TruCluster Server 5.0 Slashes Cost of Administration by Streamlining Operations

Prepared for Compaq by:

D.H. Brown Associates, Inc.
222 Grace Church Street
Port Chester, NY 10573
(914) 937-4302
Fax: (914) 937-2485

June 2000

TABLE OF CONTENTS

EXECUTIVE SUMMARY **1**

CHAPTER 1. INTRODUCTION **3**

PURPOSE OF THE REPORT 3
CURRENT STATE OF AFFAIRS 3
METHODOLOGY 3

CHAPTER 2. VERIFY AND SYNCHRONIZE **5**

OTHER MANAGEMENT TASKS 9

CHAPTER 3. FAILOVER **11**

SPONTANEOUS FAILOVER 12
NFS 13
UNSUCCESSFUL FAILOVER 13

CHAPTER 4. CLUSTER ALIAS, LOAD BALANCING, AND CLUSTER APPLICATION ARCHITECTURE (CAA) **15**

CHAPTER 5. CONCLUSION **17**

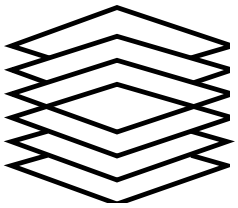

TruCluster Server 5.0 Slashes Cost of Administration by Streamlining Operations

EXECUTIVE SUMMARY

Compaq's TruCluster Server 5.0 eliminates much of the overhead associated with previous generation cluster products. As a result, TruCluster administrators save up to 82% of the time required to perform certain administrative tasks using older technology. While the current generation of competitive cluster products also address some of these issues, Compaq's solutions arise from its architectural superiority whereas other vendors provide point solutions. As a result, Compaq represents a more enduring and effective solution, delivering greater efficiency than competitors.

TruCluster Server 5.0 saves administrator time compared to previous generations of clustering products in the following ways:

- By eliminating the need to replicate changes across all nodes, TruCluster Server 5.0 reduces the maintenance effort for some tasks by 82% compared to manual replication and up to 66% compared to automated replication. This level of reduction derives not just from routine maintenance across eight nodes, but also from the iterative replication that slows down the development of customized scripts.
- TruCluster Server 5.0 reduces the likelihood of unplanned failover due to human error. It also allows an administrator to control a number of hardware-generated spontaneous failovers and can reduce NFS-related issues. These enhancements reduce the time required for failover management by up to 70% for certain failover problems.
- TruCluster Server 5.0 reduces the effort of cluster configuration by providing load balancing capability as well as policy-based failover definition. Client requests and processes need no longer be hard coded to a specific node, but can be assigned freely among all nodes in the cluster according to resource availability. These features should save a minimum of 50% of the work effort devoted to planning the configuration, configuring the cluster, and relocating resources according to changing user requirements.

As a result of these enhancements, Compaq has overcome many of the manageability issues that have hindered full utilization of cluster capability. Corporations now can look more favorably upon scaling their clusters beyond two nodes and implementing multi-instance applications.

CHAPTER 1. INTRODUCTION

PURPOSE OF THE REPORT

This report from D.H. Brown Associates, Inc. (DHBA) demonstrates the clear advantage in administration costs offered by TruCluster Server 5.0 when compared to previous generations of cluster technology and to the cluster offerings of HP, IBM, and Sun. TruCluster Server 5.0 lowers the cost of administration by reducing the time required for administration and maintenance. It also reduces complexity sufficiently for administrators to utilize the technology more broadly and effectively.

CURRENT STATE OF AFFAIRS

Today, cluster technology is underutilized. All surveyed administrators (100%) report cluster configurations of only two nodes. These configurations actively use one node for mission-critical production applications and use the second for non-critical activity such as testing and quality assurance. Over 60% of the installations reported single application usage. Of the remaining 40% that support multiple concurrent applications, half are TruCluster systems.

From these facts, one can conclude that clusters are manageable at current utilization levels (since administrators do not spend much time on them), but that the threat of heavy administrative burdens from larger configurations limits more widespread use.

METHODOLOGY

This study used data gathered during interviews with system administrators and vendor representatives, and from product analysis. Results of the interviews provided the foundation for the analysis and also supported the results of the study. DHBA did not design this study as a statistically accurate survey.

