Cost Analysis Of Sun/Oracle and Unisys/Microsoft for a BI Solution in a VLDB Environment

February 2002

Prepared by:

Walklett Group
65 Valley Stream Pkwy, Suite 250
Malvern, PA 19355
610-722-9600
www.walklett.com

DISCLAIMER

This document contains references to companies, products, and their trademarks. All trademarks remain the property of their respective trademark holders.

The information researched to create this white paper was obtained from sources considered to be reliable. Walklett Group shall have no responsibility for errors, omissions or inadequacies of this information.

All opinions expressed in this white paper are those of Walklett Group, except where the original author is credited. The opinions expressed in this white paper are subject to change without notice.

All costs were determined using the list/retail price available to the public at the time - no corporate or other discounts were applied.

Table of Contents

I. Executive Summary

Businesses today need every competitive advantage they can get, and Information Technology (IT) management is being asked to deliver solutions that give a competitive advantage to the business. One of the best ways to maintain an advantage is to understand the condition of your business, and the best way to get an understanding is to implement a Business Intelligence (BI) solution. BI is the application of business knowledge to data, which transforms it into business information. This data can come from existing CRM, ERP, or other operational systems. BI takes this data and manipulates it in a timely manner so that it can be analyzed in order to have a positive impact on the business. BI includes data warehousing, data mining, and on-line analytical processing (OLAP) that allows the user community to extract the data so it can be analyzed and viewed in a usable format.

This type of analysis requires a large amount of data about the business and its customers. Finding the most cost-effective BI solutions involving a large amount of data only increases this challenge and its cost. There are many hardware and software options to implement a BI solution. This paper was developed to explore two solutions that are available, Sun/Oracle and Unisys/Microsoft SQL Server 2000. More importantly, it examines the cost differences between these two solutions.

A. *Analysis*

For this paper, we wanted to compare two solutions that utilized different architectures. The most common architectures for BI solutions are RISC/UNIX and the Windows/Intel architectures. Sun Microsystems (Sun) is the market leader with the RISC/UNIX architecture. We selected the Sun Fire 6800 with Oracle 9*i* as the database for our RISC/UNIX architecture because it represented Sun's offering in the high availability midrange server market.

Unisys represents a company solely dedicated to the Intel/Windows/Intel architecture. For our Windows/Intel architecture we used the Unisys ES7000 with Microsoft SQL Server 2000 as the database because it is their premier server for mission-critical applications.

Our process involved identifying all the cost categories (Capital Expenses, Administration, Technical Support, and End User Operations). We then compared the costs by category and documented those categories in which a cost difference exists. Table 1 below illustrates the solution that provides the lowest cost in that category for the same function. The solution with the lowest cost in that cost category is identified.

Configuration:	Sun Fire 6800 /Oracle 9*i*	Unisys ES7000/ Microsoft SQL Server 2000
Hardware		•
Hardware Maintenance Agreement		•
Database Software		•
Multidimensional Database Software		•
Data Mining Software		•
Maintenance/Update Agreements (Operating System)	•	
Maintenance/Update Agreements (Database Software)		•
Application Developer	•	
OS Administrator		•
DBA		•

Table 1 – Cost Comparison

B. *Conclusion*

This paper looks at two alternatives offered by some of the most recognized names in IT today. These hardware and software vendors offer solutions with comparable functionality, but with significantly different costs. In some instances these costs differences are dramatic. Figure 1 – Cost Differences by Year illustrates these differences.

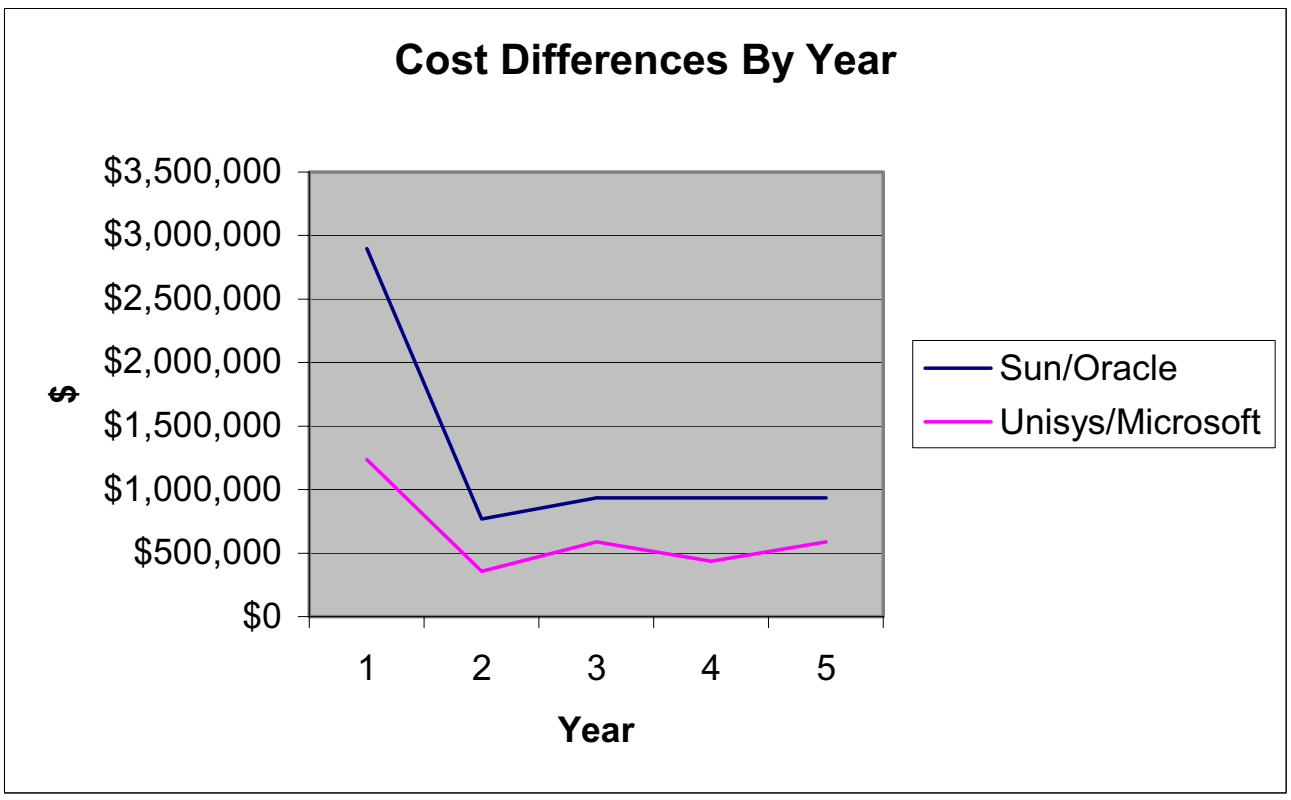

Figure 1 – Cost Differences by Year

The total savings for selecting the Unisys/Microsoft SQL Server 2000 solution is $3.2 million over 5 years. Based on the retail prices available to the public, the final conclusion can only be that the Unisys/Microsoft SQL Server 2000 BI solution is less expensive to deploy and operate than a Sun/Oracle solution over the same time period.

II. Introduction

Making the right business decisions is a means to success in business. Every day millions of business decisions are made. The one element all of these decisions have in common is that they are made with the best information available. However, many of these business decisions do not have the desired effect. In many cases, the ineffectiveness was due to incorrect information. This incorrect information could have been too old, reported incorrectly, or not reported at all. Business Intelligence is a technology solution that increases the likelihood that the pertinent information is available to make the most informed business decision possible.

Many companies use BI to make informed decisions, which leads to reduced costs and higher profits. Some BI implementations require the storage of a large amount of data about the business. This requirement has mandated the use of the Very Large Databases (VLDB) as part of the BI solution.

