

---

# **Database Storage Solutions**

**A Competitive Study on  
Total Cost of Ownership**

# About INPUT®

## **Enabling Electronic Business Decisions™**

Formed in 1974, INPUT is a Web-based IT marketing services firm.

INPUT provides support to IT buyers and vendors through subscription and custom projects:

- **Executive Guides** - decide on what to buy and who to use for help.
- **Industry Guides** - decide on plans, strategies and investment.

**Executive Guides** provide objective information on IT issues and promotional information from INPUT clients on their relevant services and products. **Buyer's Guides** are made widely and freely available to buyers, management consultants, the media and financial community through the Web.

**Industry Guides** are market research reports with analysis and 5-year market forecasts. Emphasis is on electronic business subjects. INPUT constructs and manages **Partner** and **Testimonial** databases on the Web for specific clients. It also provides Web construction and electronic commerce services for IT companies.

### **Marketing Services**

**Web-based** services in the following areas:

**Partner Promotion  
Directories and Databases**

**Testimonial and Case Study  
Databases**

**Custom Partner Promotion  
Programs and Projects**

**Total Cost of Ownership  
(TCO) studies**

**Partner/Customer  
Satisfaction studies**

**Outsourcing of Web  
marketing support activities**

**Portal and Web-based E-  
Commerce site construction**

### **Subscription Services**

**Electronic Business**

**Electronic Government**

**IT Customer Services**

**Operational/Outsourcing  
Services**

**IT Software & Services:  
Vendor Data Base  
Market Forecasts**

### **Service Features**

**Unlimited Internet** access.

**Executive Guides** on the Web.

**Industry Guides** on the Web.

**Internet-based** client support

### **Databases**

**U.S. Federal Government**

- IT Procurement Plans
- IT Awards
- Agency IT Plans

### **INPUT Offices**

[www.input.com](http://www.input.com)

**Washington, D.C.  
14900 Conference**

**Center Drive**

Suite 225

Chantilly, VA 22182 3900

U.S.A.

Tel: +1 (703) 378-8740

Fax: +1 (703) 378-8804

### **Affiliates in:**

Australia, China, HK,  
Japan, Malaysia, Singapore  
Taiwan and other countries

### **Projects**

Individual projects include:

**Guide to SAP Services  
Providers in 10 Countries**

**Guide to Baan Services  
Providers**

**Guide to Desk Top Services**

**Guide to Outsourcing  
Vendors**

### **Custom Research**

Typical projects include:

**Opportunity Analysis  
Market/Product Strategy  
Competitive Analysis**

**Due Diligence**

**Partnering/Acquisition  
Services**

INPUT offers **partnering**, and acquisition searches, particularly in the international arena.

.....contact us at:

[www.input.com](http://www.input.com)

## Abstract

In today's 24 x 7 business environment, the need for reliable, highly available storage solutions has never been greater. Companies are highly sensitized to the impact of business downtime on lost productivity and revenues, as well as on-going support & maintenance costs which can increase total ownership costs.

This report evaluates data availability and cost of ownership for two leading database storage solutions in Oracle server software environments: Network Appliance Filer and EMC Symmetrix. The study finds that the annual total cost of ownership is 75% lower for Network Appliance than for EMC database storage solutions. Data availability for Network Appliance respondents meets or exceeds 99.0%, while less than two-thirds of EMC implementations match this service level.

Published by  
INPUT  
14900 Conference Center Drive  
Chantilly, VA 20151  
United States

### **Total Cost of Ownership Studies**

#### ***Database Storage Solutions: A Competitive Study on Total Cost of Ownership***

Copyright © 2001 by INPUT. All rights reserved.  
Printed in the United States. No part of the publication  
may be reproduced or distributed in any form, or by  
any means, or stored in a database or retrieval  
system, without the prior written permission of the  
publisher.

The information provided in this report shall be used  
only by the employees of and within the current  
corporate structure of INPUT's clients, and will not be  
disclosed to any other organization or person  
including parent, subsidiary, or affiliated organization  
without prior written consent of INPUT.

INPUT exercises its best efforts in preparation of the  
information provided in this report and believes the  
information contained herein to be accurate.  
However, INPUT shall have no liability for any loss or  
expense that may result from incompleteness or  
inaccuracy of the information provided.

CYN-NA1 • 2001 • 73

# Table of Contents

---

<b>Executive Summary</b> .....	1
A. Introduction & Methodology.....	5
B. Database Storage, Growth and Scalability.....	9
C. Storage Administration Responsibilities & Opportunities.....	11
D. Database Storage and Data Availability.....	15
E. Database Storage Total Cost of Ownership.....	19

---

# List of Exhibits

---

1. Survey Respondents by Industry Segment.....	5
2. Survey Respondents by Annual Company Revenues.....	6
3. INPUT Calculation of Cost of Ownership.....	8
4. Database Size and Storage Growth.....	9
5. Database Allocation for Application Types.....	10
6. Percent of Time Available by Type of Storage Administration Task.....	11
7. Percent of Time Available for Performance Tuning Tasks.....	12
8. Percent of Time Required for Routine Administrative Tasks.....	13
9. Percent of Time Available for Architectural Improvements.....	14
10. Importance of 24 x 7 Data Availability by Storage Type.....	15
11. Data Availability Service Levels by Storage Type.....	16
12. Percent of Staff Forced to Stop Working When Data Not Available.....	18
13. Annual Cost of Ownership per GB.....	19
14. Database Storage Total Cost of Ownership per Annum.....	20
15. Staff Requirements by Storage Type.....	20
16. Satisfaction With Product Characteristics by Storage Solution.....	21