The interviewees. Initially, DHBA analysts interviewed dozens of cluster administrators to identify those with extensive understanding of and experience with clusters. DHBA also evaluated the size and complexity of candidates' installations to find those of similar size and complexity. Using these criteria, DHBA selected 12 administrators for in-depth discussion of their practices. These administrators used their clusters as back-end servers for interactive, client/server, and three-tier applications that served 250-500 users. The clusters were maintained by two to four staff members who were responsible for several additional systems, so that the total effort devoted to the cluster was approximately one man-year. Sun administrators had one to two years of cluster experience; the administrators of other systems had two to four years experience; and Compaq administrators had the greatest amount of experience overall.

This study concludes that TruCluster Server 5.0 addresses the most time-consuming aspects of today's cluster administration. On the surface, it may appear that the methods Compaq used have inadvertently introduced new time-consuming activities or areas of complexity. For example, early evaluators report TruCluster Server 5.0's Context-Dependent Symbolic Links (CDSLs) to be confusing, since it does not appear immediately clear whether one is addressing the cluster or a particular node. DHBA fully believes this issue will disappear once an administrator has worked with the system for just a few weeks.

CHAPTER 2. VERIFY AND SYNCHRONIZE

Two-thirds of surveyed administrators identify the synchronize/verify sequence of operations as a cause of frustration, potential error, and reduced efficiency. They agree that propagation of changes make the difference between a simple job and a hard job. All agree they can only trust experienced personnel with maintenance tasks because of this factor.

TruCluster Server 5.0 users derive substantial economies simply by eliminating the verify-and-synchronize routine that applies to many operations. The example of user-account maintenance illustrates the benefits of the TruCluster Server architecture over competing architectures.

TABLE 1: User Account Management

	User Accounts	Time-Consuming Element				
Wt[1]	**Activity**	**Compaq**	**Previous Generation**	**HP**	**IBM**	**Sun**
30	Initial data entry	✓	✓	✓	✓	✓
5	Ensure that all nodes are up and running	×	✓	✓	✓	✓
20	Assisted copy onto 8 nodes	×	✓	×	×	✓
40	Manual copy onto 8 nodes	×	✓	×	×	×
5	Verify that changes have been made across all nodes	×	✓	✓	✓	✓
100						
Time spent compared to previous generation		**30%**	**100%**	**40%**	**40%**	**60%**
Time saved by using Compaq			**70%**	**25%**	**25%**	**50%**

[1] Weight – the percent of total time devoted to each activity

Table 1 summarizes the steps required to maintain user accounts on each platform. It compares four competitive cluster offerings: Compaq TruCluster Server 5.0 (Compaq), Hewlett-Packard's HP-UX MC/ServiceGuard 11.0 (HP), IBM's HACMP 4.3, (IBM) and Sun SunClusters 2.2 (Sun). The column labeled "previous generation" estimates the typical activity level of a cluster running software at least one full version number behind the others included in this chart.

In real life, previous-generation clusters did not support eight nodes (except for IBM). However, if one projected previous-generation cluster operation across a hypothetical eight-node cluster, the data entry could take as little as 30% of the total operation. The rest of the time would be occupied in verifying that all nodes are online, manually copying the changes onto eight nodes, and then verifying the changes had been made. Compaq, in contrast, requires just the data entry, since

all nodes will access the same updated file. Updating user information takes only 30% of the time required by previous generation clusters.

User account information in password files and in other security files must remain consistent across nodes in order for users to experience smooth failover. If a cluster node fails, users should be able to log onto the surviving nodes without experiencing problems caused by mismatches in the user or group IDs. By the same token, users need consistent user information to allow hassle-free logon and access to their home directory from any node. This is particularly important for system administrators, who are not necessarily tied to their desks when changes must be made. Similarly, logging in as root on a cluster member gives you root privileges clusterwide.

Compaq delivers this uniform access to user information by automatically and immediately making visible a single copy of all user information, security information, and home directory on the entire cluster. In other words, TruCluster Server implements a cluster-wide /usr directory. This cluster-wide /usr sits on a cluster file system (CFS) that allows multiple nodes to share the same files and to load software from a common image. This will make a single data entry immediately available clusterwide.