The costs to initiate a BI solution in a VLDB environment can be staggering. There are many hardware and software vendors that have products to meet the challenges of BI within a VLDB environment. This paper examines the differences in costs between two BI solutions in a VLDB environment. The costs of these solutions were determined using the criteria established in the Configuration Section of this paper.

Since we are only examining the differences in the costs, items that have equal costs in both solutions will be mentioned without their respective costs.

A. *Business Intelligence (BI)*

Conducting business today requires investing in the infrastructure necessary to better understand your business, and make better business decisions. Analyzing business information derived from business data is necessary in order to remain competitive.

The hardware, software, practices, and methodologies that allow this type of analysis constitute the category of BI. Figure 2 – Elements of BI below illustrates the elements of a typical BI solution.

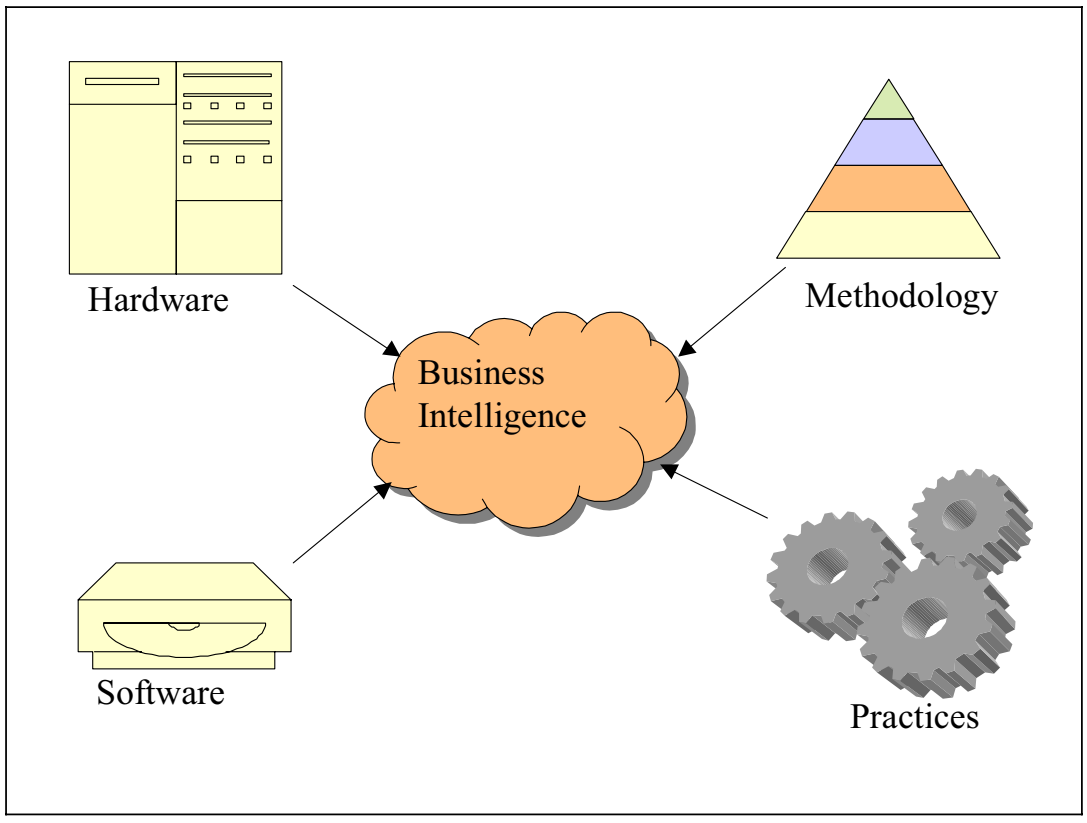

Figure 2 - Elements of BI

BI is the bridge between the user community and data in transactional systems. Transactional systems are used to conduct business on a regular basis, such as order processing, payroll, and inventory to name a few. The transactional data is collected and integrated using rules that govern how the business operates so that an enterprise-wide view of the business can be examined. It is then stored in one central location typically called a Data Warehouse. The Data Warehouse allows the user community to analyze the data from the transactional systems to make business decisions based on this single view of the enterprise.

Figure 3 – System Interaction shows how the different systems interact with one another to generate the business information needed by the user community.

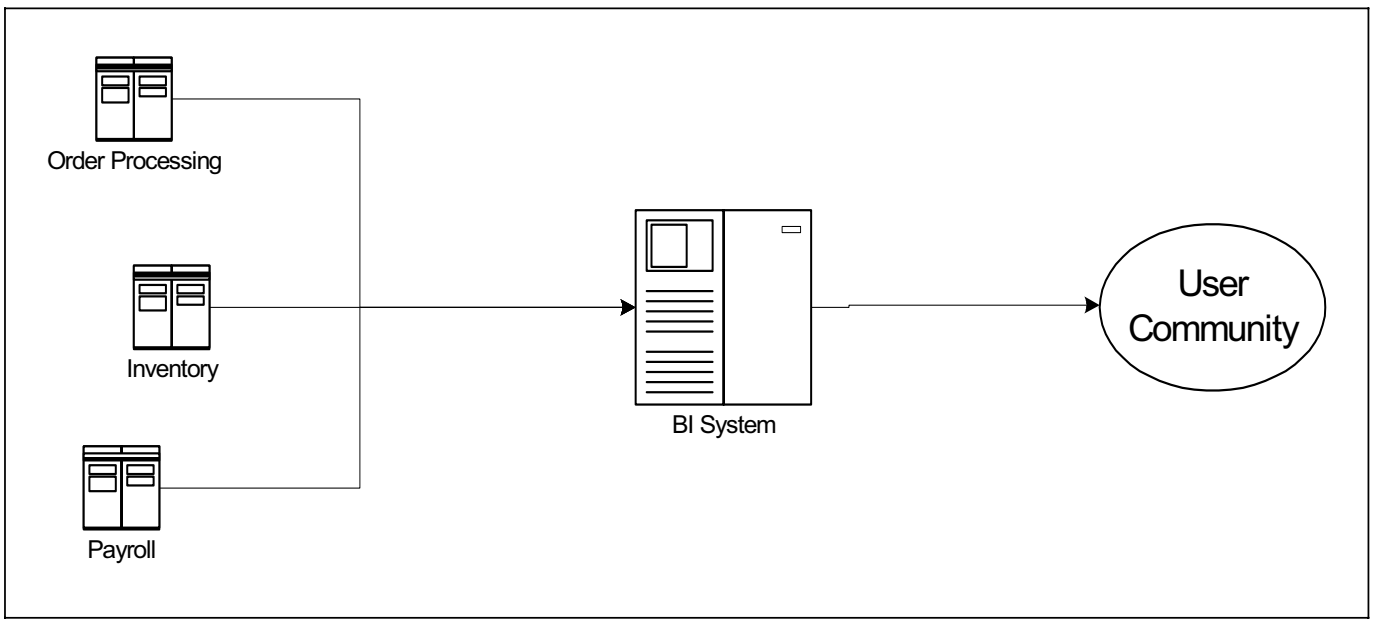

Figure 3 – System Interaction

B. *BI in Today's Environment*

Today, BI is being utilized to monitor the supply chain, build customer relationships, and create business reports. More companies are empowering their employees to be decision makers, and these new decision makers need information. This trend has increased the size of the user community from a small group of decision makers to almost everyone in an organization.

The more information the user community is able to generate, the more information they will want to examine, and greater will be the strain on the hardware managing the data. This strain can cause bottlenecks in the hardware, which will increase the amount of time the user community has to wait for their information.

To reduce the potential of bottlenecks, many businesses are turning to Symmetric Multi-Processing (SMP) hardware architecture implementations as a way to deliver this information and control the costs of BI.