---

# Database Storage Solutions

## A Competitive Study on Total Cost of Ownership

---

### Executive Summary

Data availability and total cost of ownership are widely debated metrics often used across the IT industry. Few would argue that in the era of e-business, never has maintaining 24 x 7 operations been more important, nor has the impact of business downtime been greater. With a variety of database storage architectures available in the market, this report analyzes service level and cost metrics of two types: network attached storage and storage area network, specifically Network Appliance Filers and EMC Symmetrix solutions.

While methodology and terminology used to evaluate rival products can vary dramatically from report to report, often obscured from readers' view, this report follows three straightforward guidelines in its research approach:

- First, report findings are based on interviews with actual users of each product in the evaluation to represent real-world operating environments. While some reports approximate user costs or sample a single buyer, this report is based on actual user data and a statistically significant sample size of 25.
- Second, the report focuses on access to data rather than system access at a single point on a network. This study evaluates *data availability*, which means that the entire infrastructure, including the file system and application server, must be available to count in performance levels. System availability, although near perfect across the industry, often precludes access to the actual data, as it refers to availability of the disk subsystem only.

- Finally, the report uses a straightforward approach to calculate total cost of ownership, which is explained in detail in Section A of this report. Worth noting, even though INPUT has analyzed survey data conservatively, the results show a considerably favorable tilt toward Network Appliance. Skeptics may fault the small sample size (25), but they are statistically valid. A larger sample size may yield somewhat different results; however, they are unlikely to overturn the conclusions reached in this report.

The overall conclusions of this study are:

**Database Storage, Growth, and Scalability**

- Although databases stored in Network Appliance and EMC environments are roughly equal in size, respondents reported more rapid growth in Network Appliance environments by a margin of 20%. Due in part to Network Appliance operating environments closing the gap in average database size, Network Appliance is well-positioned to step out onto the leading edge of industry averages for database size overall.
- In calculating the total cost of ownership, there is a direct relationship between the amount of downtime and total costs. One source of downtime, that encountered to scale the database, can cost valuable planned downtime. The study showed that to scale up by 200 GB, Network Appliance storage administrators experienced downtime of 10 minutes, or less, compared to four hours for EMC storage administrators.

**Storage Administration Responsibilities and Opportunities**

- Database Storage administrators spend valuable time on a wide range of tasks. With talent in short supply, maximizing the use of IT staff time is an important factor in operating the database and directly relates to ownership costs. For backup and recovery tasks, more than 35% of total time is spent on these tasks in EMC Symmetrix environments compared to less than 10% in Network Appliance Filer environments.

- On average, storage administrators devoted more than 63% of total time to value-added tasks related to performance tuning in Network Appliance environments, compared to 50% in EMC environments. This includes both database and application tuning.
- Remaining time includes time for implementing new systems, clustering, patching, user support, and version changes. On average, storage administrators in Network Appliance environments were afforded 25% of total time for these tasks, compared to 15% in EMC environments.

### Database Storage Performance and Availability

- Never before have 24 x 7 operations been more important than in today's operating environments. The Internet and e-business necessitates continuous operations and increasingly cuts into productivity and revenues when interrupted. While storage administrators in both environments rated 24 x 7 database availability at the same level of importance, Network Appliance storage administrators were *more satisfied* with their ability to deliver expected database service levels than EMC storage administrators, with satisfaction scores of 4.8 and 4.2, respectively.
- Data availability and system availability differ in that one measures the access to *data* while the other often only measures access to a single *machine*. The evaluation found that *data availability* for Network Appliance meets or exceeds 99.0%, while less than two-thirds of EMC storage administrators match this service level.

### Database Storage Total Cost of Ownership

- This study follows a straightforward, comprehensive approach to calculating total cost of ownership. All factors are included in the calculation – product; implementation; support, operations & maintenance; and downtime costs. The methodology is described in detail in Section A of this report.

- The evaluation shows that the *per GB* annual total cost of ownership is 77% lower for Network Appliance than for EMC. When the business cost of downtime is excluded from the calculation, the total cost of ownership per annum is 65% lower for Network Appliance than for EMC. The total cost of ownership per annum (including business cost of downtime) is 75% lower for Network Appliance than for EMC. Figures for total cost of ownership are presented in Section E.