In contrast, competing UNIX offerings maintain separate user information on each node. In order to make this information available clusterwide, the administrator must replicate each node's user information across all the nodes. Otherwise users will experience difficulty logging onto the takeover node after a failover and will not have access to the cluster in the event that their local node goes down. Other UNIX cluster products attempt to circumvent this problem by automatically replicating system files among all the nodes. In order for this replication to work successfully, however, the administrator must ensure that all the nodes are up and running before initiating the replication. After completion of the replication, the administrator must verify the propagation has occurred properly. These procedures require moderate time to complete, but administrators must repeat them frequently – each time user information is updated. Furthermore, the amount of time will escalate dramatically if the procedures do not go as expected.

Each of the competitive UNIX clusters automates user management in a different way.

HP has a clusterwide security file for validating internode access within the cluster and automates the propagation of changes. AIX has implemented user management under its C-SPOC facility, so that the administrator can manage users clusterwide. C-SPOC handles the replication task and can provide a single view of all users on the cluster.

Sun's replication is less refined, as it does not completely automate replication. Instead, Sun's Cluster Console interface streamlines this process. It displays one

terminal window for each cluster node, plus a small common window that can be used to control all windows simultaneously.

TABLE 2: Testing Cluster-Aware Applications

	System Programming	Time-Consuming Element				
Wt[1]	**Activity**	**Compaq**	**Previous Generation**	**HP**	**IBM**	**Sun**
18	Initial data entry	✓	✓	✓	✓	✓
6	Ensure that all nodes are up and running	×	✓	✓	✓	✓
47	Manual copy onto 8 nodes	×	✓	×	×	×
24	Assisted copy onto 8 nodes	×	✓	×	×	✓
6	Verify that changes have been made across all nodes	×	✓	✓	✓	✓
100						
Time spent compared to previous generation		**18%**	**100%**	**30%**	**30%**	**54%**
Time saved by using Compaq			**82%**	**40%**	**40%**	**66%**

[1] Weight – the percent of total time devoted to each activity

To summarize, Compaq administrators spend less time on tasks such as user administration compared to administrators of previous generation clusters and compared to administrators of competing clusters in today's market. Compaq administrators save 70% of the time required by previous generation software, 25% of the time required by HP and IBM cluster administrators, and half the time required by Sun administrators to perform user administration on an eight-node cluster.

The cluster file system provides special benefit to creating and especially to testing extended custom scripts and/or system programs. With the previous generation of UNIX clusters and with current competitors, the administrator or programmer would have to create a routine, replicate it across all nodes, and then test. The administrator or programmer would then add a few more lines to the routine, replicate across nodes, and test. If something did not work, the administrator would have to fix the problem, replicate across all nodes, and test, and so on. In such a situation, synchronize-and-verify operations can take longer than the coding itself.

For iterative tasks that required clusterwide consistency, Compaq's CFS provides the greatest benefit in terms of time saved. When testing cluster-aware applications, for example, the activities involved in replicating across nodes can take four times as long as the original task. Compaq administrators save an estimated 82% of the time required by users of previous generation software. Similarly, Compaq administrators save 40% of the time required by HP and IBM administrators and two thirds of the time required by Sun administrators.

TABLE 3: Add Node

	Add Node	Time-consuming Element				
Wt[1]	**Activity**	**Compaq**	**Previous Generation**	**HP**	**IBM**	**Sun**
5	Ensure that all nodes are up and running	×	✓	✓	✓	✓
5	Create boot disk	✓	✓	✓	✓	✓
20	Install network services	✓	✓	✓	✓	✓
20	Install operating system	×	✓	✓	✓	✓
5	Notify all nodes on cluster	×	✓	✓	✓	✓
20	Reboot node	✓	✓	✓	✓	✓
5	Verify that changes have been made across all nodes	×	✓	✓	✓	✓
20	Stop/restart cluster	×	✓	×	×	✓
100						
Time spent compared to previous generation		**45%**	**100%**	**80%**	**80%**	**100%**
Time saved by using Compaq			**55%**	**44%**	**44%**	**55%**

[1] Weight – the percent of total time devoted to each activity

The CFS is particularly valuable in keeping consistent and current versions of the operating system, clusters, and application software across all nodes. Since all systems work from the same image of the software, administrators can apply patches, upgrades, and customization features once with the assurance that the changes will take effect across all nodes and in all situations, regardless of where the process is executing and regardless of the node onto which the application has failed over. This prevents odd ends or possible glitches in the propagation, a fact that not only saves keystrokes and waiting periods as changes are propagated, but also eliminates the potential for unwitting errors in the synchronization process that will cause the system to fail. This kind of improvement proves particularly important in active environments, where new IT projects are implemented frequently.

