DOLLARS TO CENTS: TCO IN THE TRENCHES 2002

"Money speaks sense in a language all nations understand…"

- Aphra Behn 1640 – 1689

The 1999 TCO in the Trenches research note outlined our first Total Cost of Ownership (TCO) study for high-end enterprise servers. Since that time we have performed and updated over 2,000 cases and changed the name. TCO is now called CENTS (Comparative Economic Normalization Technology Study). CENTS cases are built using data input by users — absolutely no vendor data is used in the model. These cases drive our VirtualADVISOR Cost Assessment Model (CAM) database. None of the cases within the CAM database are over six months old, thus providing the latest cost data for use in building comparative systems.

In this research note the VirtualADVISOR system is used to build application models in order to compare systems. It is important to note that the Virtual ADVISOR system uses three models: Basic, which is the cost to operate systems within the computer room; Application, which is the cost to operate the application, and Downtime, which adds the cost to the application for lost revenue and production during the average yearly outage. In addition, these three models are further segmented into two categories; dedicated server and shared server. A dedicated server is a server that is dedicated to a single application, therefore, all costs are assessed to the system and the application. A shared server is a server that is not dedicated to a single application, therefore, all costs are spread over multiple applications.

As of this writing there are sixteen different system types in the VirtualADVISOR database. In this paper we show the lowest three systems and the highest three systems. The first set of cases used are "order processing" applications. The second set used are "Sales Force Automation" applications. We also used ten transactions per second at peak times and one transaction off peak for both application sets. The peak period is 3,000 hours per year out of 8,760. Using peak and off peak allows us to calculate the number of transactions lost during each period, as well as dedicated versus shared server cost.

Basic Cost - Order Processing

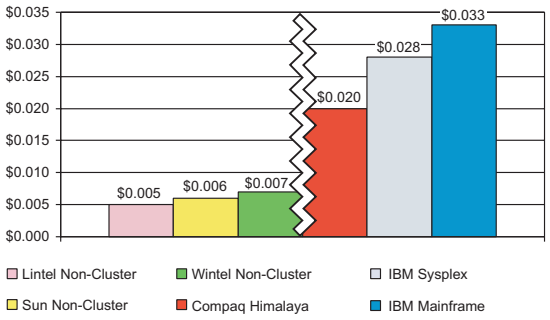

Figure 1.0

All systems, except for the Compaq Himalaya, are running Oracle with a transaction-processing monitor. The Compaq Himalaya uses NonStop SQL and Pathway. To arrive at the cost per transaction, we divided the total cost by the number of transactions that would have been processed during both peak and non-peak operational periods.

The basic cost of using a dedicated server consists of hardware and system software over a three year

open-end lease. Also included are software and hardware maintenance, operational personnel, system personnel, and other costs such as electricity, space, etc. In Figure 1.0 we show the lowest of all costs is Lintel with one half cent to create an order. This is an Intel-based system running a Linux operating system. The next is a Sun non-cluster at just over a half cent with $.006. In third place is the Wintel non-cluster, which is an Intel-based computer running WindowsNT/2000. On the high water mark is the IBM 390 type mainframe at 3.3 cents

Yearly Application Cost - Order Processing
(Per $1,000)

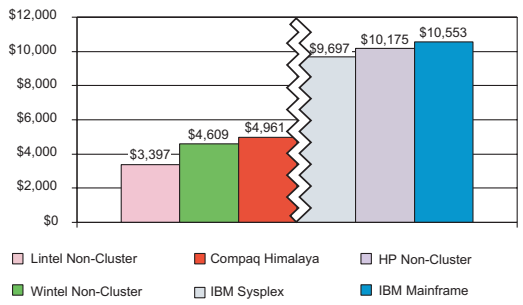

Figure 2.0

per order. This is followed by the IBM Sysplex with just under three cents at $.028. Showing in third to last is the Compaq Himalaya at 2 cents.

In Figure 2.0 we show the total operational cost to operate each given system based on an order processing application running at ten transactions per second at peak period. Again, the Lintel system comes in first with a yearly cost of $3.4 million per year, followed by the Wintel non-cluster. Compaq Himalaya, which was third last in the basic cost, is third best in the application cost – a mere 8% more than the Wintel system. The IBM mainframe costs

Application Cost - Order Processing

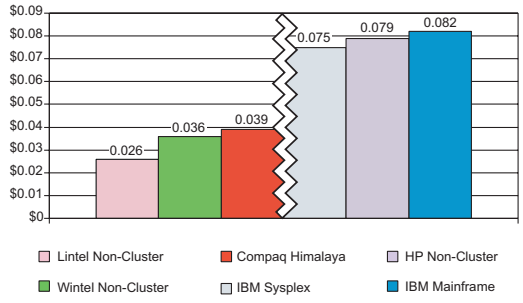

Figure 3.0

three times the Lintel system at $10.5 million, followed closely by HP non-clusters and IBM Sysplex.