**A**

**Introduction & Methodology**

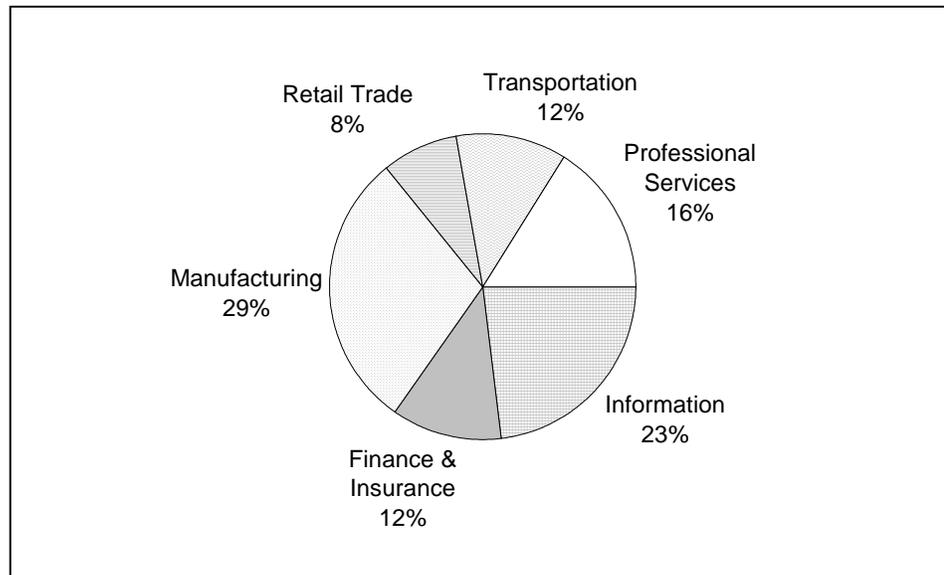
This report compares network attached storage and storage area network database storage architectures. INPUT surveyed 25 users of Network Appliance Filer and EMC Symmetrix solutions on data availability and cost factors of total cost of ownership. Respondents were asked to provide data availability service levels, quantify time spent on database administration tasks, and describe backup and recovery procedures. Additionally, respondents were asked to provide product; implementation; support, operations & maintenance; and downtime costs.

Interviews were carried out across a broad range of enterprise customers in various industry sectors to give a representative mix of industries.

Exhibit 1 shows the distribution of respondents by NAICS industry sector.

Exhibit 1

**Survey Respondents by Industry Segment**



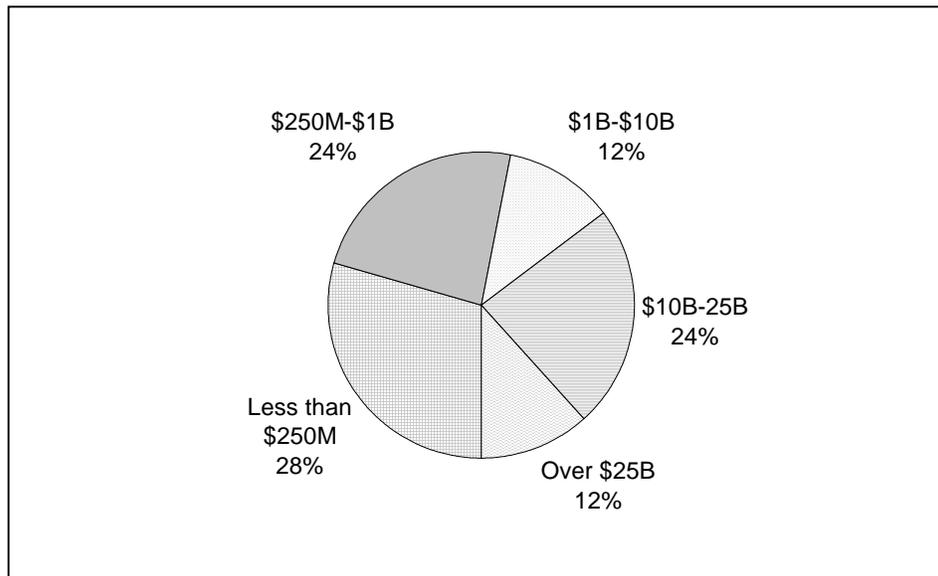
Source: INPUT

The project scope required that all databases included in the study operate on Oracle database server software, and respondents purchased either EMC Symmetrix or Network Appliance Filer storage solutions. The distribution of EMC to Network Appliance respondents was approximately 50:50.

Exhibit 2 shows the distribution of survey respondents by size of annual revenues.

Exhibit 2

**Survey Respondents by Annual Company Revenues**



Source: INPUT

**Methodology**

To ascertain the costs of ownership for each database storage solution, 25 phone interviews were carried out with Database and Systems Administrators. Each interview referred to an individual database within an organization and collected measures of cost of ownership and effectiveness for that database and storage solution.

Data collected within the total cost of ownership included:

Product costs: In addition to database product costs, these included all related product costs such as cost of underlying hardware and expenditure on database tools. Also included within this category are:

- Maintenance costs (annual expenditure on maintenance for all hardware and software associated with the database)
- Upgrade costs (includes the cost of hardware upgrades, additional license fees, and any additional implementation and training costs resulting from the upgrades to this database)

Database implementation costs (costs, mainly manpower, associated with the initial database implementation both internal and external).