Comparing procedures for adding a node to the cluster clearly illustrates CFS efficiency compared to competing approaches. HP and IBM have implemented dynamic reconfiguration of this task to eliminate rebooting the cluster, providing some relief compared to the previous generation of software. All vendors make it possible to clone the installation of network services and operating system, reducing the weight assigned to these activities lower than it otherwise might be. Compaq streamlines the process of adding a node to less than half the levels required by previous generation software or by Sun software. It reduces the levels by almost half compared with HP and IBM.

Compaq administrators spend less than half the time adding a node compared to competing cluster systems. They save 44% compared with HP and IBM administrators and 55% compared with Sun administrators.

OTHER MANAGEMENT TASKS

The benefits derived by the cluster file system – along with cluster-wide /usr and /var – extend across the range of cluster operations and administration tasks. Variations in percentages of time saved will depend on the proportion of initial entry to the total job. For tasks that need considerable up-front planning, the efficiency derived from eliminating the synchronize-and-verify operations declines, though it never disappears. For these tasks, Compaq's advantages in planning and configuring, which derive from the cluster alias, load balancing, and policy-based configuration, often become more prominent.

CHAPTER 3. FAILOVER

TABLE 4: Failover

	Event or Activity	Penalty Hours				
Wt	**Type of Failover**	**Compaq**	**Previous Generation**	**HP**	**IBM**	**Sun**
10	Hardware-related failover (spontaneous failover)	5	10	5	5	5
10	NFS-related failover	4	10	5	5	7
10	Inconsistent configuration files across nodes	1	10	2	2	4
30						
Time spent compared to previous generation		**10**	**30**	**12**	**12**	**16**
Time saved by using Compaq			**66%**	**17%**	**17%**	**37%**

Failover is a time-consuming process that involves not only relocating the applications but also unmounting filesystems, importing and exporting file systems, starting and stopping the fault monitor, and so on. Even getting users to log off is a time-consuming process. About a third of administrators surveyed reported greater than 12 failovers a year. Over half reported that failover took 30 minutes or longer. This means reducing the number of failovers can result in significant savings in administrative time.

Hardware considerations such as the storage subsystem and the interconnect are important determinants in failover time. A standard Ethernet connection can take 10 times longer to fail over compared to a high-speed interconnect for the same size database. Failover of a 45GB database can take from two to three minutes over a high-speed interconnect such as Compaq's Memory Channel, versus half an hour using standard Ethernet connections. However, even on a high-speed interconnect, glitches in the failover process can offset any potential gain in high-speed equipment, since they must be diagnosed and fixed. All vendors offer a high-speed option.

Application state is another determinant in the time required to fail over. One administrator, for example, reported that if his database were in backup mode at the time of failure, it did not come up cleanly and may even require a restore of the archive logs, obviously a time-consuming effort. This type of application-specific problem must be addressed in cooperation with the ISV.

Table 4 summarizes DHBA's findings with respect to the three most common types of failover:

- Spontaneous hardware-related failover
- NFS-related failover
- Failover that results when configuration files do not properly synchronize across nodes.

DHBA cannot predict the percentage of occurrences for each failover type based on current data. For this reason, all received an arbitrary weight of 10 in Table 4. Individual cluster administrators may adjust these weights to represent the characteristics of their own environment.

SPONTANEOUS FAILOVER

Some users experience spontaneous failover from one node to another with no apparent cause beyond a noisy network or a spike in a piece of equipment that would otherwise have caused no harm. When administrators report repetitive failures, sometimes as frequent as once or twice a month, the cause is most often hardware related, such as a faulty processor, loose connection, or noisy network.

With each software release, all vendors are improving their capability to reduce the number of unnecessary spontaneous failovers. They are fine-tuning system thresholds to avoid unnecessary failover and providing mechanisms that enable the user to gain control over idiosyncrasies in their individual installations. Compaq provides several mechanisms for detecting failure and controlling the trigger point for failover, enabling the administrator to:

- Automate responses to system and cluster events through the Event Manager (EVM)
- Specify a check interval for the system to check itself for pending failures
- Specify a failure threshold – a tolerance level of failures before triggering a failover
- Specify a failover delay – an interval of time to wait before triggering a failover, allowing the node a chance recover spontaneously

Competitors offer equivalent capability.