The SMP architecture allows a single application to use multiple processors and the assignment of applications to one or more processors within the same machine. This gives an application the flexibility to use as much processing power as it needs to be effective and efficient. For example, if there are 32 processors, 24 can be assigned to the database and 8 can be assigned to the Web server application within the same machine. The SMP architecture also allows for larger memory configurations. These larger memory configurations improve performance and reduce the risks of a bottleneck.

The software approach to improve performance to the user community is to create multi-dimensional databases (MDDB). An MDDB stores the data in predefined summaries. Summaries are comprised of regularly requested data from the user community. Predefined summaries reduce the time needed to get the data to the user community.

An MDDB is often referred to as a cube. Cubes are constructed using a ROLAP, MOLAP, or HOLAP format. The particular format employed is determined based on the analytical requirements of the user community.

The various hardware and software approaches allow for a large amount of data to be available to the user community in a timely manner.

C. *Characteristics of a Very Large Database*

The demand for more information, and the data needed to create it, has increased the use of a VLDB in a BI solution. The term "Large" is relative. For some, several

gigabytes is large, whereas others feel that anything more than a terabyte (1,000 gigabytes) is large. Most experts agree that one terabyte or greater falls into the VLDB realm.

The user community expects information instantly. The very nature of the VLDB makes it a challenge to meet these expectations. The challenge in successfully implementing and using a VLDB is to find the optimum hardware and software configuration, along with sound BI implementation techniques, which will avoid slow response to requests for data.

III. Configuration Assumptions

For our scenario, we defined the parameters for a hypothetical BI implementation with characteristics consistent with a VLDB for a large organization. This large organization needs a BI solution that is available 24 hours a day, 7 days a week. The organization does not have any resources – hardware, software, or staff – to support this BI endeavor. Therefore, it has to purchase new hardware, hire new staff members, and purchase all the database software. We will measure the costs for a life cycle of five years. We believe that five years would make our cost assumptions more reliable.

Figure 4 – Configuration Assumptions illustrates the main configuration assumptions we used to define our scenario. The rest of this section outlines these configuration assumptions in detail.

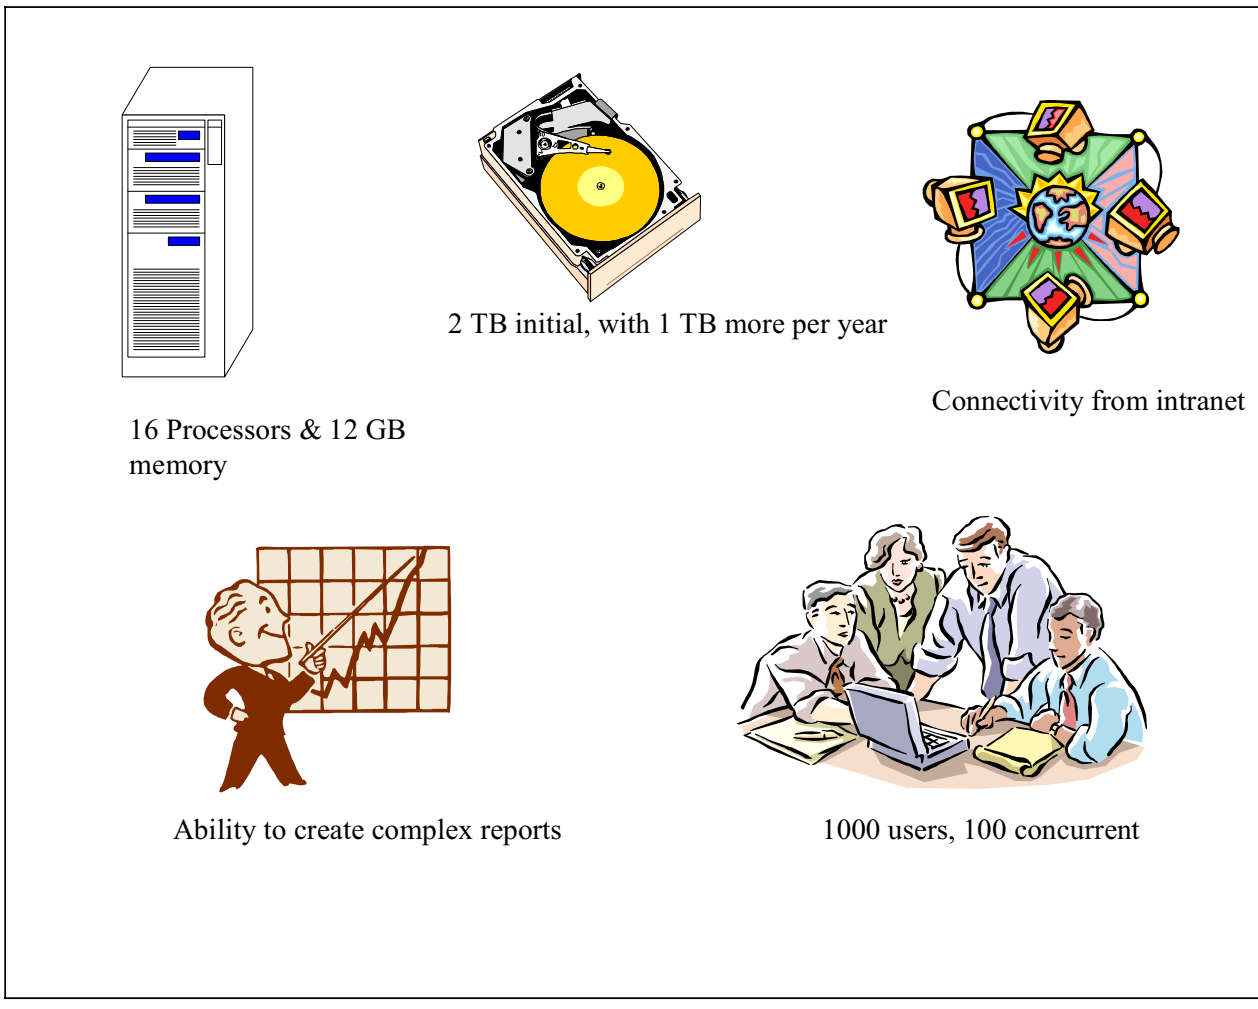

Figure 4 – Configuration Assumptions

A. *Data Warehouse*

Our initial data warehouse would be two terabytes, which captures the last two years of existing data. The data warehouse would grow by one terabyte annually, resulting in a six-terabyte warehouse for the life cycle that we are going to measure. To enhance performance we would create an MDDB from the data warehouse. This MDDB would reduce the time needed to extract regularly requested sets of data.

B. *User Community*

To develop a cost for the hardware and software licenses, the user community is made up of 1,000 authorized users. We expect only 100 of those users will be accessing the data concurrently.

C. *Business Needs*

The user community is interested in creating their own reports and doing some advanced manipulation of the data. Their own reports would be the usual weekly and monthly reports they need to monitor the business. The advanced manipulation includes creating forecasts, budgets, and complex financial calculations. The user community is also interested in finding patterns in the sales data.

D. *User Interface*

The user community will interface with the databases using a Web-based tool on their company intranet. For our scenario, the data will be accessed only through the company intranet. We chose this constraint due to the security issues that surround an Internet deployment. The hardware and software we reviewed is capable of connecting to the Internet, but the cost for establishing and maintaining a secured link is beyond the scope of this paper.

E. *Hardware Requirements*

We believe that the hardware requirements listed below are the minimum needed to meet the needs of the BI solution envisioned for our scenario.

1. Processors

Our hardware solution should have at least 16 processors. This would give us sufficient processing power using the faster microchips of today. It should be field upgradeable to accept more processors if needed in the future.

2. Memory

The main memory support should be at least 12 gigabytes. This ensures that there is enough memory for the processors to run the database supporting the data warehouse, the MDDB, and Data Mining applications. It should be field upgradeable to accept more memory if needed in the future.