Also included within this category are:

- Training costs (includes the cost of training both IT and user personnel associated with the database)

Operations costs (including all costs of IT personnel, user personnel and external expenditure incurred on server and network operations associated with the database, whether on central or local sites). Also included within this category are:

- Support costs (including the cost of personnel and external expenditure associated with performance monitoring and tuning for the database, and help-desk costs associated with the database)

Cost of downtime (includes the direct cost of downtime in terms of additional support cost and estimates of the indirect cost of downtime in terms of cost to the business).

Annual costs of ownership were calculated as shown in Exhibit 3.

Exhibit 3

**INPUT Calculation of Cost of Ownership**

Item	Calculation of annual cost of ownership
Software costs	Total software cost divided by expected life of database in years
Hardware costs	Total hardware cost divided by expected replacement cycle of hardware in years
Implementation costs	Total implementation cost divided by expected life of database in years
Training costs	Total training cost divided by expected life of database in years
Upgrade costs	Total upgrade cost to date divided by current age of database in years
Maintenance costs	Calculated as annual figures
Operations costs	Calculated as annual figures
Support costs	Calculated as annual figures
Downtime costs	Calculated as annual figures
Total annual cost of ownership	Calculated as sum of above components

*Source: INPUT*

The following sections of the report provide an overview of the survey findings, including data availability and total cost of ownership comparisons.

**B**

**Database Storage, Growth and Scalability**

Storage administrators operating databases in Network Appliance storage environments manage databases of 610 GB in size on average, with actual size ranging from 40 GB to more than 1 TB. These databases are expected to grow by 50% by year-end 2001.

For EMC solutions, respondents reported average database size of 540 GB, with a range of 250 GB to more than 1 TB. EMC respondents reported a 30% growth rate by year-end.

It comes as no surprise that Network Appliance databases grew at a faster rate. Network Appliance respondents reported higher growth rates due in part to more rapid scaling of smaller databases to a size on par with industry averages.<sup>1</sup> *However, as the current growth trend continues, Network Appliance databases will assume a solid position on the leading edge of industry averages overall.*

The size and growth of database storage environments have a significant impact on database total cost of ownership. In Section E of this report, Exhibit 10 shows that Annual per GB costs are 77% lower for Network Appliance environments than for EMC, as shown in the following exhibit.

Exhibit 4 shows database size and storage growth data for EMC and Network Appliance.

Exhibit 4

**Database Size and Storage Growth**

	<b>EMC</b>	<b>NetApp</b>
Average size of database	540 GB	610 GB
Average size of database at year-end (2001)	702 GB	915 GB
Growth rate, 2000-2001	30%	50%

Source: INPUT

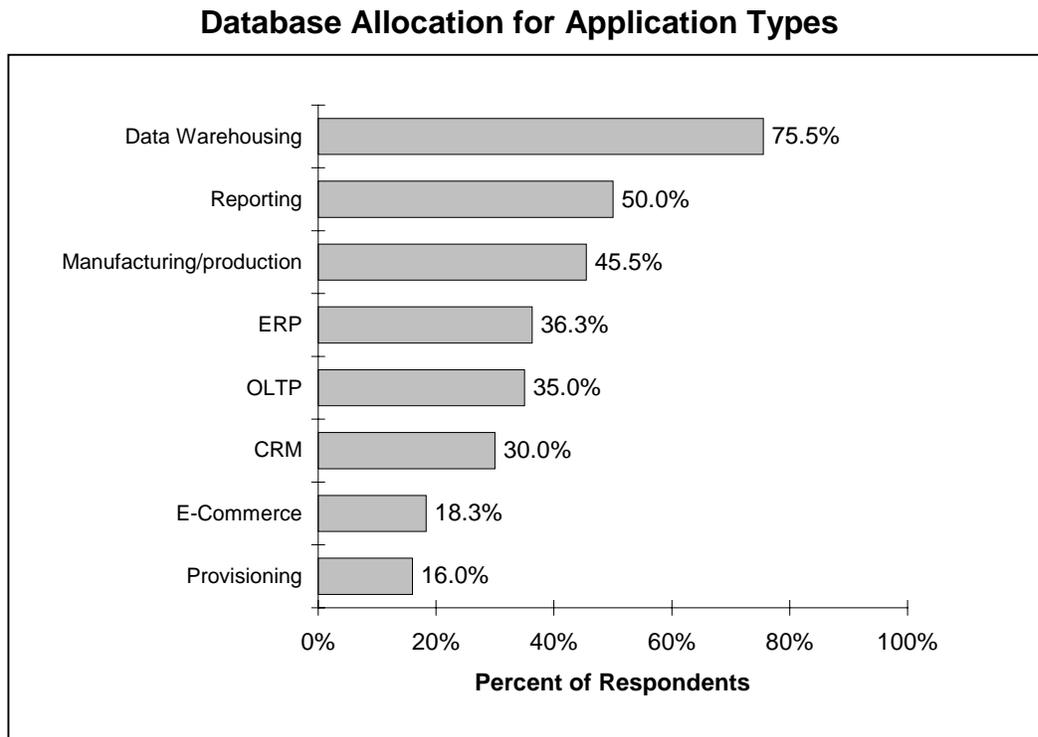
<sup>1</sup> In a 2000 report, Strategic Research Corp. reported average database size of 591 GB, with an average annual growth rate of 24.1% for mission critical databases.

