Quantifying the Total Cost of Ownership for Entry-Level and Mid-Range Server Clusters

A Detailed Analysis of the Total Cost of Ownership of HP OpenVMS, IBM AIX and Sun Solaris server clusters.

June 2007
Version 1.0

© 2007, TechWise Research, Inc. All rights reserved. No portion of this paper may be reproduced or distributed in any manner without the express written permission of TechWise Research.

Total Cost of Ownership for Server Clusters

Executive Summary

The overall concept of Total Cost of Ownership for server clusters has been an important issue within the IT community for many years. Numerous studies and customer experiences have proven that purchase price alone is not an adequate measurement to compare server clusters from various vendors. Usually other factors, including the costs to manage and maintain the servers, as well as downtime, have a greater financial impact on an organization than just the system's purchase price. Recognizing these factors, TechWise Research developed an analytical approach in 1999 called Reliability-Adjusted Total Cost of Ownership™ that incorporates management costs and application availability in the TCO analysis. As part of this analysis, server clusters from different manufacturers are compared in terms of the actual number of downtime hours per year that customers typically experience. These downtime findings are then converted into a monetary measurement of the cost differences between clusters which TechWise Research refers to as the "Availability Advantage™."

This study focuses on the following three brands of entry-level and mid-range server clusters: HP OpenVMS Integrity, IBM AIX System p5, and Sun Solaris Sun-Fire clusters. TechWise Research last studied these clusters in 2004. Since that time, HP, IBM and Sun have all introduced new server hardware as well as updated operating system and cluster software. These changes all have a direct impact on the Total Cost of Ownership for these clusters. TechWise Research conducted a new study, as a result, to provide more current information on this important topic.

This study involves data from customers as well as publicly available pricing information. TechWise Research surveyed customers that have one of the three types of server clusters installed at their site. The purpose of these interviews was to collect data on the operational costs associated with installing, managing, and maintaining their clusters. Information was also collected on the number of hours and associated costs for various downtime events each company experienced over a twelve-month period. All of these operational data were then combined with current system and service pricing (from IDEAS International) as well as energy costs to calculate the Reliability-Adjusted TCO™ - an analytical approach that factors downtime costs/rates into the TCO analysis. TechWise Research included four main cost components in the TCO analysis. These are the costs to (1) buy the servers and service contract, (2) install and configure the cluster, (3) manage and maintain the cluster over three years (including energy and cooling costs), and, (4) the costs associated with application downtime over three years. Four different cluster configurations from the manufacturers were analyzed. For each configuration, the Reliability-Adjusted TCO™ was calculated at various downtime costs to allow readers the ability to compare the different platforms at a downtime cost rate that most applies for their firm.

Study Results: Five potential causes of downtime were used in the analyses in this paper. When all five causes of downtime are considered, **HP OpenVMS Integrity clusters averaged the fewest total number of annual downtime hours (1.96), followed by IBM AIX System p5 clusters (6.21), then by Sun Solaris Sun-Fire clusters (9.53)**. In terms of availability, the HP, IBM, and Sun clusters are 99.978%, 99.929%, and 99.891% respectively. These differences may seem trivial, but in fact they are not.

The reliability findings were further analyzed to determine hardware and software reliability. Isolating crashes caused by the server hardware itself showed that there is little difference in reliability between HP Integrity and IBM System p5 server hardware. Sun Sun-Fire servers, however, were found to have more than double the downtime of HP or IBM. Isolating crashes caused by operating system and clustering software showed clear differences between the brands. **HP OpenVMS averages only 0.19 hours of downtime due to the operating system and clustering software. This works out to 99.998% availability**. IBM's AIX and HACMP average 4.36 hours (99.950% availability) while Sun's Solaris and Sun Cluster average 4.97 hours (99.943% availability).

Respondents reported that when all costs associated with downtime were considered, including lost sales, lost production and productivity, the average cost per hour of downtime is $145,000. Nearly half of the study's respondents report that their company loses as least $10,000 for each hour their cluster is down. The cost of having a cluster down a few extra hours per year, therefore, can be quite significant. This is why for most companies, management and downtime costs are significant contributors to the TCO for entry-level and mid-range clusters.

The four different cluster configurations analyzed range in price from $100,000 to over $1 million. **For all four configurations studied, HP OpenVMS Integrity clusters offer the lowest TCO *regardless* of the cost associated with an hour of downtime.** This represents a change from past studies where Sun clusters used to offer the lowest TCO when there are no costs associated with downtime. Over the past two years HP has been transitioning its OpenVMS business from AlphaServer to Integrity. One reason for this shift was that Integrity servers offered lower cost solutions thanks to economies of scale. This transition has helped to vault HP to be the clear leader in TCO for entry-level and mid-range servers. In all but Cluster Configuration #4, Sun Solaris Sun-Fire clusters offer lower TCO than IBM AIX System p5 clusters at low downtime hourly rates. Once the cost per hour of downtime exceeds a certain level, IBM's better reliability compared to Sun gives IBM a lower TCO than Sun. Cluster Configuration #4 represents the largest systems studied in this paper. These clusters, which are one step below enterprise-class clusters, have price tags over $1 million. In this configuration, Sun clusters have the highest TCO of the three brands regardless of the cost per hour of downtime.

This study showed that the differences in TCO between these three cluster brands can be quite significant and can easily amount to hundreds of thousands of dollars over three years. **Applying typical hourly rates for downtime costs, HP OpenVMS Integrity cluster's cost between $290,000 and $1.5 million less than IBM AIX System p5 clusters over a three-year period. Similarly, HP's TCO advantage over Sun ranges from $254,000 to $2.8 million over three years**.

The History of TCO

The overall concept of Total Cost of Ownership continues to be an important issue within the IT community, as evidenced by the many published studies on the topic. Some TCO analyses performed today focus only on the prices that manufacturers charge for their clusters and the service contracts purchased with them. Many studies have shown that this simplistic approach can be misleading because it ignores other costs that often exceed the initial system price. TechWise Research has developed a robust approach that accounts for these other costs. It is our belief that this approach provides a more accurate picture of the true total cost of ownership for these server clusters.

Two key costs are often ignored in TCO analysis: management and downtime costs. TechWise has conducted a number of studies over the past nine years which show management costs (the time spent managing and maintaining a server cluster over its lifetime) contribute a significant amount to a cluster's TCO. Most companies' IT staff are stretched thin with many responsibilities. Clusters that require more time to manage on a day-to-day basis take important resources away from other tasks. As important as management costs are in calculating TCO, they are often overshadowed by downtime costs. Clusters by definition are usually deployed in situations where application availability is important. One of the primary functions of a server cluster is to provide high availability through automated failover. In the event that one server crashes, the other servers in the cluster should keep the applications available so that end users are unaffected. The cost per hour of downtime varies considerably for companies. ***It is not uncommon for this cost to be measured in the tens, if not hundreds of thousands of dollars per hour***. Clearly any TCO analysis that ignores actual availability data will be lacking a key cost component. Simply purchasing the lowest cost hardware may be a crucial mistake if that hardware were to be unreliable. TechWise recognized this back in 1999, and as a result developed a proprietary technique called Reliability-Adjusted TCO™. Using this technique, downtime costs have been an integral part of all of our subsequent TCO studies.