Comparing transaction costs again in Figure 3.0, we can see Lintel comes in at 2.6 cents a transaction, Wintel non-cluster at 3.6 cents and the Compaq Himalaya at 3.9 cents per order. On the high side, IBM mainframes are 8.2 cents, HP non-cluster are 7.9 cents followed by IBM Sysplex at 7.5 cents.

Downtime Cost - Order Processing

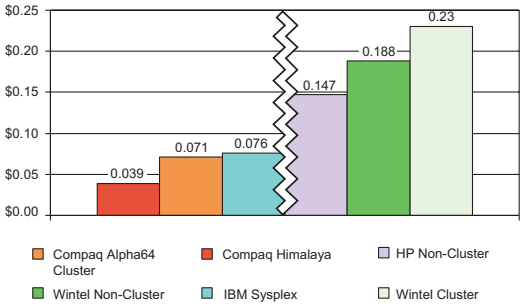

Figure 4.0

It is important to note that for our purposes, a transaction is a complete business deal or a completed unit of work. Therefore, an order and a transaction are the same. They tend to be heavily weighted. This is much different than a message or a piece of a transaction*.

Adding the cost of downtime to a cost model changes things dramatically. The cost of downtime is fairly personal and specific to an application and company. The Standish Group has seen applications valued in the millions of dollars per minute of uptime. In the VirtualADVISOR system, the user

Basic Cost - SFA Server

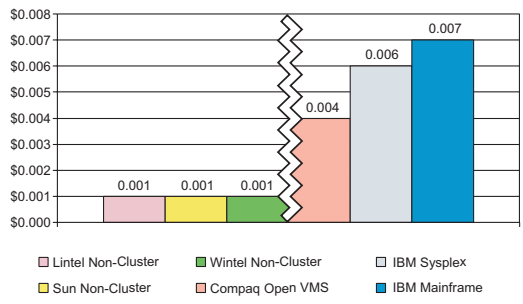

Figure 5.0

can input their cost of downtime by the loss of the ability to execute a transaction. In Figure 4.0 we use $10.00 as an example figure. The VirtualADVISOR system then multiplied the number of transactions lost at peak and non-peak periods based on the downtime reported on a given platform. Doing so shows a reversal of fortunes. The Wintel non-cluster, which previously rated second best, is now second worst. In comparison, the IBM Sysplex went from third worst to third best.

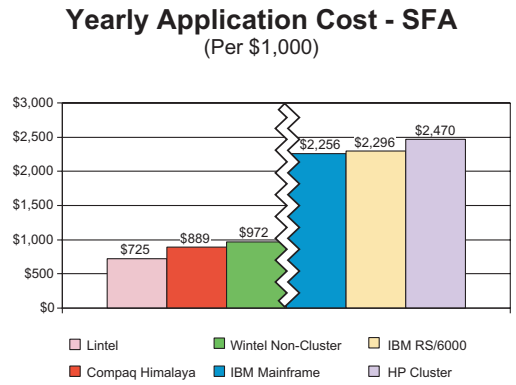

Figure 6.0

The dedicated server model uses calculations based on no other major applications being run on the same system. The shared server, however, looks at the utilization of the server based on multiple applications running on the same system. In 1999 IBM Sysplex respondents claimed an average of almost 44 mission-critical applications ran on their same server. In Figure 5.0 the Sales Force Automation (SFA) applications are used for cost examples. Again, the system executes ten transactions per second at peak and one off-peak. As shown, the yearly cost to operate a system on a Lintel machine is $125,000 or .001 cents. While the Sun non-cluster is

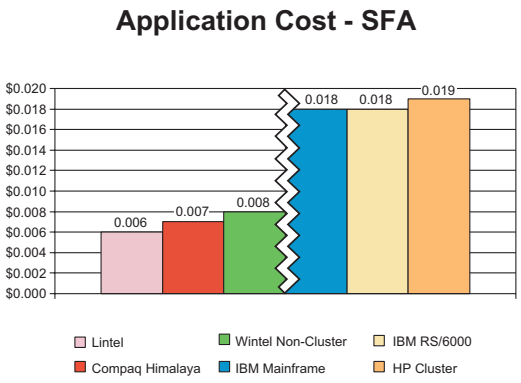

Figure 7.0

next, at a cost of $161,000, it also averages .001 cents. The IBM mainframe is the most costly at $883,000 per year or .007 cents.

In Figure 6.0 we show the total cost to operate each given system based on a SFA application running at ten transactions per second at peak. The Lintel system comes in first with a yearly cost of $725,000 per year, followed by Compaq Himalaya and then the Wintel non-cluster. Dead last on this application is the HP Cluster at almost $2.5 million to support the application – three times the cost of the Lintel system.