Database size and growth also significantly impact scalability and availability. Storage administrators were asked about downtime involved in scaling databases by 200 GB. Network Appliance storage administrators reported 10 minutes of downtime, or less, to complete the scenario, compared to four hours, or more, for EMC.

A significant number of respondents allocated database storage for data warehousing, with 75% of respondents allocating space for this purpose. One-half of respondents allocated a portion of total database storage for reporting and manufacturing.

Exhibit 5 shows the percentage of users allocating database resources by application.

Exhibit 5



Source: INPUT

Respondents also allocated storage for GIS software, document fulfillment, financials, and Web content.

**C**

**Storage Administration Responsibilities and Opportunities**

Database storage administrators spend valuable time on a wide range of tasks. With talent in short supply, maximizing the use of IT staff time is an important factor in operating the database, and directly relates to ownership costs. The more time spent on higher value tasks not only improves cost effectiveness, but also increases staff productivity and effectiveness.

INPUT asked respondents to approximate time spent on various tasks, including back up and recovery, application tuning, database tuning, version changes and other tasks. The next exhibit shows the responses divided into two types of tasks: value-added tasks, including application tuning, database tuning and other tasks; and routine tasks, including back up, recovery, and version changes.

Exhibit 6 shows the distribution of tasks grouped by value-added and routine tasks.

Exhibit 6

**Percent of Time Available by Type of Storage Administration Task**

Type of Task	EMC	NetApp
Value-added Tasks	59%	89%
Routine Tasks	41%	11%

*Source: INPUT*

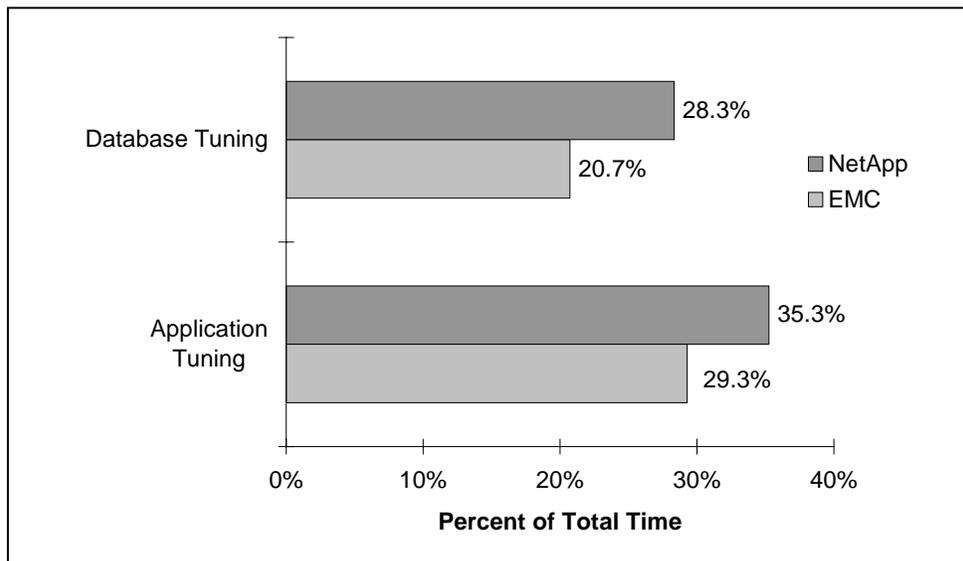
Network Appliance storage administrators spend nearly 90% of their time on value-added tasks, compared to nearly 60% for EMC storage administrators. The remainder of available time is spent on routine tasks. The breakdown of tasks is further analyzed in this section and presented in the following exhibits.

Performance tuning includes both database and application tuning. Time spent on these tasks contributes significantly to the overall effectiveness and stability of the database environment. On average, Network Appliance storage administrators were afforded more time, on average 63% of total time, for these critical, high-impact tasks. This compares to only 50% of average total time for EMC storage administrators.

Exhibit 7 shows the percent of total time available for performance tuning tasks for each product type.

Exhibit 7

**Percent of Time Available for Performance Tuning Tasks**



Source: INPUT

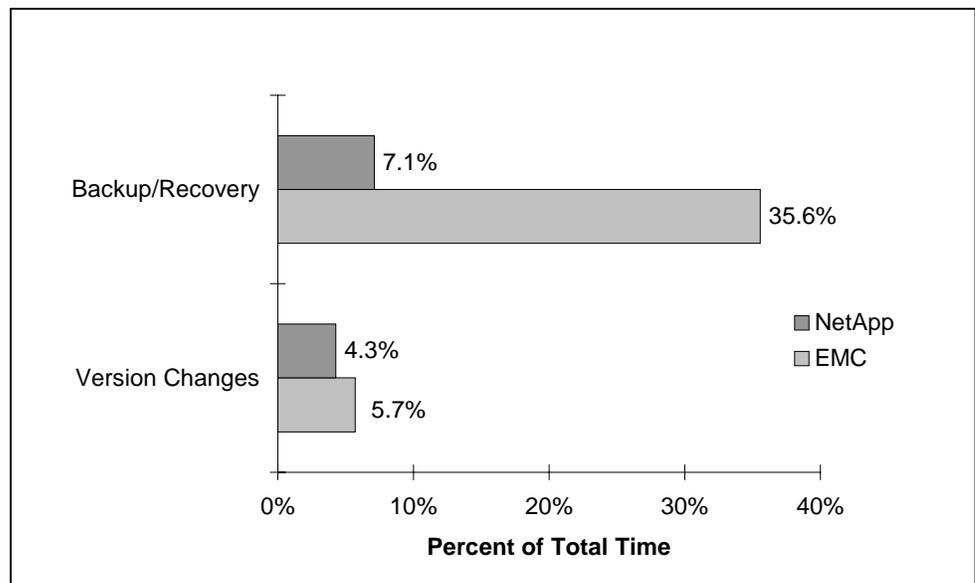
Database administrators must also spend a portion of their valuable time performing routine tasks, such as back up, recovery, and version changes. While critical to business continuance, the amount of time required here can impact the ability to leverage more valuable database administrator skills. Survey respondents indicated that to the extent they were able to reduce time spent on routine tasks, they could spend more time on other more satisfying, higher value tasks.