Background on This Paper

TechWise Research's most recent study that focused on entry-level and mid-range servers was published in February, 2004. That study focused on RISC-based servers from HP, IBM and Sun. Since that report was published, all three manufacturers have introduced new server models based on new processors and architecture. IBM, for instance, introduced the higher performance dual core POWER5+ processor in 2005. In that same year HP released the first version of OpenVMS to run on its Integrity platform, based on the Intel® Itanium® 2 microprocessor. In 2005 Sun introduced two new processors, the UltraSPARC IV+ and the multicore-multithreaded UltraSPARC T1 processor. All of these new systems offered higher performance at a lower price than the systems analyzed in the 2004 study. Because the findings from the February 2004 report are no longer current, TechWise Research completed this new study to provide updated information on server cluster TCO. For this particular paper, TechWise Research decided to focus on HP's Integrity platform and not AlphaServer, since HP stopped shipping AlphaServers in April, 2007.

TechWise Research's TCO Model

There are four components in TechWise Research Reliability-Adjusted TCO™ model, or TCO for short:

- **Acquisition Costs** – The costs to purchase the hardware and software from the manufacturer along with a three-year support agreement.
- **Start-Up Costs** – The costs to install and configure the cluster, as well as any time and money spent to train staff on the cluster.
- **Management Costs** – The ongoing costs associated with managing and operating the cluster on an annual basis (including energy and cooling costs).
- **Downtime Costs** – The number of hours, and resulting costs for cluster downtime, on an annual basis.

The acquisition costs for the clusters were calculated using current system and service pricing from IDEAS International. IDEAS International is recognized worldwide as a leading authority on systems technology, specializing in the research of comparative information on computer systems. Their current system and service pricing is updated daily with new product and price announcements from server manufacturers. When buying servers, two customers can pay very different prices for two identical systems depending on when they buy them, and on the level of discount they can negotiate from their channel. By using current list prices from IDEAS International, this time and purchasing power bias was eliminated from the TCO analyses.

TechWise Research conducted a market research study to collect information for the other three components. Management costs have two main components: Costs for companies to hire third-parties to manage their cluster on an ongoing basis, and costs for managing the cluster "in-house." In the former situation, respondents provided the actual costs for outsourcing cluster management activities. In the latter case, we collected the number of hours "internal staff" spent on all management activities associated with the cluster. TechWise Research converted internal hours spent into a cost figure by using staff salary data provided by respondents. When calculating internal management costs, the TechWise model factored in the hours spent, if any, managing and maintaining: servers in the cluster, the cluster's storage array, cluster software and operating system, end-user applications, and network permissions. For the first time TechWise included energy costs as part of the cluster's annual management costs. The cost to run and cool servers continues to rise. Cooling costs can be very significant for companies with large server rooms. As it turns out, in this particular study the differences between brands in energy and cooling were negligible and did not affect the overall findings.

Downtime costs were calculated in a multi-step process based on findings from a number of studies TechWise Research has conducted on cluster availability and downtime over the past five years. TechWise defines a cluster "crash" as any situation that causes the cluster's primary application(s) to become unavailable to end-users. Downtime hours are the number of hours per year, if any, when the cluster's primary application(s) were not available for end-users to access. TechWise Research developed a list of five potential causes of downtime. These are explained in detail in the section entitled: Possible Reasons for Crashes and Downtime.

This study focuses on three platforms: HP Integrity servers running OpenVMS, IBM System p5 servers running AIX, and Sun Sun-Fire servers running Solaris. The latter two platforms were studied in great detail in the 2004 paper. TechWise surveyed companies in 2006 and determined that on average the number of hours spent managing a server cluster declined by 25% between 2004 and 2006. Similarly, the number of hours of downtime declined by 30% during that same period. In order to calculate the management and downtime hours for IBM and Sun, TechWise reduced IBM's and Sun's 2004 operational costs by 50%. By applying a greater reduction than what was measured, TechWise is using a conservative approach that is likely understating IBM's and Sun's management and downtime costs. The HP Integrity platform for OpenVMS, on the other hand, was not available in 2004. TechWise conducted surveys in 2006 with companies that have these clusters to calculate its management and downtime costs. TechWise Research used a web survey to collect operational and profiling data about all three brands of clusters. The web survey lasted between 25 and 30 minutes. Throughout the survey, respondents were given several opportunities to clarify any answers they provided. One of TechWise Research's senior analysts, who specializes in server clusters, personally reviewed each completed survey and followed-up with respondents by phone if any answers needed clarification. Approximately 30 surveys were completed with each of the three cluster brands.

Previous TechWise cluster studies have indicated that a three-year time frame is appropriate to evaluate entry-level and mid-range clusters' TCO. Once again, this same time period was applied in the analysis for this current study. Since each company will have different costs associated with downtime, TechWise also calculated the three-year TCO at various hourly rates of downtime costs.

Configurations Analyzed

TechWise Research analyzed four different cluster configurations for this paper. All the clusters were two node clusters. They differed in terms of the server models, memory, and storage. These configurations were selected because they represent the manufacturers' full range of entry-level and midrange servers. The specific server models in each configuration were selected because they represent comparable machines in terms of performance and expandability. The chart to the right describes the first two configurations tested. These two configurations represent the entry-level clusters in the study. Some servers have options for different CPUs. For this reason we included the specific CPU in our descriptions. Almost all of the servers utilize dual core processors, a technology that was not widely available when TechWise last studied these clusters. In some cases it was not possible to have all the servers configured with the same number of processor cores. In these cases a configuration was selected that best matched the desired target.