IV. Marketplace Solutions

Many solutions are available in the marketplace that meet the Hardware Requirements. The basic components for the BI solution are shown in Figure 5 – Solution Hardware.

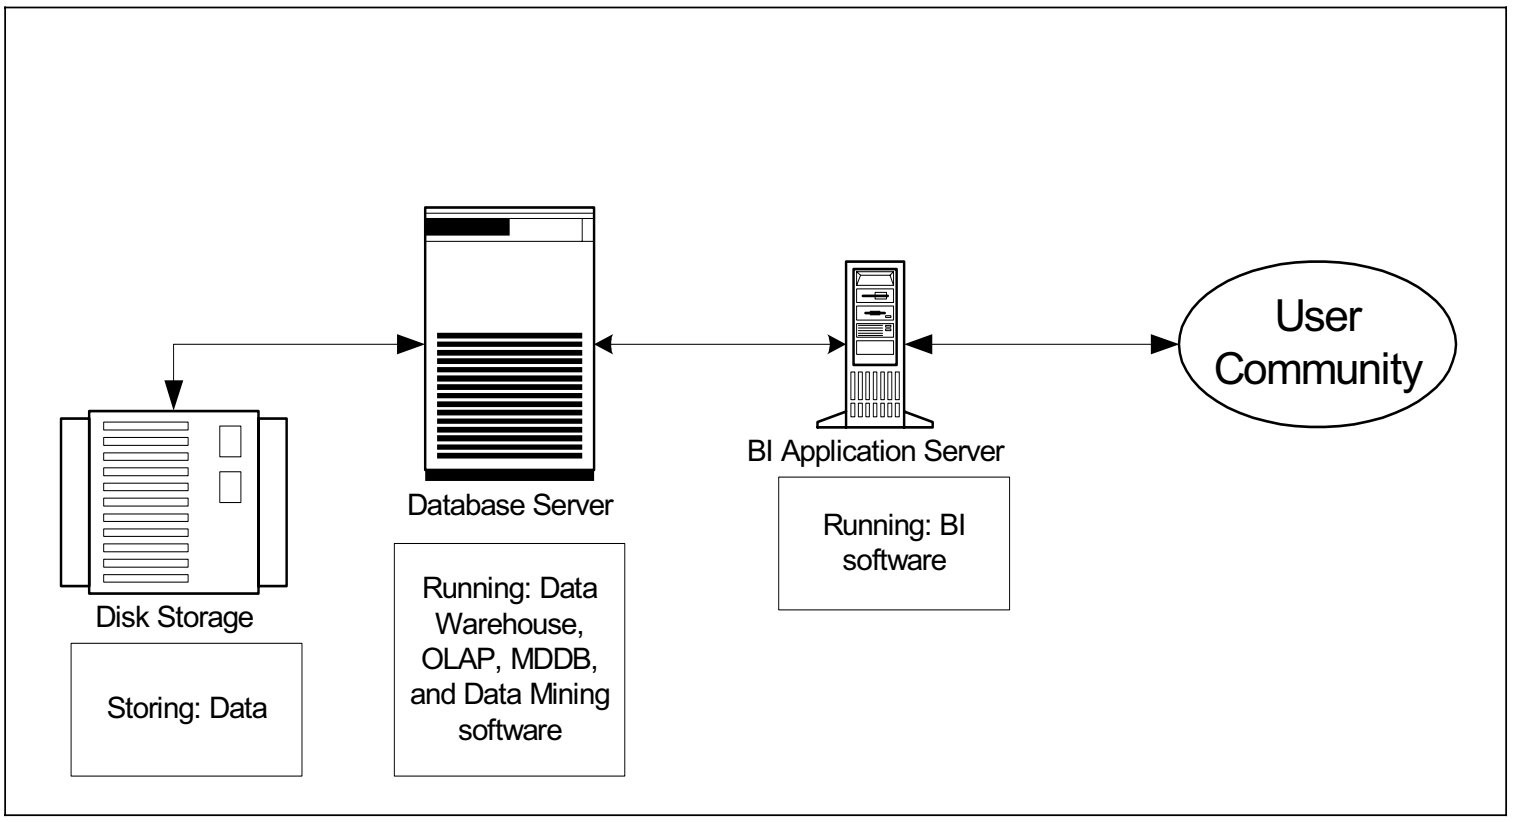

Figure 5 – Solution Hardware

This paper will examine two solutions that meet the Business Needs and the Hardware Requirements. The hardware components for these solutions provide comparable performance under the Configuration Assumptions, permitting a reasonable cost comparison. We have selected configurations that are capable of delivering a BI solution in a VLDB environment.

A. *Sun/Oracle*

One of the most recognized names in the high-end, non-mainframe server market is Sun Microsystems (Sun). Sun is one of the market leaders in Reduced Instruction Set Computing (RISC)/UNIX systems. RISC was developed in the early 1990s as a way to improve the speed of the microprocessor. All Sun machines use the RISC architecture and the Solaris operating system (their version of UNIX). For our scenario, the Sun Fire 6800 Server with 16 processors was selected to provide the best solution. The Sun server would be running Solaris operating system (OS).

We selected an Oracle Database Management System (DBMS) on this machine, since Oracle is the market leader of DBMS on the UNIX platform. Oracle 9*i* is the latest version of their DBMS.

B. *Unisys/Microsoft SQL Server 2000*

The Unisys e-@ction Enterprise Server ES7000 provides an equivalent platform in terms of its hardware configuration in a BI deployment in the Windows/Intel environment. Windows/Intel machines use Intel processors and a Microsoft Windows operating system.

The ES7000 would be running Microsoft Windows 2000 Datacenter Server as its OS. Microsoft SQL Server 2000 Enterprise Edition will be deployed as the database.

C. *Disk Storage*

All of the data for the database will be stored on an external drive array, which is the separate machine labeled as Disk Storage in Figure 5. This type of external storage can grow to meet the needs of either solution.

D. *BI Application Server*

The user community wants to create their own reports. This can be done using a BI application. The BI application would reside on the BI Application Server, which bridges the gap between the database and the user community. The BI application must be able to operate on the company intranet and access the data warehouse and the MDDB.

E. *MDDB Software*

In order to meet the Business Needs, additional software might need to be purchased. The user community wants to be able to create forecasts, budgets, and complex financial calculations, and we want to create a MDDB to help with performance. In order to meet this requirement, an MDDB solution will need to be included in our costs.

F. *Data Mining*

The user community is also interested in finding patterns in the sales data. This type of function is called Data Mining. Data Mining is the process of reviewing the data using statistical methods that allows patterns in the data to become evident. In order to meet this requirement, Data Mining software will also be needed.

V. Overview of Total Cost of Ownership

Investment in an SMP architecture along with a VLDB is an expensive proposition. Beyond the initial costs, there are operational and managerial costs once deployed. The calculation of the initial costs and ongoing costs over the life cycle of a solution is called Total Cost of Ownership (TCO).

Many companies use TCO models to understand the budgeted and unbudgeted costs associated with an IT solution throughout its useful life. Several TCO models are available. Rather than select one, we used the underlining concepts of all of them.

This section will outline the basic structure of the TCO model, and show how some of its characteristics were used as a guide for our analysis.

A. *The TCO Model*

The basic TCO model divides the costs of IT solutions into two different groups: Direct (budgeted) and Indirect (unbudgeted). Direct costs attempt to capture all the expenses related to the IT solution. The Indirect costs capture the efficiency of the delivered solution. Indirect costs are harder to calculate. For our purposes, we used only the high-level cost categories of the basic TCO model.

These groups are further divided into categories with their corresponding subcomponents. The Direct categories are:

- Capital Expenses – Hardware and Software costs
- Administration – Cost of Operating System Management, Application Management, Database Management
- Technical Support – Internal Help Desk, Developers, Other Support Staff

The Indirect category is End-User Operations, which consists of user self support and downtime.