In EMC storage environments, more than 35% of the administrators' time is spent on back up and recovery tasks, compared to less than 10% in Network Appliance environments. For version changes, although the difference is minimal, the time spent on these tasks is somewhat less for Network Appliance than for EMC.

Exhibit 8 shows the percent of total time required for routine storage administration tasks.

Exhibit 8

**Percent of Time Required for Routine Administrative Tasks**



Source: INPUT

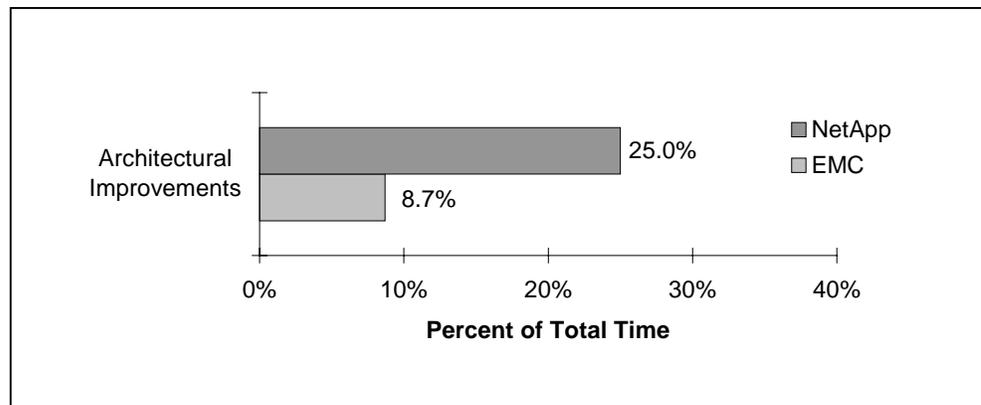
Despite a lesser time requirement for routine administration activities, Network Appliance administrators achieve a higher level of business continuance, as presented in the next section on data availability. Further, Network Appliance administrators were able to run automated backup tasks more frequently than EMC administrators, on average five times per day compared to 1.5 times per day. EMC storage administrators reported spending considerable backup and recovery time on file system recovery and reconfiguration.

Importantly, Network Appliance storage administrators reported 25% of total time available for other tasks, including clustering, implementing new systems, projects, planning, patching and security. All of these tasks contribute to improving the storage architecture and advancing the database infrastructure.

Exhibit 9 shows the percent of time available for database storage architectural improvements.

Exhibit 9

**Percent of Time Available for Architectural Improvements**



Source: INPUT

An information industry dotcom respondent said this of Network Appliance:

*“We spend 90% of our time testing, and 10% monitoring. We are mostly focused on projects and assisting developers. The production machines are stable.”*

Attesting to its reliability, the respondent added, “(The Filer) runs itself. We’ve never lost a file that was stored on a NetApp.”

The next section covers database storage performance and availability.

**D**

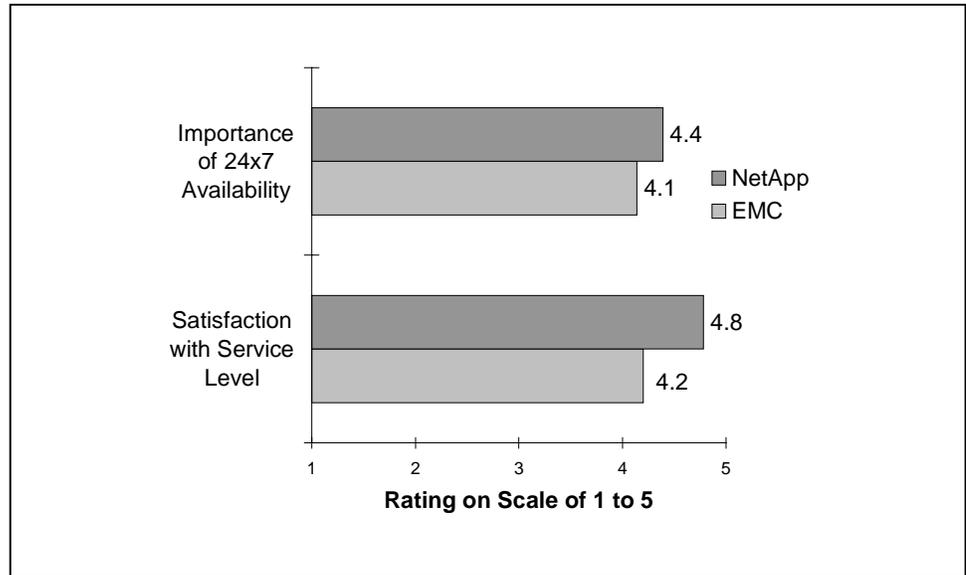
**Database Storage and Data Availability**