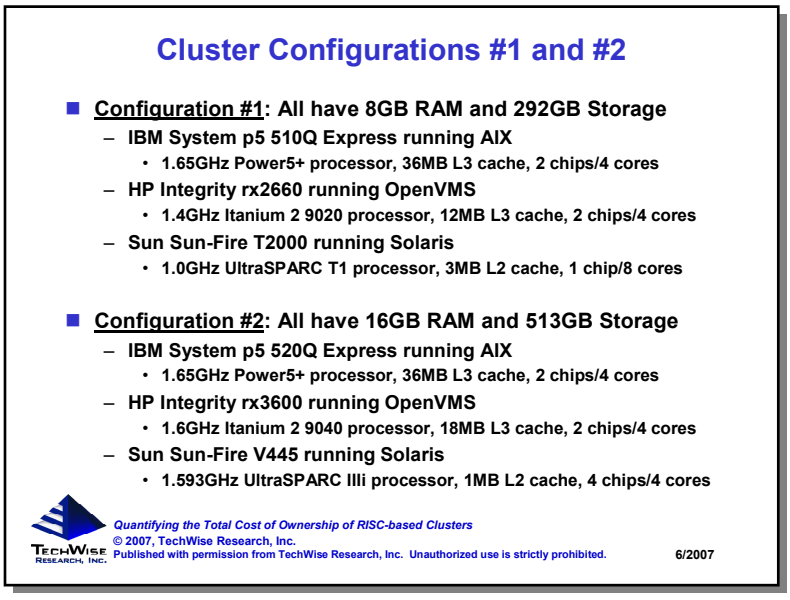

Cluster Configurations #3 and #4

- **Configuration #3: All have 16GB RAM and 1,168GB Storage**
  - IBM System p5 550Q Express running AIX
    - 1.65GHz Power5+ processor, 36MB L3 cache, 4 chips/8 cores
  - HP Integrity rx6600 running OpenVMS
    - 1.6GHz Itanium 2 9040 processor, 18MB L3 cache, 4 chips/8 cores
  - Sun Sun-Fire V490 running Solaris
    - 2.1GHz UltraSPARC IV+ processor-, 32MB L3 cache, 4 chips/8 cores
- **Configuration #4: All have 64GB RAM and 1,200GB Storage**
  - IBM System p5 570 running AIX
    - 1.9GHz Power5+ processor, 36MB L3 cache, 8 chips/16 cores
  - HP Integrity rx7640 running OpenVMS
    - 1.6GHz Itanium 2 9040 processor, 18MB L3 cache, 8 chips/16 cores
  - Sun Sun-Fire E4900 running Solaris
    - 1.95GHz UltraSPARC IV+ processor, 32MB L3 cache, 8 chips/16 cores

Quantifying the Total Cost of Ownership of RISC-based Clusters
© 2007, TechWise Research, Inc.
Published with permission from TechWise Research, Inc. Unauthorized use is strictly prohibited. 6/2007

The chart to the left shows the other two configurations studied. These mid-range clusters are configured with more memory and storage than the entry-level configurations.

TCO Findings Excluding Downtime

In Configurations 1, 2, and 3, IBM offers the lowest acquisition costs of all three cluster brands. In these cases, the price to purchase the IBM cluster complete with a three-year service plan typically costs 10% to 30% less than HP or Sun. The only exception to this is Configuration #4 where the HP OpenVMS rx7640 cluster costs 5% less than the IBM AIX 570 cluster and 29% less than the Sun Solaris E4900 cluster. IBM's lower acquisition costs, however, do not tell a complete picture of the cost of owning a cluster. The relative TCO between the three brands changes when start-up and management costs are included in the analyses.

The chart to the right shows the breakdown of these three TCO components for Cluster Configuration #3. All numbers have been normalized to the acquisition cost of the IBM cluster. The blue bars on the chart show that both the Sun Solaris V490 cluster and the HP OpenVMS rx6600 cluster acquisition costs are almost 140%, or 40% greater, than the IBM AIX 550Q cluster. Start-up costs, shown in orange, are a relatively minor component of TCO. Management costs, shown in green, however, are significant. The cost to manage this IBM cluster over three years is more than twice the costs to manage the HP cluster. When all three factors are added together, the HP OpenVMS rx6600 cluster has a lower TCO than the Sun Solaris V490 and IBM AIX 550Q clusters. The next sections address what happens when downtime is also included in the analyses.

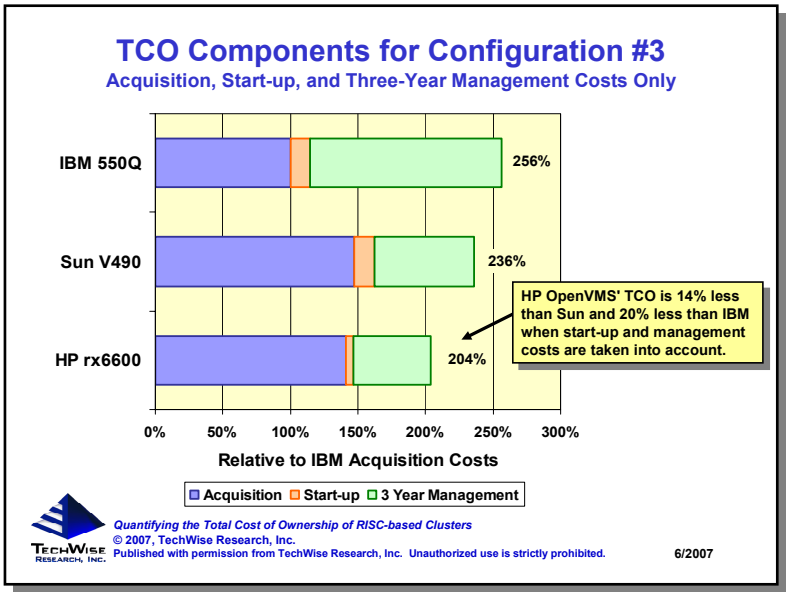

The Value of an Hour

In past studies respondents rated how important nine different factors would be in a future cluster purchase decision. The two most important factors included (1) the cluster's overall reliability, and (2) how well the cluster software performs when there is a failure. These two factors were rated as more important than overall performance, security features built into the operating system, applications supported on the platform, ability to add more servers to the cluster, ease of management, and disaster tolerance abilities. The fact that reliability and cluster software are rated highest demonstrates the overall importance of availability to cluster users, and further re-enforces the primary reason for establishing a cluster - to ensure that applications are available to end-users 24x7.