We only captured the Direct costs for this paper, since they can be obtained using public information. Indirect costs can only be drawn from actual experience. No actual experience could be developed since this is a hypothetical situation. TCO models also consider these costs through the entire life cycle of the IT solution. The stages of the life cycle are:

- Planning – initial planning and development of a solution
- Acquisition – purchase of the hardware and software and hiring of staff to administer the solution
- Deployment – execution of the plan with the hardware and software by the staff
- Operation – ongoing maintenance of the solution
- Replacement – replacement of the solution

B. *Application of TCO Model for Analysis*

We assume that the Planning portion of the solution is cost neutral, so we did not include any costs from that part of the life cycle. We focused our attention on the Direct costs so that we could present the most accurate findings. The Direct costs were acquired from the companies' Web sites and other public sources, therefore no discounts or other offerings were considered. We used this TCO model to compare our selected Market Solutions. Any formulas for calculating our findings can be found in the Formulas Section of this paper.

VI. Solution Economics

Now that we have identified the solutions and the methodology to capture the costs of the solutions, we can apply the actual costs from all the available sources. All of our sources are identified in the Resource Section of this paper. This section will explore the costs by TCO categories for the Sun Fire 6800 and the Unisys ES7000 environments.

A. *Capital Expenses*

1. Hardware

a. *Equipment Costs*

The initial equipment costs were derived from pricing generally available on the Internet. They include the Operating System (OS) and the hardware for that particular machine. The Sun machine will operate using Solaris 8. The Unisys machine will operate using Windows 2000 Datacenter Server (DCS).

Server with Operating System*	# Of Processors	Costs
Sun Fire 6800 Server/Solaris	16	$787,245
Unisys ES7000 /Windows 2000 DCS	16	$406,925
Difference		$380,320

Table 2 – Initial Hardware Cost

*details of the hardware can be found in the Hardware Profile Section of this paper.

According to these figures, the Sun Fire 6800 Server solution is $380,320 more than the Unisys solution.

b. *Infrastructure Costs*

Under our schema, we have an additional commodity server for the BI application. Since this can serve both Sun and Unisys servers, there is no cost difference so we will not include it in our calculation.

c. *Disk Storage*

Under our schema, we need at least six terabytes of storage. We could accomplish this by attaching the database server to a disk array. A disk array is a machine that contains enough space to hold many hard drives for storage. Both Sun and Unisys can use a disk array from a third-party vendor. Since this item is interchangeable, there is no cost difference so we will not include it in our calculation.

2. Software

For our scenarios, we assumed that both the DBMS can handle large amounts of data. Based on the Business Needs, there is some additional software to purchase. DBMS vendors offer different functionality in their bundled software. In order to make the comparison equal, and to achieve the functionality needed, any required software will be itemized and its cost captured in separate sections.

The licensing of software is based on one of two models, Named User or Per Processor. Our scenario calls for 1,000 named users with 16 processors. In order to get the best price, we placed 1,000 named users and 16 processors into the respective pricing models from both Oracle and Microsoft. The model that produced the lowest cost for each company is the cost that we compare.

a. *DBMS License*

We determined that the best pricing model would be Per Processor. Oracle offers licensing terms that limit the time the software can be used. They offer 3 plans: 2-year, 4-year, and perpetual. Microsoft only offers perpetual licenses. The cost for a perpetual license for each DBMS is listed in Table 3 below:

DBMS	Cost/Processor	# Of Processors	Annual Costs
Oracle Enterprise Edition	$40,000	16.0	$640,000
SQL Server 2000 Enterprise Edition	$19,999	16.0	$319,984
Difference			$320,016

Table 3 – DBMS License Costs

Based on these figures, the Oracle Enterprise Edition costs $320,016 more to purchase than SQL Server 2000 Enterprise Edition.

b. *MDDB Tool*

Creating forecasts, budgets, and complex financial calculations is an excellent application for an MDDB. We can also use the MDDB to reduce the time it takes for the user community to get the information they want. SQL Server 2000's Analysis Services has the MDDB built into it. For the Sun/Oracle solution, we must purchase the software. Oracle offers Express Service to build MDDBs.

The cost for this additional software is listed in Table 4 below:

DBMS	Cost/Processor	# Of Processors	Final Cost
Express Services	$24,000	16.0	$384,000
Analysis Services	Included	Included	$0
Difference			$384,000

Table 4 – MDDB Software Cost

The MDDB software purchase reflects an additional software cost of $384,000 for the Sun/Oracle solution.

c. *Data Mining*

Microsoft SQL Server 2000's Analysis Services has Data Mining capabilities built into it. This function is sold separately by Oracle as Oracle Data Mining and must be separately licensed. The cost for this additional software is listed in Table 5 below.

DBMS	Cost/Processor	# Of Processors	Final Cost
Oracle Data Mining	$20,000	16.0	$320,000
SQL Server 2000 Data Mining	Included	Included	$0
Difference			$320,000

Table 5 – Data Mining Software Cost

The Data Mining software purchase reflects an additional software cost of $320,000 for the Sun/Oracle solution.

d. *Maintenance Agreements (Vendor Support)*

The DBMS, Hardware, and OS vendors offer maintenance on their products that allows them to be contacted if our internal support team cannot solve a problem related to their products. These agreements must be renewed each year. The prices for the agreements are listed below.

i. Maintenance on OS and Hardware

Sun is the manufacturer of the hardware and the OS for the Sun/Oracle solution. They offer a maintenance agreement that covers the hardware and OS. The Unisys solution is the combination of two different vendors, Unisys and Microsoft. Unisys offers a maintenance agreement that supports both hardware and OS.

Both the Sun and Unisys agreements allow them to be contacted for help with any problems associated with the OS or the hardware. This price also

includes any updates that are available for the OS. This agreement must be renewed every year in order to continue the upgrade service. The annual price is listed below. For our scenario, we assumed these costs would remain the same for each year.

Hardware	Agreement Cost per Year	Operating System	Agreement Cost per Year	Totals
Sun Fire 6800	$15,708	**Solaris**	Included	$15,708
ES7000 Server	$12,480	**Windows 2000 DCS**	$31,044	$43,524
Difference	$3,228	***Difference***	$31,044	$27,816

Table 6 – Maintenance Agreement for hardware and OS

These figures demonstrate that the Sun/Oracle solution is $27,816 less per year.

ii. Maintenance on DBMS, MDDB, and Data Mining Software

Oracle maintenance agreements are priced based on the DBMS pricing model, which in this case is Per Processor. The Oracle agreement must be renewed each year. For our scenario, we wanted to renew it for the life cycle we were measuring.

Microsoft supports SQL Server 2000 on a per incident basis through its Microsoft Professional Support service. The cost for this service is $245 per incident. Based on our experience, we considered 2 incidents per month more than reasonable (24 incidents per year). We used this assumption for every year that was measured. The maintenance agreement cost for one year is listed below in Table 7a for Oracle 9*i*, and Table 7b for Microsoft SQL Server 2000:

DBMS	Maintenance Cost/Processor	Processors	Total
Oracle Enterprise Edition	$2,800	16	$44,800
Oracle Express + Data Mining	$4,200	16	$67,200
Oracle Total			$112,000

Table 7a – Oracle 9*i* Maintenance Agreement Cost

DBMS	Maintenance Cost/Incident	Incidents per Year	Total
SQL Server 2000 Enterprise Edition	$245	24	$5,880
SQL Server Total			$5,880
Difference			$106,120

Table 7b – Microsoft SQL Server 2000 Maintenance Agreement Cost

According to these figures, if all 24 incidents occur, the Oracle agreement would be at least $106,120 more than the Microsoft SQL Server 2000 environment per year.

e. Updates

Once deployed, there are agreements that allow the OS and DBMS software to be upgraded to the latest versions so the solution continues to perform at the optimum level thought out its life cycle.

i. Updates to OS Software

The update contract costs for the OS is captured in the Maintenance cost in Table 6 above.

ii. Updates to DBMS, MDDB, and Data Mining Software

These agreements are priced based on the DBMS pricing model for both Oracle and Microsoft. The Per Processor pricing model is used for both. Oracle requires the update agreement be renewed each year. Microsoft's update agreement must be renewed every 2 years. In order to get an annual figure for comparison, the Microsoft cost was divided in half. For the Per Processor option, the cost for the upgrade agreement for one year is listed below in Table 8a for Oracle 9*i* and in Table 8b for Microsoft SQL Server 2000.

