Retrieval Finished.");
                        }
                        catch (Exception e) {
                                e.printStackTrace();
                        }
                }

                this.logFile.addEntry("QS (" + Permutation + ") Done.", true);
                this.db.disconnect();
                this.logFile.close();
        }
}
```