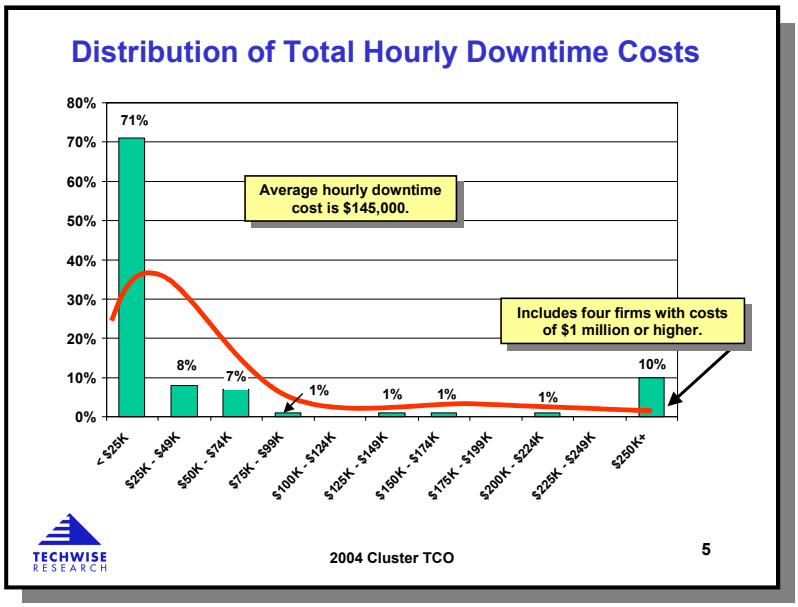

As reported in prior TCO studies conducted by TechWise Research, each company has a unique situation that determines the financial impact of downtime. For some, when primary applications are not available to end-users, the impact is lost sales. For others, it means lost employee productivity or a decline in manufacturing production. Many firms are affected in multiple ways. TechWise Research asked each respondent to quantify the financial impact per hour of downtime. The chart to the left shows the distribution of hourly downtime costs. Seventy-one percent of respondents said their costs are less than $25,000 per hour. Thirteen percent estimate their hourly costs at $125,000 or more. Four of the companies surveyed indicated that they lose **$1 million or more per hour**! On average, however, respondents report that **each hour of downtime costs their firm a total of $145,000** when the costs associated with lost sales, wages, and production are considered. These downtime cost figures demonstrate the important role availability plays in calculating the true total cost of ownership of a cluster.

Possible Reasons for Crashes and Downtime

TechWise Research defines a cluster crash as any event that caused one or more of the cluster's primary applications to become unavailable to end-users. Some crashes result in only a few seconds of downtime as the cluster software "fails over" to another node. Other crashes can cause applications to be down for minutes or even hours. As in past studies, TechWise Research collected downtime information on all crashes, no matter how short or long in duration.

There are a variety of potential causes for crashes and downtime. In this TCO analysis, TechWise Research included downtime from the following five different causes:

1. **Server hardware failure during normal cluster operation:** These are crashes caused when one or more servers failed, when no maintenance was being performed on any of the servers.
2. **Server hardware failure during planned server maintenance:** This type of crash occurs when one server failed while planned maintenance is being performed on the other server in the cluster. In these cases, the cluster crashes despite its automated failover configuration.
3. **Operating system or cluster software problems:** These are any crashes caused by the operating system or clustering software.
4. **Software virus or worm:** Hackers and malware can and do cause clusters to crash. TechWise Research measured how much downtime, if any, was caused by software viruses and worms.
5. **System management application problem:** This includes any crash caused by a problem with any of the system management application(s) running on this cluster.

In past studies TechWise also included crashes caused by end-user application problems and hardware failures in a storage array. Every cluster is running a unique set of end user applications. Furthermore, virtually all of these end user applications come from third-parties or are custom written. Since any crashes caused by third-party applications are beyond the control of HP, IBM, and Sun, these crashes were excluded from the analyses. TechWise has noticed that an increasing number of cluster customers are using third-party storage arrays, such as from EMC. Once all the surveys were collected, there was insufficient data to analyze the reliability of HP, IBM, and Sun's own storage arrays. For this reason crashes caused by the cluster's storage array were also excluded from the analyses. By using the above factors, the reliability analysis is based exclusively on each manufacturer's controllable hardware and software solutions.

Average Annual Downtime Hours
All Downtime Factors

Cause of Downtime	HP Integrity	IBM System p5	Sun Sun-Fire
Hardware Failure - Normal Operation	1.59	1.37	2.57
Hardware Failure - Planned Maintenance	0.18	0.48	1.99
Operating System or Cluster Software	0.19	0.58	1.23
Software Virus or Worm	0.00	2.87	2.16
System Management Application	0.00	0.91	1.58
TOTAL	1.96	6.21	9.53

Note: A cluster was considered "down" when one or more of its primary applications was not available for end-users. Best in class results are highlighted in green.

Quantifying the Total Cost of Ownership of RISC-based Clusters
© 2007, TechWise Research, Inc.
Published with permission from TechWise Research, Inc. Unauthorized use is strictly prohibited.
6/2007

Cluster Reliability Findings

Each respondent reported the total number of hours per year, if any; their cluster's primary applications were offline due to each of the five types of problems previously described. The table to the left summarizes these findings by platform. "Best in class" findings for each category are highlighted in green. **HP OpenVMS Integrity clusters averaged the fewest total number of annual downtime hours of all three brands tested**. HP's average downtime of 1.96 hours is 68% less than IBM's System p5 clusters (6.21 hours), and 79% less than Sun's Sun-Fire clusters (9.53 hours). In four of the five categories, HP won best in class. For the remaining category, IBM clusters won best in class.

The first two downtime categories listed in the table above are crashes caused by server hardware failures. For server hardware failures occurring during normal cluster operation, IBM averaged the fewest number of downtime hours followed closely by HP. For server hardware failures occurring during planned maintenance, the results switch whereby HP averaged virtually no downtime hours, followed closely by IBM. Combining these two categories show that there is little difference between the reliability of HP Integrity and IBM System p5 server hardware (1.77 hours and 1.85 hours, respectively). Sun's Sun-Fire hardware is somewhat less reliable, averaging more than double the downtime (4.56 hours).

Combining the last three categories gives a complete measurement of the reliability of the three operating system/clustering software combinations. This is because any crash caused by a software virus or worm is likely due to vulnerability in the operating system or clustering software. Furthermore, TechWise designed the survey to differentiate between third-party and OEM system management applications. This analysis focused only on downtime caused by the manufacturers' system management applications. OpenVMS clearly separates itself from AIX/HACMP and Solaris/Sun Cluster. **HP OpenVMS averages only 0.19 hours of downtime due to the operating system and clustering software. This works out to 99.998% availability**. IBM's AIX and HACMP average 4.36 hours (99.950% availability) while Sun's Solaris and Sun Cluster average 4.97 hours (99.943% availability). These differences may seem small. However, they are in fact significant in cases where downtime costs are measured in tens of thousands of dollars.

In summary, when all five types of crashes are considered, IBM System p5 clusters averaged 4.25 more downtime hours per year than HP Integrity OpenVMS clusters. Furthermore, Sun Sun-Fire clusters averaged 7.57 more downtime hours per year than HP Integrity OpenVMS clusters. Downtime costs vary from company to company. **For each $10,000 of hourly downtime costs, HP's total Availability Advantage™ over IBM and Sun translates into $127,500 and $227,100, respectively, over three years**.