DBMS	Update/Processor	Processors	Total	Annual Total
Oracle Enterprise Edition	$6,000	16	$96,000	$96,000
Oracle Express + Data Mining	$9,000	16	$144,000	$144,000
Oracle Total	$15,000			$240,000

Table 8a – Oracle 9*i* Update Agreement Cost

DBMS	Update/Processor	Processors	Bi-Annual Total	Annual Total
SQL Server 2000Enterprise Edition	$9,564	16	$153,024	$76,512
Difference				$163,488

Table 8b – Microsoft SQL Server 2000 Update Agreement Cost

According to these figures, the Oracle update agreement is $163,488 more than the Microsoft SQL Server 2000 per year.

f. BI Application

A variety of third-party BI applications are available on the market today that meet our Business Needs and User Interface requirements. Any one of these BI

applications would be able to interface with either vendor's solution. Since this application is independent of the solutions, there is no cost difference, so we will not include it in our calculation.

g. *Upgrades to the BI Application Server*

Our schema includes the BI Application Server. This server would need to have its software upgraded and maintained. However, since this would impact each solution equally, there is no price difference, so these costs will not be calculated.

h. *Security Software*

We will not be using any additional security software for the BI application implementation. Any additional security software would cost the same for either platform; there is no cost difference so we will not include it in our calculation.

i. *Software Total*

All of the items listed in the above sections need to be purchased in order to compare the software packages with equal functionality. The total of all the software with their upgrade and maintenance contracts in the first year are listed in Table 9 below.

Oracle	Cost/Processor	SQL Server 2000	Cost/Processor
DBMS	$640,000	**DBMS**	$319,984
Oracle Express + Data Mining	$704,000	**Analysis Server + Data Mining + OLAP**	$0
Software Costs	$1,344,000	**Software Costs**	$319,984
Maintenance (DBMS + Oracle Express +Data Mining	$112,000	**Maintenance (DBMS + Analysis Server)**	$5,880
Upgrades (DBMS + Oracle Express +Data Mining	$240,000	**Upgrades (DBMS + Analysis Server)**	$153,024
Maintenance and Upgrade Costs	$352,000	**Maintenance Costs**	$158,904
Total Costs	$1,696,000	**Total Costs**	$478,888
Difference:	**$1,217,112**		

Table 9 – Database Software Costs

Achieving the required functionality would require an additional cost of $1,217,112 for the Sun/Oracle solution in the first year. The maintenance and update agreements are charges that would continue through the measuring period.

3. Setting Up the User Environment

Once the hardware and software have been installed, the databases and access to the databases must be created. This process takes place in the beginning of the project. The staff members gather the business requirements, create the databases, maintain the environment, and develop programs. Some of the staff would continue after deployment to keep the solution operational.

We determined the salaries by capturing salary information from a variety of public sources. (See the Formula Section for more details on our salary calculations.) We then applied our own experience to this information in order to determine an average base salary for each position. This section takes the base salary and applies it to the number of staff members needed based on our own staffing experience. Any difference in staffing costs for our BI solution will be discussed in this section.

a. Database Modeling

A Data Modeler and Business Analyst will be needed to bridge the gap between the user community and the rest of the Staff. In fact, most of the up-front project time will be used by these Staff members as they define the business objectives and solutions. The Business Analyst is responsible for identifying the business specifications that the BI solution is to provide. The Data Modeler outlines the structure of database according to these business specifications from the Business Analyst. Since Data Modelers and Business Analysts are able to create their output independently of the DBMS, and they do not maintain anything after deployment, there is no cost difference, so we will not include them in our calculation.

b. Extraction, Translation, and Load Developers

All BI implementations require some form of application development to ensure that the data gets from the transactional systems to the data warehouse. These developers are responsible for the extraction of data from the transactional systems. This data is then translated and loaded into the database structure created by the Data Modeler. This function can be performed by the developer using a third-party vendor tool. Since this tool could function in either environment, there would be no cost difference, so we will not include it in our calculations.

c. *Application Development*

The user community will want to create ad hoc reports that are not part of the pre-defined environment. Ad hoc reports can be complex, requiring a developer to create custom programs to manipulate data and deliver these reports. Since the user community will have access to the data through a BI application, the number of ad hoc reports should be rather small. Our experience suggests that one developer could meet the demand.

Our investigation of the average annual base salary for the Microsoft SQL Server 2000 programmer showed it to be slightly higher than for the Oracle programmer. We believe this difference is due to the skill set needed for the Microsoft SQL Server 2000 programmer. In many cases, this skill set would include popular programming languages such as Visual Basic and Java (used for Web development). These additional languages make the Microsoft SQL Server 2000 developer more versatile; therefore there is more value added and a higher salary to reflect this added value.

We needed to factor in the annual salaries for this developer to create any programs that would deliver these reports. Table 10 represents these costs.

DBMS	Avg. Annual Base Salary	Human Resources	Annual Costs
Oracle 9*i*	$70,907	1.0	$70,907
SQL Server 2000	$77,407	1.0	$77,407
Difference			$6,500

Table 10 – Application Developer Cost

According to these figures, the Oracle environment is $6,500 less/year than the SQL Server environment.

B. *Administrative Expenses*

1. OS Administrator

The OS Administrator is an ongoing cost for a BI endeavor. Any system requires some form of human intervention on a regular basis. An OS Administrator is responsible for installing software releases, operating system upgrades, evaluating and installing patches to the OS, and resolving OS related problems. We believe that the number of administrators needed is equal in both environments. In our experience only one full time OS administrator was needed. As a backup we will have a person in the split role

of both OS administrator and DBA. They will be represented as .5 of an administrator. This would make the total number of administrators 1.5 for each solution.

Our examination of the annual base salary showed that administrators for the Solaris OS were better paid than Windows 2000 administrators. We attribute this difference to the market demand for Solaris administrators. There is a smaller supply of Solaris administrators in the market, so their salaries are in higher to reflect this demand. Based on current average annual base salaries, the cost for one year of service is listed in Table 11 below:

Operating System	Avg. Annual Base Salary	Human Resources	Annual Costs
Solaris	$78,144	1.5	$117,217
Windows 2000 DCS	$72,613	1.5	$108,920
Difference			$8,297

Table 11 – OS Administrator Cost

According to these figures, the Solaris environment is $8,297 more per year than the Windows 2000 DCS environment.

2. Database Administration

Some of the activities of the Database Administrators (DBAs) are create/update/delete tables, query tuning, backup/recovery of data, and assistance during the data extraction. The number of DBAs needed to manage an environment can be linked to the size of the database.

According to generally available data, 3 to 4 Oracle DBAs are optimum for each terabyte (TB) of data. It is our opinion that this quantity is a little extreme for an Oracle database. We have developed Table 12 from our own experience as our guide to the number of Oracle and SQL Server 2000 DBAs needed to manage the DBMS.

We believe more Oracle DBAs are needed due to the amount of manual processes that are required in maintaining an Oracle database. These processes include query tuning, backup and restore, and sizing the amount of space needed by the database. These processes are automated in a Microsoft SQL Server 2000 environment; therefore the number of people needed to perform them is reduced.