DHBA estimates these features reduce hardware-related failovers by 50% compared to previous generation clusters for all vendors.

NFS

All NFS users reported that NFS filesystems cause problems in failover. A main cause of difficulties in failover situations derives from the reluctance of the failing system to relinquish resources, especially mount points, so that the administrator must unmount them manually. This costs precious time when the administrator is under pressure. The survey identified this problem as the single most common cause for delayed failover.

TruCluster Server 5.0 addresses this dilemma at the architectural level by implementing a cluster file system. Filesystems remain mounted even when a node goes down and visible to all other nodes in the cluster, enabling automatic failover. Filesystem mount points are visible to all nodes in the cluster. As a result, the failover is smooth.

HP has recently replaced its NFS automount with a new facility. IBM has also redesigned its NFS capability (not to be confused with its HANFS product, which is not compatible with HACMP). The jury is still out on this latest generation of clustered NFS software, since not many systems have been put in production mode. However, Compaq seems to provide the most direct and efficient solution, whereas other vendors deliver workarounds. Sun keeps the export file on the shared cluster filesystem, but customers still complain of NFS problems.

TruCluster Server 5.0 delivers the additional benefit of the cluster alias to NFS servers. Clients use the default cluster alias as the name of the NFS server when mounting file systems exported by the cluster. That is, clients have access to NFS filesystem both before and after failover. If a node goes down, the system automatically and transparently routes subsequent requests to the failover node; users remain connected to the cluster and do not need to re-establish their connection.

DHBA estimates that increasing familiarity with this issue and regular releases of new solutions have reduced occurrences of NFS-related failover by as much as 66% compared to the previous generation of cluster software. Compaq delivers the greatest relief, HP and IBM tie for second place, and Sun places last.

UNSUCCESSFUL FAILOVER

The penalty for unsuccessful failover is considerably higher than for successful failovers. When configuration files are out of sync, the failover will not succeed. Almost every administrator interviewed admitted to unsuccessful failovers that brought their system down due to errors in synchronizing the data across nodes. Resolution of unsuccessful failover can take from three hours to five days. In one instance, the administrator even misdiagnosed and bought a new power supply, erroneously believing that to be the cause. System administrators report that

errors in the propagation process have led to unsuccessful failovers that took over an hour to fix.

While the risk of unsuccessful failover is low – averaging fewer than five per year once the system is installed and stabilized – the penalty for downtime is high. Administrators respond to this situation with rigorous monitoring and checking procedures, which can be manual or scripted. Some immediately call the vendor.

Approximately 60% of installations perform routine maintenance to minimize the risk of unsuccessful failover. These installations check the logs daily for unusual events (15 minutes per day for a two-node cluster), and they conduct a monthly exercise of the failover process – along with other maintenance and routine checks – to ensure it is working properly (four hours per month).

IBM attempts to control synchronization through its verification process. A window of vulnerability still exists, however, since the operator may forget to perform verification. Similarly, HP addresses the synchronization problem through automation via SAM. Nevertheless, unsuccessful failover due to faulty synchronization remains a threat whenever a system requires multiple copies of a file. CFS reduces the potential for human error compared to clusters that require the synchronization process. CFS does this by definition, since it requires just one copy of the configuration information and the software.

DHBA estimates that Compaq TruCluster Server 5.0 eliminates 90% of unsuccessful failovers due to faulty synchronization. The remaining 10% comes from the several network services, such as NIS and DHCP, that remain node-specific. Compared with HP and IBM, Compaq reduces the work effort arising from failovers by 17%. Compared with Sun, it reduces the work effort by 37%.

CHAPTER 4. CLUSTER ALIAS, LOAD BALANCING, AND CLUSTER APPLICATION ARCHITECTURE (CAA)

Most of this study focuses on ways in which TruCluster Server 5.0 streamlines cluster operations. TruCluster Server can also automate configuration planning, however, to minimize hard-coding applications with specific hardware. DHBA estimates this capability will initially reduce the planning and definition of the cluster by 30%. It will further eliminate the need to re-configure the cluster to accommodate changing user requests. Even when a new node is added, it will be assigned processes according to the user-selected placement policy in effect. Administrators do not need to explicitly assign applications or resources to the new node.