Detailed Findings: Reliability-Adjusted TCO™ for Configuration #1

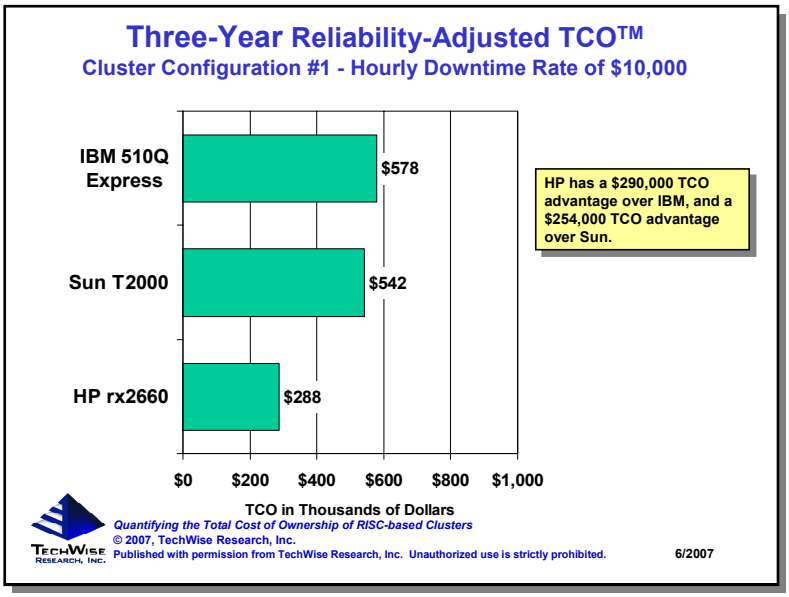

The chart to the left summarizes the Three-Year Reliability-Adjusted TCO™ for the three entry-level clusters studied. The bars represent the total cost of owning each cluster including acquisition, service, installation, training, management and downtime over a three-year period. Downtime costs were calculated by applying a rate of $10,000 per hour. The **HP OpenVMS rx2660 cluster has the lowest TCO. Over a three-year period, HP's TCO advantage is $290,000 over IBM AIX 510Q Express clusters and $254,000 over Sun Solaris T2000 clusters.**

The application of a $10,000 per hour downtime rate may be too low for some firms, or too high for others. For this reason, TechWise Research provides a more detailed graph below. This chart summarizes the three-year Reliability-Adjusted TCO™ for the three clusters studied at various hourly downtime rates. In cases where there are no costs associated with downtime, there is little difference between this HP OpenVMS Integrity cluster and the comparable Sun Solaris Sun-Fire cluster. However, as the cost per hour of downtime increases, HP's cluster develops a clear TCO advantage.

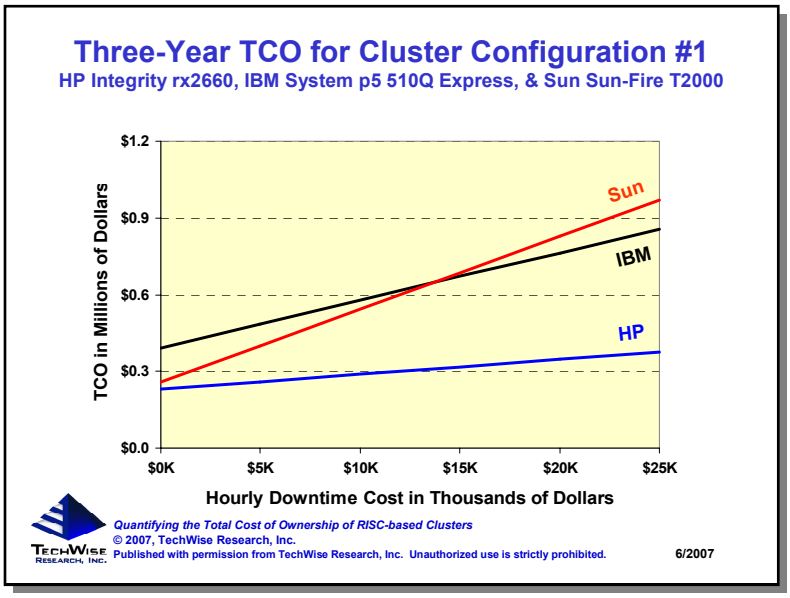

At lower hourly downtime rates the key drivers to TCO are acquisition and management costs. The IBM AIX 510Q Express cluster has the lowest acquisition cost of the three clusters. However, this is more than offset by the high cost to manage the IBM cluster on an ongoing basis. IBM's high management costs are why the IBM cluster has the highest TCO at low hourly downtime rates. Once the hourly downtime rate exceeds $15,000 this IBM cluster has a lower TCO than the Sun cluster. This is thanks to the fact that **IBM clusters are, on average, more reliable than Sun clusters**.

Detailed Findings: Reliability-Adjusted TCO™ for Configuration #2

The second cluster configuration studied involved the HP rx3600, IBM 520Q Express, and Sun Sun-Fire V445 servers. The chart to the right summarizes the Three-Year Reliability-Adjusted TCO™ for these three clusters. A higher hourly downtime rate of $25,000 was used for these clusters. As seen before, HP OpenVMS Integrity clusters have the lowest TCO. **HP's TCO is $628,000 less than Sun Solaris Sun-Fire clusters and $470,000 less than IBM AIX System p5 clusters over a three-year period.** The chart below summarizes the three-year Reliability-Adjusted TCO™ for these three clusters studied at various hourly downtime rates. Compared to Configuration #1, these clusters have higher acquisition and service costs. However, despite the higher acquisition costs, the findings are little changed thanks to the impact management and downtime costs have on TCO.