DBMS	Size of DB in TB	Number of DBAs Needed
Oracle 9*i*	1-3	2
Oracle 9*i*	>3	4
SQL Server 2000	1-3	1
SQL Server 2000	>3	2

Table 12 – DBAs per TB of Data

The Oracle DBMS requires 2 DBAs according to Table 12. We used 2.5 for our calculation to represent the split role administrator mentioned in the OS Administrator section.

Our experience has demonstrated that SQL Server requires 1 DBA. However, to incorporate the split role administrator mentioned in the OS Administrator section, we set the number of SQL Server 2000 DBAs at 1.5. We used these figures to calculate Table 13, which contains the costs for years 1 and 2.

DBMS	Avg. Annual Base Salary	Human Resources	Annual Costs
Oracle 9*i*	$84,522	2.5	$211,306
SQL Server 2000	$80,522	1.5	$120,784
Difference			$90,522

Table 13 – DBA Cost Years 1-2

Initially, the cost difference is $90,522 more for the Oracle *9i* solution, however the database will grow by one terabyte per year. This growth will require the addition of 2 Oracle DBAs and 1 SQL Server 2000 DBA in year 3. Table 14 illustrates the cost difference when the database grows over three terabytes.

DBMS	Avg. Annual Salary	Human Resources	Annual Costs
Oracle 9*i*	$84,522	4.5	$380,351
SQL Server 2000	$80,522	2.5	$201,306
Difference			$179,045

Table 14 – DBA Cost Years 3-5

For years 3-5, the Oracle solution is $179,045 more than the Unisys/Microsoft SQL Server 2000 solution.

VII. Conclusions

A. *Overall TCO Savings*

Now that all of the costs have been identified, we can look at the total cost for the entire lifetime of the solutions. Table 15 represents the 5-year Pro Forma cost sheet for the Sun/Oracle solution. Note that the OS Maintenance/Update agreement is $0. The cost for this agreement is included in the Hardware Agreement cost.

5-Year Pro Forma Cost Sheet - Sun/Oracle Solution					
Cost Category	**Year 1**	**Year 2**	**Year 3**	**Year 4**	**Year 5**
Hardware	$787,245	$0	$0	$0	$0
Hardware Agreement	$15,708	$15,708	$15,708	$15,708	$15,708
DBMS	$640,000	$0	$0	$0	$0
MDDB	$384,000	$0	$0	$0	$0
Data Mining	$320,000	$0	$0	$0	$0
OS Maintenance/Update Agreement (Solaris)	$0	$0	$0	$0	$0
Database Maintenance/Update Agreements (Oracle)	$352,000	$352,000	$352,000	$352,000	$352,000
Application Developer	$70,907	$70,907	$70,907	$70,907	$70,907
OS Administrator	$117,217	$117,217	$117,217	$117,217	$117,217
DBA	$211,306	$211,306	$380,351	$380,351	$380,351
Total Cost By Year	$2,898,383	$767,138	$936,182	$936,182	$936,182
				TCO	**$6,474,066**

Table 15 – Sun/Oracle 9*i* TCO

5-Year Pro Forma Cost Sheet - Unisys/SQL Server 2000 Solution					
Cost Category	**Year 1**	**Year 2**	**Year 3**	**Year 4**	**Year 5**
Hardware	$406,925	$0	$0	$0	$0
Hardware Agreement	$12,480	$12,480	$12,480	$12,480	$12,480
DBMS	$319,984	$0	$0	$0	$0
MDDB	$0	$0	$0	$0	$0
Data Mining	$0	$0	$0	$0	$0
OS Maintenance/Update Agreements (DCS)	$31,044	$31,044	$31,044	$31,044	$31,044
Database Maintenance/Update Agreements (SQL Server 2000)	$158,904	$5,880	$158,904	$5,880	$158,904
Application Developer	$77,407	$77,407	$77,407	$77,407	$77,407
OS Administrator	$108,920	$108,920	$108,920	$108,920	$108,920
DBA	$120,784	$120,784	$201,306	$201,306	$201,306
Total Cost By Year	$1,236,447	$356,514	$590,061	$437,037	$590,061
				TCO	**$3,210,119**

Table 16 – Unisys/SQL Server 2000 TCO

Table 16 represents the Unisys/SQL Server 2000 TCO for the same categories. Please note that the Database Update Agreement is renewed every two years, so its cost only appears in years 1, 3, and 5.

These figures clearly demonstrate that the Sun/Oracle solution costs $3,263,947 more than the Unisys/Microsoft SQL Server 2000 solution. That is a 50.42% cost difference in favor of the Unisys/Microsoft SQL Server 2000 solution.

B. *Costs Comparison by Category*

When the categories are placed next to each other, the differences are easier to see. The graph below demonstrates that the Sun/Oracle solution costs more in every category over the life cycle. Figure 6 – 5-Year Pro Forma Cost by Category represents the cost differences by category for each solution.

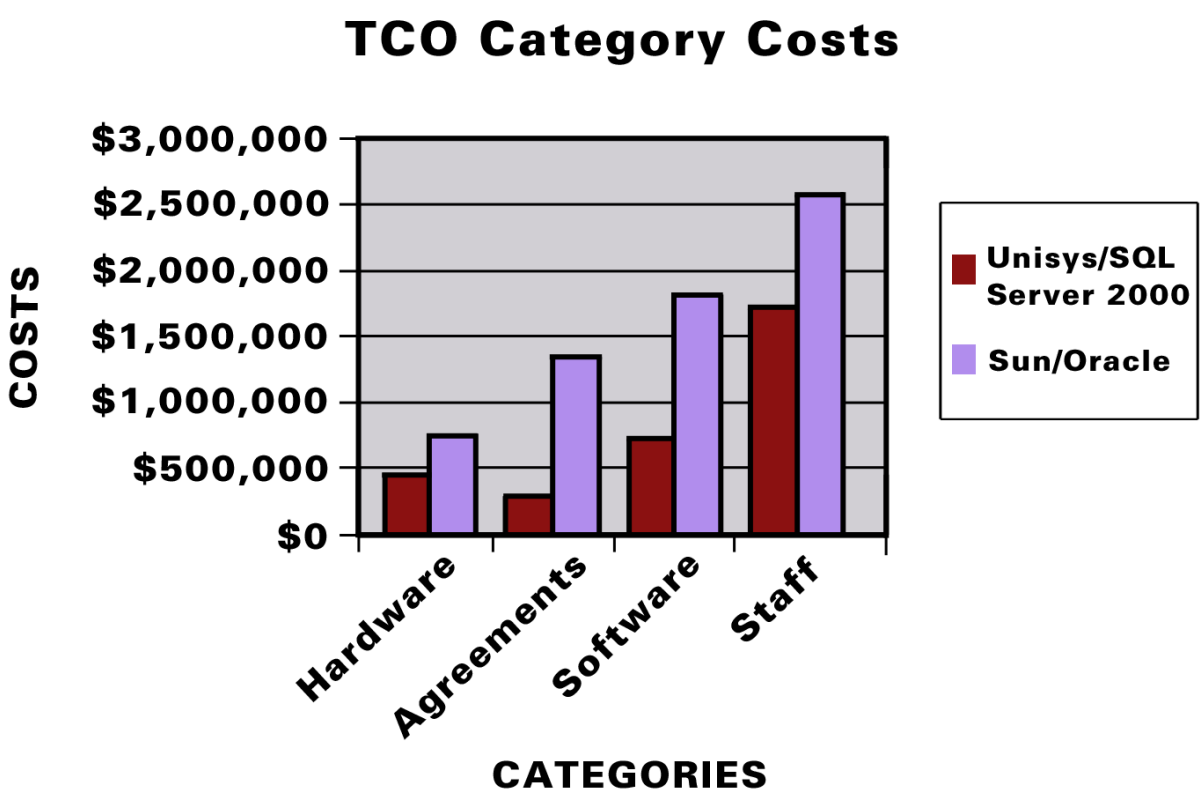

Figure 6 – 5-Year Pro Forma Cost by Category

Based on this information, we can conclude that the Unisys ES7000 with Microsoft SQL Server 2000 is significantly less expensive to operate over 5 years.