In two-node clusters, the cluster alias simplifies the configuration of failover, eliminating a potential source of error. Clients access specific applications using the cluster alias, rather than the individual node names, and then the cluster can identify the appropriate resources no matter which node they reside on. As larger clusters of three or more nodes become common, the cluster alias feature will assume increased significance.

For the system administrator, the cluster alias simplifies both script writing and configuration planning. It finds resources, so that the administrator need not assign a specific node. It also balances the workload among processors, reducing the criticality of capacity planning. Further, it simplifies the process of making changes to accommodate changing user requirements. The cluster alias takes a step toward reducing the level of expertise required for the task, freeing the supervisor to devote more time to strategic planning for the corporation.

Cluster aliases prove most useful for clusters with more than two nodes, so this survey could not provide insight into the amount of time saved by using them. Nevertheless, it is clear that TruCluster Server eases this process by automating some decision making so the administrator does not have to assign applications to specific nodes ahead of time. Cluster aliases make initial configuration easier and eliminates the necessity to move applications around at a later time as usage patterns become more apparent. They also eliminate the necessity to hard-code cascading failover or distributed failover schemas. In all likelihood, aliases reduce configuration planning and maintenance by a minimum of 50%.

CHAPTER 5. CONCLUSION

TABLE 5: Time allocation to selected tasks

		Hours per year			
	Annual Effort	**Compaq**	**HP**	**IBM**	**Sun**
Maintenance					
User accounts	4%	26	35	35	52
Respond to failovers	10%	200	240	240	320
monitor and tune	20%	416	520	520	520
Special Projects					
Small program-testing project	6%	120	180	180	300
Add 6 nodes	1%	5	8	8	10
Plan and configure 6 nodes	1%	0	16	16	16
Other special projects	2%	36	60	60	108
Meetings	10%	210	210	210	210
Total hours	**54%**	**1,013**	**1,269**	**1,269**	**1,536**
Total days (based on an 8-hour day)	**54%**	**127**	**159**	**159**	**192**
Total Days (based on a 6-hour day)	**54%**	**169**	**212**	**212**	**256**

- User accounts: Annual effort is based on estimated one-half hour per week to key in an average of 10 updates per week.
- Respond to failovers: Annual effort is based on interview information.
- Monitor and tune: Annual effort is based on interview information. Compaq efficiencies derive from its event management system and integrated log capabilities. DHBA's Operating System Function Review presents a more detailed competitive evaluation of these capabilities.

To illustrate the impact of these efficiencies on the daily routine of system administrators, consider how the results of this study apply to a hypothetical cluster installation with 500 GB database that supports 1,000 users (20% turnover per year plus 30% changes in other user information such as privileges, passwords, and file access). During the current year, the IT department plans to expand the cluster from two to eight nodes and to write some custom system-level programs to make a special piece of equipment available to applications. The cluster is staffed as follows:

- One high-level system administrator, who spends 50% of her time meeting with business users for strategic IT planning and coordinating with vendors, 30% writing proposals and documentation, and 20% working with the hands-on administrators. She is responsible for the entire IT installation of 10 UNIX systems from Compaq, HP, IBM, and Sun, plus eight NT servers.
- Three first-level system administrators who have hands-on responsibility for the cluster, plus three additional UNIX systems. Together, these administrators devote a total of 10 man-months to the cluster.

The following analysis applies to the schedule of the first-level administrators with hands-on responsibility for the cluster.

As indicated in Table 5, when the efficiencies outlined in this report are applied to the above scenario, a Compaq administrator would complete the selected tasks in 127 days, whereas the HP or IBM administrator would require 159 days to perform the same tasks and a Sun administrator would require 192 days to complete these tasks. This translates to a savings of 32 man-days, or over 1.5 man-months compared to HP and IBM, and a savings of 70 man-days, or 3.5 man-months, compared to Sun. Note that these estimates are conservative, since they are based on an eight-hour day. If we base our calculation on the more realistic assumption that only six hours a day are truly productive, then the TruCluster Server 5.0 advantage increases by an additional 25% to 43 days, or about 8.6 weeks compared with IBM and HP, and 87 days, or over 17 weeks, compared with Sun.

Results will vary from installation to installation and year to year, based on individual corporate requirements, the size of the installation, and special projects. In all situations, however, administrators of TruCluster Servers will reap tangible benefits from the simpler system image and especially from the reduced number of files they need to maintain.