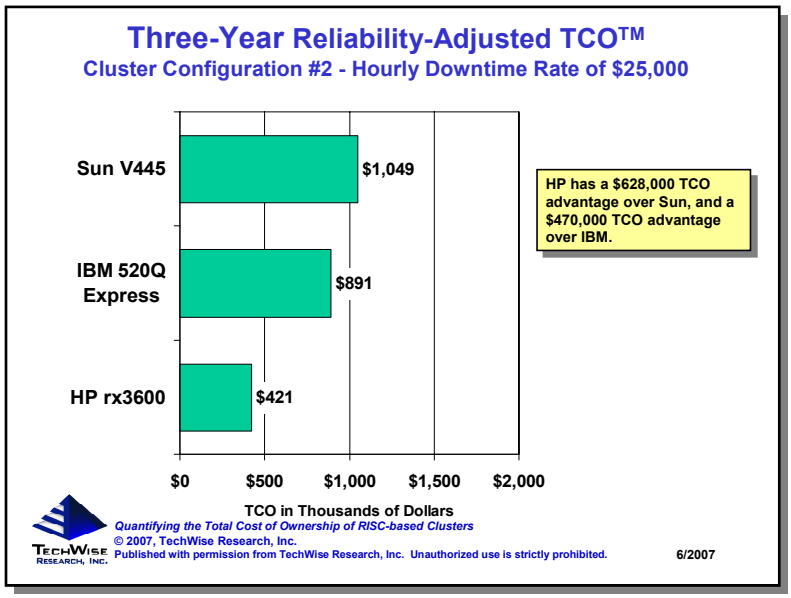

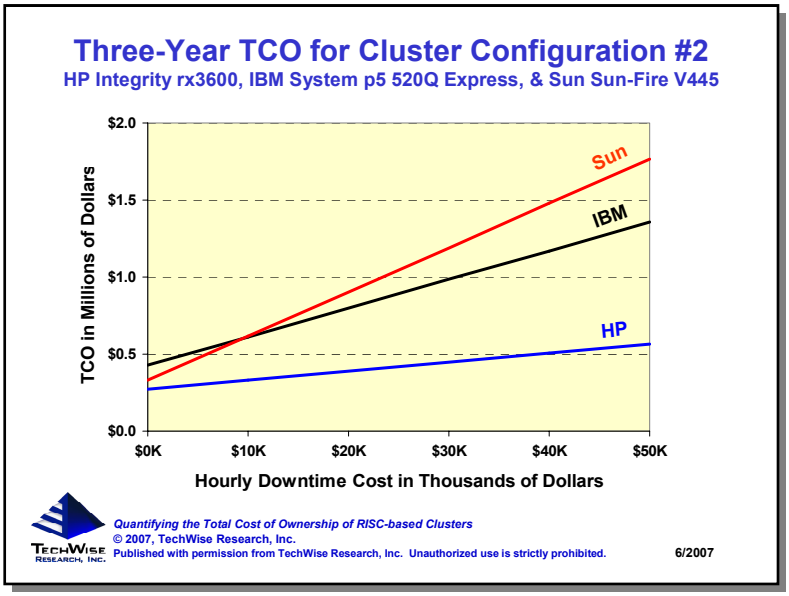

Again, at all downtime rates the HP OpenVMS rx3600 cluster has the lowest TCO of the three brands studied. In this scenario, Sun Solaris V445 clusters beat IBM AIX 520Q Express clusters for second place in situations where downtime costs are below $10,000 per hour. At higher downtime rates IBM's Availability Advantage over Sun makes their clusters less expensive than Sun. In situations where there are no costs associated with downtime, the differences between the three clusters is relatively small at $164,000. This figure includes all acquisition, set-up, and management costs. As downtime costs increase, the difference between the brands grows to as much as $1.2 million. In this configuration, the system cost (hardware, software, and service contract) for all three brands is approximately $100,000. **Another way to look at these findings is that for every $5,000 in hourly downtime costs, companies can save enough money with the HP OpenVMS rx3600 cluster to purchase another two-node cluster complete with a three-year service contract**.

Detailed Findings: Reliability-Adjusted TCO™ for Configuration #3

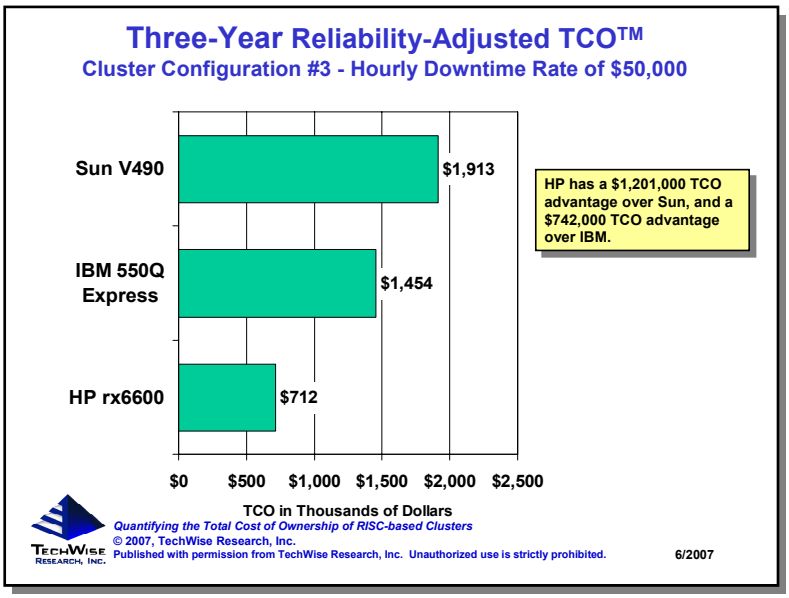

The chart to the left summarizes the Three-Year Reliability-Adjusted TCO™ for Configuration #3. Given that this configuration involves more expensive mid-range systems, these findings were calculated using a higher hourly downtime rate of $50,000. As with the other two cluster configurations, HP OpenVMS Integrity clusters have the lowest TCO. **The three-year TCO of the HP OpenVMS rx6600 cluster is $1.2 million less than the Sun Solaris V490 cluster and $0.7 million less than the IBM AIX 550Q Express cluster**.

The chart to the right summarizes the three-year Reliability-Adjusted TCO™ for these three mid-range clusters at various hourly downtime rates. These three clusters have very similar acquisition costs as the differences between them are less than $100,000. However, as the cost per hour of downtime increases, there is a significant difference in their three-year TCO. As with the other two configurations studied, the HP OpenVMS rx6600 cluster has the lowest TCO no matter what hourly downtime cost is used.

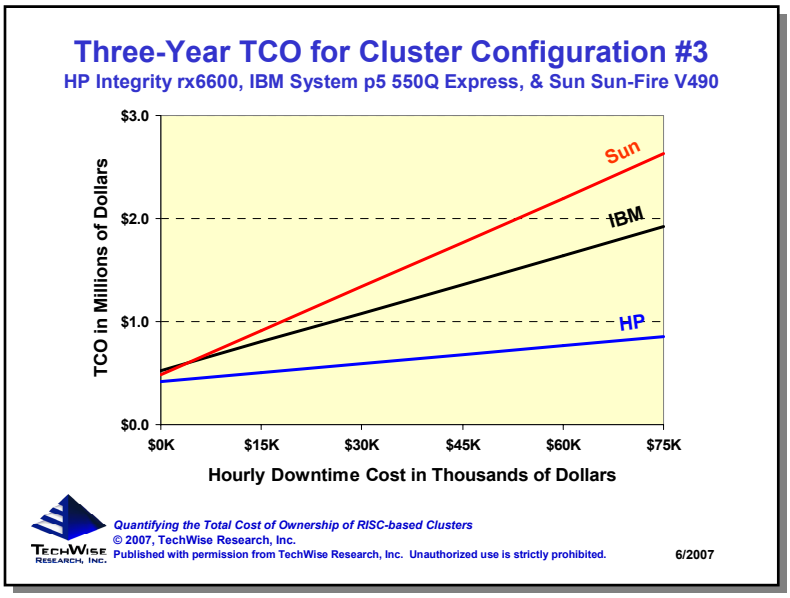

In virtually all cases, IBM's AIX 550Q Express cluster has the next lowest TCO followed by Sun's Solaris V490 cluster. At the highest downtime rate analyzed, $75,000 per hour, the **Sun Solaris V490's TCO is more than triple the TCO of the HP OpenVMS rx6600 cluster while IBM's AIX 550Q Express cluster costs more than double the HP OpenVMS rx6600 cluster.** HP's advantage over IBM and Sun with these clusters will only get bigger as the cost per hour of downtime increases.