This paper looked at two alternatives offered by some of the most recognized names in IT today. What we knew was that the cost of implementing a BI solution in a VLDB environment is an expensive proposition. But what we discovered is that the current state of Windows/Intel solutions offers considerable savings and an attractive alternative to the classic RISC/UNIX solutions for a BI implementation.

VIII. Appendix

A. *Hardware Profiles*

1. Sun

The Sun Fire 6800 consists of the following:

Description	Qty	Ext List Price
Sun Fire Server 6800 Rackmount (28U) includes 1 x system controller, 2 x power supplies, 2 x fan trays	1	$150,000
Sun Fire Server CPU/MEM Board Bundle includes 4 x UltraSPARC-III 900MHz processors with 4 licenses, 4 x 1GB memory, 4 x 8MB Ecache,	4	$588,000
PCI I/O Assembly for Sun Fire 4800, 4810, 6800 servers	4	$32,000
StorEdge D240 Media Tray - Rackmount with two disk drives, one DVD-ROM and one tape drive (2 x 18.2GB disk drives, 1 x DVD-ROM, 1 x DDS-4 tape drive)	1	$5,800
PCI – 10/100BaseT F/W UltraSCSI PCI Adapter	1	$995
Sun Fire Server cabinet includes 2 x rack fan trays, 1 RTU with 1 x RTS	1	$10,450
Warranty Upgrade to SunSpectrum Platinum	4	$6,336
Warranty Upgrade to SunSpectrum Platinum	1	$9,372
Total		**$802,953**

2. Unisys

The Unisys e-@ction Enterprise Server ES7000 consists of the following:

Description	Qty	Ext List Price
ES7000 Enterprise Server with 16 Intel Pentium III Xeon 900MHz CPUs 2MB L2 cache	1	$251,200
Module, PCI Enclosure 3.3V, 4 Card	2	$800
Dual Channel SCSI controller	4	$4,400
LP8000-F1 fiber controllers for ESM7900	2	$4,000
GB Ethernet Optical Adapter	2	$1,700
36 GB RAID1 SCSI disk drives for boot drives	1	$3,850
Module, PCI Enclosure w/BID Motherboard	2	$6,600
Memory 8GB, 1 GB DIMMs	1	$48,000
Memory 4GB, 1 GB DIMMs	1	$26,000
ES7000 Value Add Software	1	$5,000
ES7000 Value Add Software - Upgrade	1	$0
ES7000 Call Home Feature	1	$0
Keyboard, Mouse, Monitor	1	$375
MS Windows 2000 Datacenter Server Operating System	1	$55,000
TOTAL:		**$406,925**

B. *Formulas*

The salary calculations were developed using a variety of public resources. First we determined the average salary. This was calculated using the median salary for each position from computerjob.com from three different cities: Los Angeles, Chicago, and New York. The results of these calculations are listed in Table 17 below.

Occupation	Median LA	Median CHI	Median NY	Avg. Salary
Senior System Administrator	$80,261	$78,105	$82,992	$80,453
DBA	$81,894	$79,695	$84,680	$82,090
Programmer III	$74,450	$72,450	$76,982	$74,627

Table 17- Avg. Salary

We then located the specific occupation on computerjobs.com and monster.com and selected the highest published salary we could find from a location other than Los Angeles, Chicago or New York. We then applied our own salary knowledge to form Table 18 below.

Occupation	Avg. Salary	Computerjobs.com	Monster.com	Walklett Experience	Avg. Annual Base Salary
Unix System Administrator	$80,453	$70,000	$87,125	$75,000	$78,144
Datacenter Administrator	$80,453	$75,000	$65,000	$70,000	$72,613
Oracle DBA	$82,090	$85,000	$85,000	$86,000	$84,522
SQL Server DBA	$82,090	$80,000	$80,000	$80,000	$80,522
Oracle Programmer	$74,627	$69,000	$70,000	$70,000	$70,907
SQL Server Programmer	$74,627	$80,000	$75,000	$80,000	$77,407

Table 18 – Avg. Annual Base Salary

The result of Table 18 is the average annual base salary. The average annual base salary was used to calculate all the salary information for this white paper.

C. *Definitions*

The definitions listed below are only applicable to this white paper. They are intended to define certain terms used in the context of this white paper.

Term	Definition
Architecture	The methodology and components used to build the hardware.
Base Salary	Calculation of a national average salary with specific salary information to determine an approximate salary for an occupation.
Data	Facts as a result of business activities.
Gigabyte	1,024 megabytes
Hardware	The equipment used to hold data and generate information.
HOLAP	Abbreviation for Hybrid Online Analytical Processing. This MDDB format utilizes the best of ROLAP and MOLAP and can use pre-calculated data sets or subsets of tables to respond to requests for data.
Information	The results of analyzing data.
Interface	Application that allows the user to work with the program.
Intranet	A collection of computers connected together within the same company.
MDDB	Abbreviation for Multi-Dimensional Database. An MDDB stores regularly requested subsets of data from a data warehouse in a format that makes accessing the data easy for the user community. By storing these subsets of data, the overall time to get data to the user community is reduced.
Marginal Cost	Difference in the costs between the two solutions.
MOLAP	Abbreviation for Multi-Dimensional Online Analytical Processing. This MDDB format utilizes a pre-calculated set of data that contains all the possible answers to business questions.
OLAP	Abbreviation for Online Analytical Processing. It is technology that enables business users to examine transactional data to gain insight into the condition of the business.
Processors	Microchips inside the computer that provide computational function.
ROLAP	Abbreviation for Relational Online Analytical Processing. This MDDB format utilizes a subset of the tables arranged so that any question could be asked of the database, and the resulting data set can be re-queried for more detailed data (AKA drill down).
Software	The programs that reside on the hardware that allows the hardware to function.
Terabyte	1,024 gigabytes or 1,048,576 megabytes
Vendors	Companies who manufacture and sell their hardware and/or software.

IX. Resources

Unisys ES7000 information
http://www.unisys.com/hw/servers/es7000/
Sun Server information:
http://www.sun.com/servers/;$sessionid$PHDMDNKFHZ3JVAMTA1LU5YQ
Oracle information:
http://www.oracle.com/
SQL Server 2000 information:
www.microsoft.com/sql/default.asp
Server performance information:
www.tpc.org
Salary information was developed using our own experience and these Web resources:
www.salary.com
www.comperterjob.com
www.monster.com
TCO information was found on:
http://www.compaq.com/tco/models.html
www.gartner.com
www.computerworld.com
www.cio.com
www.dmreview.com
www.intelligententerprise.com
www.micromation.com
www.itpapers.com
www.input.com/buyers_guide/yntco/yntco_main.cfm

Walklett Group

The Walklett Group has assisted many Fortune 1000 corporations in developing business-critical information systems since 1985. Our strength is in the strategic planning, design, development, and implementation of solutions using industry-leading technologies, with a focus on enhancing our clients' use of their corporate business intelligence. Over the past few years our Business Intelligence practice has established itself as one of the key providers of successful Business Intelligence solutions in the mid-Atlantic region. Through our relationship with leading experts and vendors, and a focus on excellence, we have developed a successful approach to building quality business intelligence solutions.

As a leading project delivery and project management company, Walklett Group has been designing and developing mission-critical information systems for more than a decade for such key business industries as the financial, chemical, pharmaceutical, manufacturing, and retail industries. Central to Walklett Group's successful delivery of business solutions is our unique ability to solve our customers' business problems by combining state-of-the-art technologies with industry-proven methodologies and highly skilled professionals.