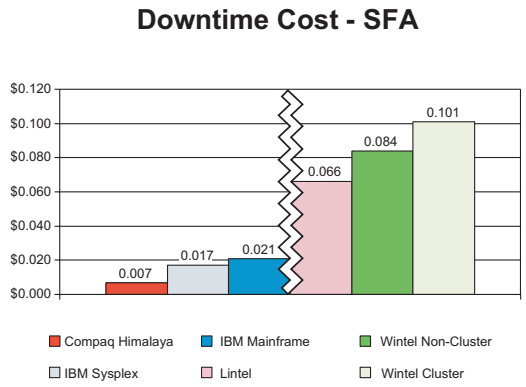

Figure 8.0

One of the most difficult problems with any measurement is to set a level playing field. In order to accomplish this, we established a model transaction. Many of the questions in the case survey are about the transaction detail within the application. The application cost per transaction is the purest cost. This is calculated by taking the total cost to operate the application at peak and dividing by the number of transactions processed during the yearly peak period, then adding the applications at off-peak. We show in Figure 7.0, the lowest three are Lintel, Compaq Himalaya and Wintel Non Cluster. The most costly are HP Clusters, IBM RS/6000 Clusters and IBM Mainframe.

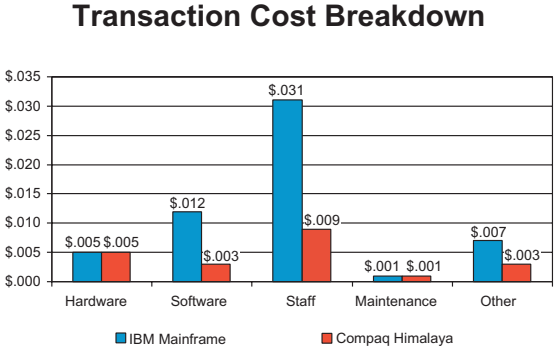

Figure 9.0

There are many causes of downtime and many ways to look at downtime numbers. Most vendors will quote their numbers in relationship to their hardware and software. This is a relatively narrow view. For years, The Standish Group has reported downtime numbers in relationship to applications, not systems. In this research we collected downtime data every month for systems and applications downtime, and break down the causes. In the SFA application we used a cost of $5.00 per lost transaction. In these models we see Compaq Himalaya, the least costly, at half the cost of the next system the IBM Sysplex. Both Wintel models add over $10 million in downtime, driving them to be the most costly.

Summary

By breaking down the cost in each model, it is clear that manpower costs overshadow all costs except downtime. For example, if we look at both the Compaq Himalaya and IBM Mainframe (Figure 9) we see stark differences. Looking at a billing application with the same profile of the other two we see the cost to execute a bill on the IBM Mainframe is 5.6 cents and Compaq Himalaya is 2.1 cents. Of the 2.1 cents per bill on the Compaq Himalaya, only .09 cents (43%) goes to manpower. On the other hand, 3.1 cents of the 5.6 cents (56%) for a bill using the IBM mainframe is manpower cost.

The main purpose of our CENTS research and the VirtualADVISOR is to study and compare operational and cost environments on the more popular systems. The research uses an "If Bought Today" (IBT) theory. IBT works this way: Systems and software are all bought at different times. We asked each user respondent to estimate what they would pay today for a new system with the same capacity and performance. In this way we are measuring the latest "street" cost. To ensure the latest cost estimates, one-sixth of all cases (over 300) are updated every month and no case is older than six months. All data is now collected online. To view the questions go to the Standish User Research Forum (SURF) on our web site, www.standishgroup.com, and use "demo" as the user name and "CENTS" as the password.

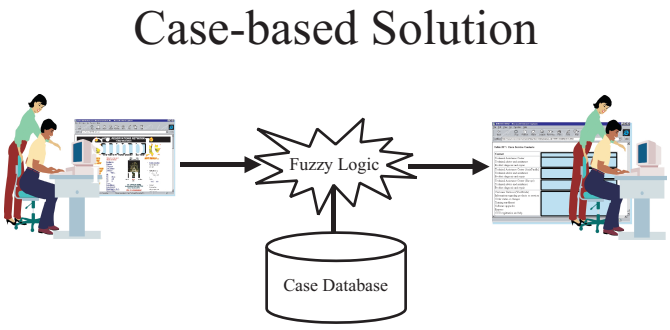

Figure 10.0

Access to the Virtual ADVISOR is available over the Web through our annual subscription service. Users simply input the application they are considering, their transaction rate, hours of peak and off-peak traffic, plus products they wish to review and the VirtualADVISOR does the rest. Users can benchmark their own cost or use the tool for server consolidation as well as server or database selection. For more information on the VirtualADVISOR visit our website.

**We are currently developing a message-based CENTS model, to be released 2nd quarter 2002.*

The Standish Group International, Inc. produces research advice based on extensive primary research in the area of mission-critical applications. The Standish Group provides this advice through our continuous information service. This research service studies the requirements for developing, implementing and maintaining mission-critical applications.

THE STANDISH GROUP

The Standish Group International, Inc. • 196 Old Townhouse Road • West Yarmouth, MA 02673
508-760-3600 • FAX: 508-760-0036 • www.standishgroup.com