Detailed Findings: Reliability-Adjusted TCO™ for Configuration #4

The final configuration studied includes mid-range systems that are one step below enterprise class. All three clusters, the HP OpenVMS rx7640, IBM AIX 570, and Sun Solaris E4900 have acquisition costs exceeding $1 million. Certainly these million-dollar clusters are usually deployed in instances where downtime is important. For this reason, TechWise used a $100,000 hourly downtime rate when calculating the Three-Year Reliability-Adjusted TCO™ shown in the chart to the right. **The HP OpenVMS rx7640 cluster has a three-year TCO that is $2.8 million less than the Sun Solaris E4900 cluster and $1.5 million less than the IBM AIX 570 cluster.**

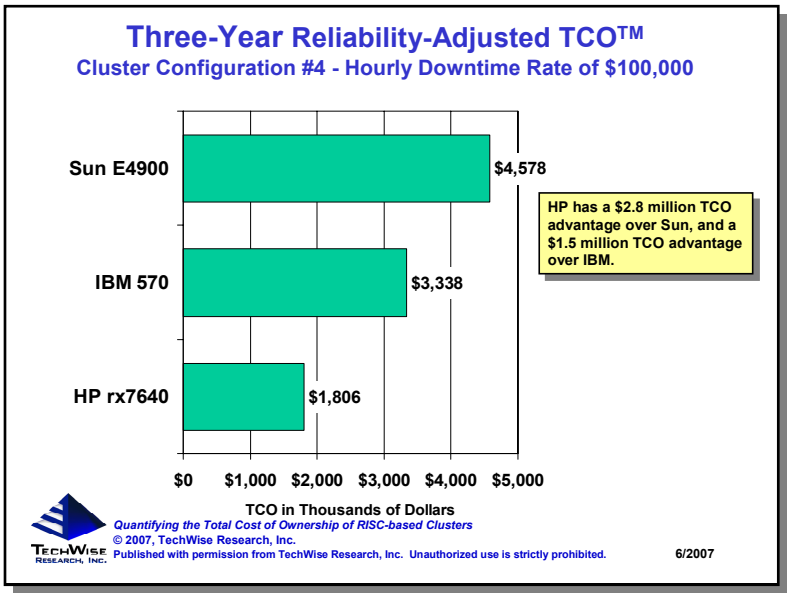

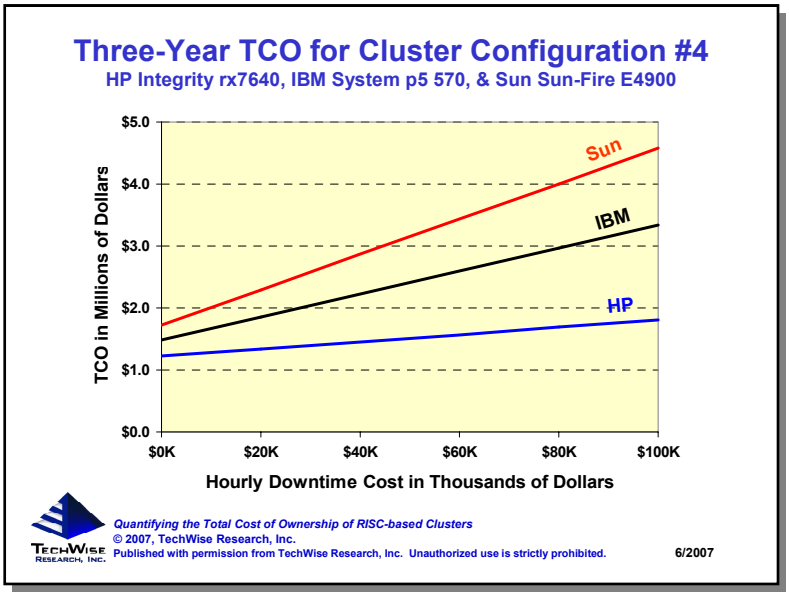

The chart to the left summarizes the three-year Reliability-Adjusted TCO™ for these "million-dollar clusters" at various hourly downtime rates. This chart shows that for some clusters at higher downtime rates the acquisition costs can represent less than 25% of the total TCO. The differences in acquisition costs are greatest in this configuration. There is nearly a $500,000 difference between Sun and HP. However, unlike all previous studies conducted by TechWise, *the HP cluster is less expensive than the Sun cluster in terms of system and service pricing*. In all previous studies the HP OpenVMS cluster consisted of AlphaServers which tended to have higher system prices than IBM and Sun. HP transitioned its OpenVMS business from AlphaServer to Integrity. One of the reasons for this transition was to be more price competitive with Sun and IBM. It is clear from the above that HP has met this goal with this configuration. HP's significant TCO advantage over Sun (and IBM too) only grows as downtime costs increase. In this configuration it can be said that choosing the **HP OpenVMS rx7640 cluster over IBM AIX 570 cluster and Sun Solaris E4900 cluster can save companies millions of dollars over the long-term.**

Differences between Short and Long-Term Evaluations

The analyses on the previous pages show how three-year TCO varies depending on the cluster configuration and the costs incurred per hour of downtime. A three-year period was chosen for several reasons. First, manufacturers offer three year service contracts when purchasing a new server cluster. These service contracts can be extended, of course. Out of warranty support costs, however, are more difficult to obtain and frequently change. In addition, it is difficult to predict what manufacturers will charge four years in the future for service for a server purchased today. The second reason for choosing three years is that TechWise has seen in past studies that after three years many companies have changed the number and/or types of servers in their cluster. Not all companies, however, look beyond the first year or two when comparing server clusters on cost and TCO. TechWise Research conducted additional analysis on Cluster Configuration #3 to see how TCO varies over time.

It is interesting to note that when simply looking at acquisition costs (purchase price plus service contract), the relative pricing between the HP OpenVMS Integrity, IBM AIX System p5, and Sun Solaris Sun-Fire clusters is not always the same. In Configuration #1 IBM has the lowest acquisition costs followed by Sun then HP. In Configurations #2 and #3, IBM has the lowest acquisition costs followed by HP then Sun. In Configuration #4 HP has the lowest acquisition cost followed by IBM and Sun.