Respondents were asked to rate the importance of 24 x 7 data availability to their organization on a scale of 1 to 5, with 1 being of no importance and 5 of extremely high importance. Both EMC and Network Appliance administrators alike rated the importance high. In a few instances where access patterns followed typical office hours, respondents still had service level requirements as if they were operating 24 x 7.

Exhibit 10 shows the importance of 24 x 7 data availability by storage type.

Exhibit 10

**Importance of 24 x 7 Data Availability by Storage Type**



Source: INPUT

Although 24 x 7 rated highly important, *overall satisfaction with actual service levels varied for Network Appliance and EMC*. Network Appliance administrators rated a near-perfect satisfaction level of 4.8, compared to an above average satisfactory level of 4.2 for EMC.

Respondents also provided information about actual database service levels for data availability.

Here, it is important to note the distinction between “system” availability and “data” availability. *System availability* refers to availability of the disk subsystem only, while *data availability* refers to the availability of the entire infrastructure, including the file system and application server. For purposes of this study, a system is counted as available if the disks have power and spin, even though the host server or file system may not be available. *Data availability requires end-to-end availability, including all hardware, software, attached storage and networks.*

For Network Appliance environments, data availability service levels ranged from 99%, and above. In EMC environments, respondents reported data availability service levels ranging from 95-99.5%.

Exhibit 11 shows database service levels for EMC and Network Appliance administrators participating in the survey.

Exhibit 11

**Data Availability Service Levels by Storage Type**

Database Service Level	EMC	NetApp
Over 99.5%	-	67%
99% to 99.5%	57%	33%
97% to 99%	29%	-
95% to 97%	14%	-
Less than 95%	-	-

Source: INPUT

Less than two-thirds of EMC implementations matched Network Appliance data availability metrics of 99% and above.

The impact of downtime on an organization can range from a mere disruption to productivity internal to the organization, to a complete work stoppage that extends beyond the boundaries of the organization to clients and suppliers.

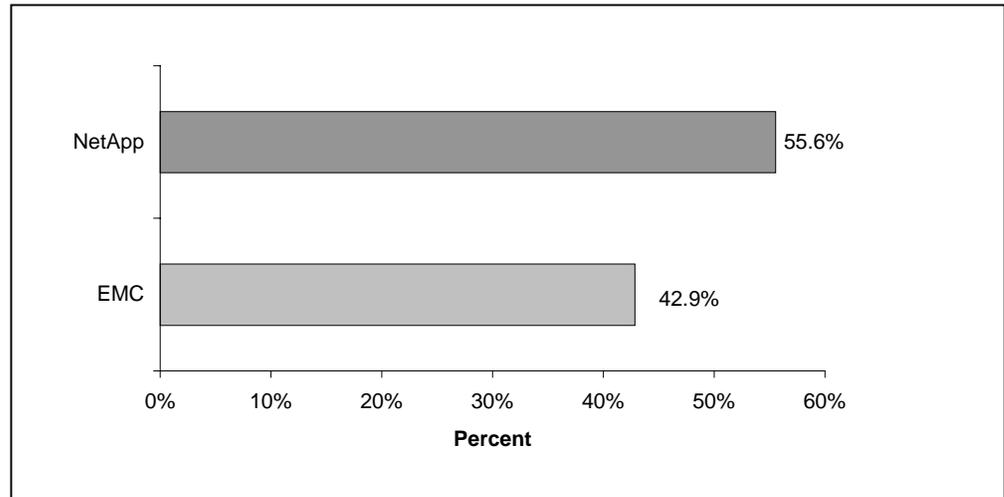
The trend is toward the later scenario, with the Internet creating new ways of doing business every day and increasing reliance on information technology. Averting downtime becomes increasingly important in interconnected E-business environments.

INPUT asked respondents to indicate whether downtime would merely be an inconvenience or if it would force people to stop working. The greater the impact on productivity, the greater the need for highly reliable database storage solutions that provide high data availability.

Exhibit 12 shows the impact of downtime on productivity.

Exhibit 12

**Percent of Staff Forced to Stop Working When Data Not Available**



Source: INPUT

More Network Appliance solutions are implemented to meet requirements that extend beyond the boundaries of the company and that directly affect staff productivity. Of each product type implemented, more than 55% of Network Appliance storage solutions were implemented for this purpose, compared to 43% of EMC storage solutions.