The chart to the right shows how the TCO for Cluster Configuration #3 changes over time. A $25,000 per hour downtime cost was used in this analysis. At the time of purchase the only costs incurred are acquisition and start-up costs. The IBM AIX 550Q Express cluster is more than 20% less expensive than the HP OpenVMS rx6600 and Sun Solaris V490 clusters. This situation changes, however, after one year of management and downtime costs is included in the analysis. After one year, the HP OpenVMS rx6600 cluster's TCO is 20% lower than the IBM AIX 550Q Express cluster and 37% lower than the Sun Solaris V490 cluster. These differences become more pronounced after two years where HP's TCO is 35% and 48% lower than IBM's and Sun's, respectively. This analysis shows the danger of looking strictly at acquisition costs when evaluating clusters on a financial perspective. Based on acquisition costs it appears that the IBM cluster will cost 20% less than HP and Sun. After one year the IBM cluster still does cost 20% less than Sun. However, the HP cluster costs 20% less than the IBM cluster. As shown before, **looking strictly at acquisition costs can greatly mislead companies about the true costs of owning a cluster.**

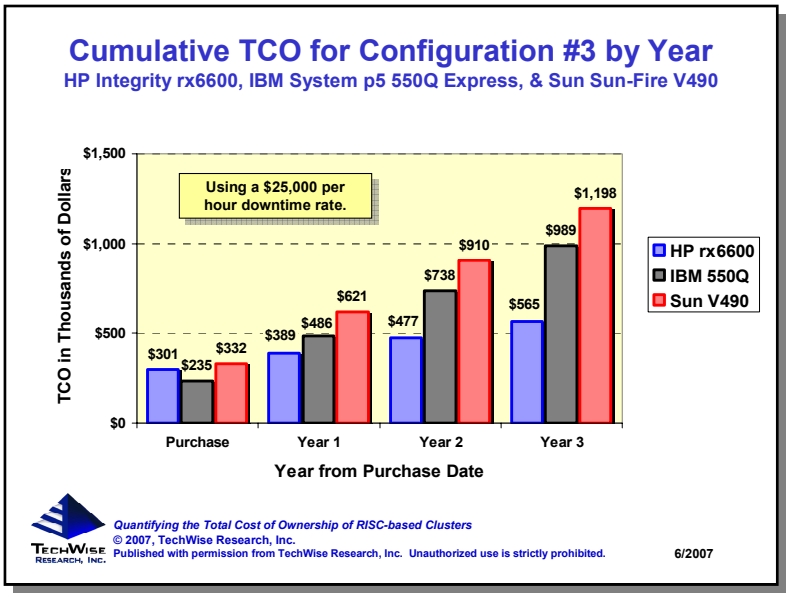

Conclusion

This study focused on the three-year total cost of ownership for three different server clusters: HP OpenVMS Integrity, IBM AIX System p5, and Sun Solaris Sun-Fire clusters. Four different costs were factored into this TCO analysis including the acquisitions costs, start-up costs (i.e., training and installation), management, and downtime costs. Most of the TCO calculations were based on a three-year time period.

The cluster configurations used in the analysis were selected to represent the manufacturers' entry-level and midrange product lines. Four different cluster configurations were analyzed. Comparable server models from each manufacturer were selected to use for the analysis (see section entitled Configurations Analyzed).

Acquisition costs alone, namely the list prices of the cluster and related service agreements, represent only a portion of the true Total Cost of Ownership. Management costs, which include any fees paid to external parties and the time IT staff spends managing the cluster on an ongoing basis, are significant contributors to TCO. However, in many cases the most important driver of TCO is related to availability. Every company has different costs associated with an hour of downtime. Respondents reported that their average cost per hour of downtime was $145,000. Nearly half indicated that their cost per hour of downtime is $10,000 or greater. In situations where downtime costs are not trivial, downtime costs are the single biggest factor of TCO. Given the impact of downtime costs, the Reliability-Adjusted TCO™ was calculated for each of the four cluster configurations at several different hourly downtime rates. This way, readers have the ability to apply the results to their company's unique situation by selecting the appropriate hourly downtime rate. Downtime hours were analyzed for five different causes of downtime. When all five causes of downtime are considered, **HP OpenVMS Integrity clusters averaged the fewest total number of annual downtime hours (1.96), followed by IBM AIX System p5 clusters (6.21), then by Sun Solaris Sun-Fire clusters (9.53)**. HP cluster's have a significant Availability Advantage™ over IBM and Sun. Much of this advantage is thanks to the reliability of OpenVMS. Isolating crashes caused by operating system and clustering software showed that **HP OpenVMS averages only 0.19 hours of downtime compared to 4.36 hours for IBM AIX/HACMP and 4.97 hours for Sun Solaris/Sun Cluster**.

For all four configurations studied, HP OpenVMS Integrity clusters offer the lowest TCO *regardless* of the cost associated with an hour of downtime. This represents a change from past studies where Sun clusters offered the lowest TCO when there are no costs associated with downtime. HP's emergence as the clear leader in TCO in all scenarios is thanks to their transition from AlphaServer to Integrity. The Integrity servers have a much lower acquisition cost than the AlphaServers they replace. This study showed that despite the lower price point, Integrity servers are every bit as reliable as AlphaServers. HP OpenVMS Integrity servers offer the combination of low acquisition costs and best-in-class reliability. These two factors contribute to HP's TCO leadership. In all but Configuration #4, Sun Solaris Sun-Fire clusters offer lower TCO than IBM AIX System p5 clusters at low downtime hourly rates. Once the cost per hour of downtime exceeds a certain level, IBM's better reliability compared to Sun gives IBM a lower TCO than Sun. Cluster Configuration #4 represents the largest systems studied in this paper. These clusters, which are one step below enterprise-class clusters, have price tags over $1

million. In this configuration, Sun clusters have the highest TCO regardless of the cost per hour of downtime.

The differences in TCO between these three cluster brands can be quite significant. Over a three-year period the cost of owning one brand of cluster may be hundreds of thousands of dollars less than another brand. The HP OpenVMS Integrity clusters consistently offered the lowest TCO. **Applying typical hourly rates for downtime costs, HP OpenVMS Integrity cluster's cost between $290,000 and $1.5 million less than IBM AIX System p5 clusters over a three-year period. Similarly, HP's TCO advantage over Sun ranges from $254,000 to $2.8 million over three years**.

Several factors outside the scope of this study are also important in the purchase decision for entry-level and mid-range clusters. These include cluster performance, software features, application availability, and quality of service and support. However, when comparing clusters from a financial perspective, it is critical that IT managers factor in management costs as well as the cost of downtime when calculating the TCO.

TechWise Research is an independent primary market research firm that specializes in the computer industry. If you have any questions regarding this research, please contact us at:

TCO2007@TechWise-Research.com.

AlphaServer, Integrity, and OpenVMS are trademarks of Hewlett-Packard. System p5, AIX, and HACMP are trademarks of IBM. Sun-Fire, Solaris, and Sun Cluster are trademarks of Sun Microsystems. Intel and Itanium are trademarks of Intel. Reliability-Adjusted TCO, Availability Advantage, and TechWise are trademarks of TechWise Research. For more information about TechWise Research, visit our website at:

http://www.techwise-research.com.