A large retail trade industry user of EMC storage said this about restoring the database:

*“We would only know after deluge of calls to helpdesk reporting problems. Once the 'down' server is pinpointed only then do we know about it. This often takes a long time, generally a restore can be over in 3 hours.”*

On average, time to restore from a system outage is 1.5 hours for NetApp and four hours for EMC. Typically, the time advantage resulted from the use of Network Appliance product features such as Snapshot and SnapRestore.

The previous three sections covered issues of database scalability, availability and downtime, as well as the time storage administrators spend on various tasks. The next section covers database storage total cost of ownership.

**E**

**Database Storage Total Cost of Ownership**

The methodology for calculating the total cost of ownership is detailed in Section A. The calculation includes product, implementation, operations/maintenance/support, and downtime costs.

In the survey, respondents were asked to estimate actual costs for each component. The results reflect actual results in user environments, and closely approximate industry averages.

Exhibit 13 shows the annual cost of ownership per GB for each storage solution.

Exhibit 13

**Annual Cost of Ownership per GB**

	EMC	NetApp
Average size of database	540 GB	610 GB
Total annual cost per GB, including business cost of downtime	\$31.0K	\$7.0K
Total annual cost per GB, minus business cost of downtime	\$7.7K	\$2.8K

Source: INPUT

The total cost of ownership per GB per annum (including business cost of downtime) is *77% lower for Network Appliance than for EMC*. TCO per GB per annum is calculated by dividing the annual TCO and by the average database size.

The total cost of ownership per GB per annum (excluding business cost of downtime) is *65% lower for Network Appliance than for EMC*. This calculation excludes the downtime component.

Exhibit 14 shows the annual total cost of ownership by TCO component. The results are actual results based on user estimates collected in the survey.

Exhibit 14

**Database Storage Total Cost of Ownership per Annum**

	EMC	NetApp
Product (Hardware, Software, Upgrades)	\$1,575K	\$750K
Implementation (Installation, Training)	\$425K	\$110K
Support, Operations & Maintenance	\$2,100K	\$875K
Downtime <sup>2</sup>	\$12,500K	\$2,500

Source: INPUT

*The total annual cost of ownership is 75% lower for Network Appliance than for EMC.*

Significant factors contributing to the lower TCO includes database size, implementation time, maintenance costs, downtime, and number of full-time equivalent administrators.

Exhibit 15 shows the average full-time storage administrator requirements for each solution.

Exhibit 15

**Staff Requirements by Storage Type**

	EMC	NetApp
Average full-time DBA and Systems Administrators	10.4	3.3

Source: INPUT

<sup>2</sup> Assumes 20 hours of annual downtime in EMC environments and 4 hours of annual downtime in Network Appliance environments.

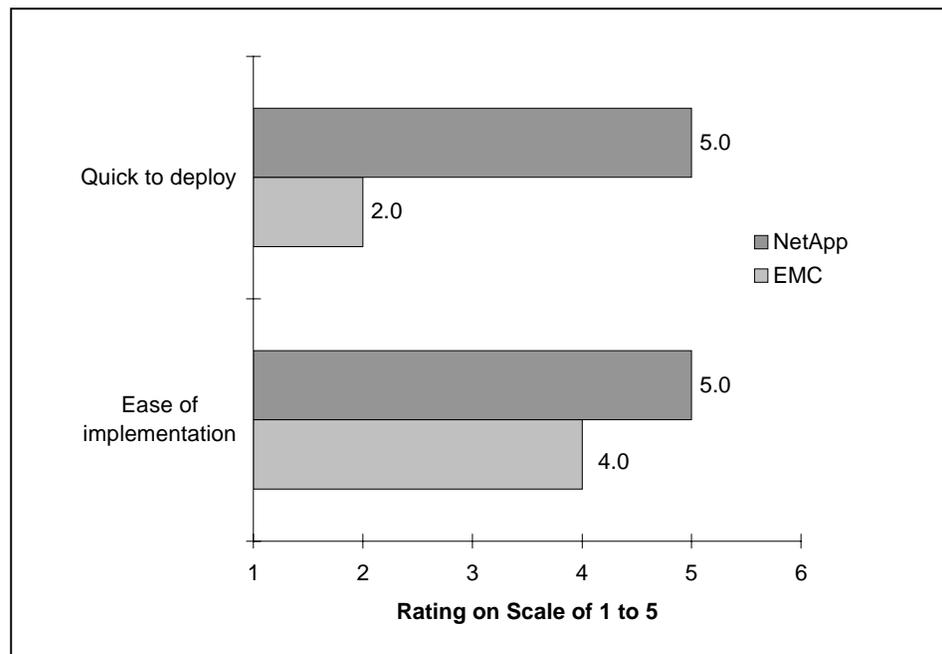
The number of full-time equivalent storage administrators is on average 68% lower for Network Appliance environments than for EMC environments.

Implementation time in EMC environments averages 420 days, compared to 15 days in Network Appliance environments.

For each of the following implementation characteristics, respondents provided the ratings shown in Exhibit 16.

Exhibit 16

**Satisfaction with Product Characteristics by Storage Solution**



Source: INPUT

The most dramatic difference between the two products was with the time to deploy the solution, with Network Appliance receiving a perfect 5.0 rating and EMC a less than satisfactory rating of 2.0.