Green Data Center Storage
Part I

Sustained Cost-Effectiveness in Transitioning Times

Bruce Naegel
John Colgrove
W. David Schwaderer

September, 07

Executive Summary

Why You Should Read This Paper

This is the first part of a three-part series dealing with data center power and cooling issues that either are presently, or will soon, affect virtually every enterprise data center. The series outlines the challenges and Symantec solutions.

In May, 2006, IDC reported its 2005 data center survey results in a telebriefing titled *End User Perspectives on Server Blade Adoption*[1]. This telebriefing indicated that, in order, the top three data center issues survey qualified respondents felt they faced were[2]:

1. Power provisioning
2. Floor Space
3. Power Consumption

Exacerbating matters, respondents also expected both processing requirements and power requirements to increase – at 25 percent for power by 2009 while other independent estimates suggest approximately a doubling every five years.

Subsequently, on November 29, 2006, Gartner, Inc. issued a news release[3] titled:

50 Percent of Data Centers Will Have Insufficient Power and Cooling Capacity by 2008.

Power, Space, and Cooling – exceed any of these three constraints within data centers and enterprises may well face building new data centers, sometimes for as much as $1,000 per square foot.

Hardware power and cooling problems are clearly a major and growing concern.

Yet, today, Gartner research vice president Rakesh Kumar estimates traditional data centers typically waste more than 60% of the energy they use to cool just their equipment.[4]

The good news is that effective software utilization can significantly help reduce hardware power consumption and consequent cooling requirements. While it appears that more energy-efficient power supplies, processors, chipsets, and cooling solutions are beginning to, will eventually, address these near-term problems, in the short term, enterprises must vigilantly oversee data center power and cooling issues. To that end:

Symantec contends its software products can help customers reduce data storage energy consumption by as much as 50 percent and total data center energy consumption by as much as 25 percent. Moreover, Symantec software allows enterprises to exploit new energy-efficient servers and storage more easily, enabling them to continue conserving energy optimally.

[1] http://www.idc.com/getdoc.jsp?containerId=TB20060525
[2] http://poland.emc.com/techlib/pdf/H2402_power_efficiency_storage_array_wp_ldv.pdf
[3] http://www.gartner.com/it/page.jsp?id=499090
[4] http://www.gartner.com/it/page.jsp?id=498224

This paper explains how CIOs can potentially save tens of millions of dollars in advance of potential governmental data center energy conservation mandates and taxation penalizations.

Symantec invites enterprises to request a free ROI analysis on the total savings Symantec software can bring to their specific data center facilities. We look forward to the ensuing engagements.

Executive Overview

Today's data center professionals are contending with unprecedented challenges that were largely unforeseen a short time ago. Collectively, the challenges invariably involve Data Center electrical energy, namely its:

- Decreasing Reliability
- Decreasing Availability
- Increasing Cost Per Watt
- Consumption Patterns
- Generated Heat Disposal

The situation is so pressing that enterprises increasingly find they cannot obtain additional power at any cost.[5]

Indeed, power and cooling limitations are now usually the greatest data center growth impediment in addition to floor space availability. The problem is highlighted by a November 7, 2006 *Business Week* article titled *Energy Could Eat Up to 40% of IT Budgets*:

> *Large organisations spend between four and eight per cent - and sometimes as much as 10 per cent - of their IT budget on energy and Gartner predicts this will rise by up to four times within the next five years.*[6]

One 2006 data center survey conducted by power supply vendor Liebert found:

- 33% of respondents expect to be out of power and cooling capacity by the end of 2007
- 96% stated they would be out of capacity by 2011.[7]

Unfortunately, traditional data center provisioning methodologies and simplistic rules-of-thumb that worked in the past are now proving inadequate – both economically and operationally. Worse, they likely provide a direct path to unpleasant near-term enterprise surprise as well as long-term disruptions that could be highly visible both internally and externally.

Industry consensus suggests that, if not efficiently and effectively addressed, these challenges will inevitably limit data center growth and flexibility that enterprise competitiveness requires. Indeed, the IT industry must now solve these many significant energy problems using a different level of thinking than the one used to create them.

That's some of the news. Daily headlines suggest more – namely, the situation is already volatile and rapidly getting worse. Corporate compliance regulators and adversarial, litigious interests would likely agree, at your enterprise's expense. Because significant change and its careful planning can require substantial lead-time to effect, it is important to begin addressing these obvious challenges now, particularly within data storage systems.

This Symantec white paper presents a preliminary storage energy analysis and conservation methodologies for both data center servers and storage subsystems. The discussion enables IT professionals to identify and avoid unnecessary energy consumption, with a focus on data center storage subsystems. This

[5] http://www.businessweek.com/technology/content/may2007/tc20070514_003603.htm
[6] http://www.businessweek.com/globalbiz/content/nov2006/gb20061107_447569.htm
[7] http://www.enterprisestorageforum.com/management/features/article.php/3639286

discussion enables readers to determine how much energy they can potentially save within their data centers. The *Symantec Technology Network*[8] will post subsequent updates to this paper.

Clearly, significant change is at hand, requiring strong leadership in unstructured, transitional times. The good news is that efficient energy consumption and conservation is not only economically beneficial and possible, it also provides indisputable evidence of executive commitment to social responsibility in a world of increasing ecological awareness.

[8] http://www.symantec.com/enterprise/stn/index.jsp

Table of Contents

A Brief Instructional Retrospective

From 1949 to 1966, Joseph Eichler's company, Eichler Homes, built approximately 11,000 Northern and Southern California homes. Achieving remarkable architectural acclaim and commercial success, these uncommonly attractive homes were spacious - often bathed inside with natural light from floor-to-ceiling glass rear and inner-courtyard walls. Even today, many passers-by find the homes visually appealing, if not stunning.

Beyond looks, Eichler homes included numerous innovations, including concrete floors with embedded copper pipes to circulate heated water for unobtrusive radiant heating. Best of all, the acquisition entry price was relatively modest considering the truly superb livability and privacy. To many, Eichler Homes indisputably built best-of-class dwellings in their day. Fast forward to today.

Unfortunately, Eichler Homes constructed homes in an era when monthly home electrical bills were often ten dollars or less. Moreover, contemporary nuclear power pundits suggested that nuclear energy's unlimited potential and availability would make energy so inexpensive that it would cost too much to meter it. But, as Neils Bohr once observed:

Prediction is very difficult, especially about the future.

Today, an original unmodified Eichler Home comprises an unqualified energy efficiency and monthly electric bill disaster. With no motivation to conserve energy costs, Eichler homes had roofs without insulation and used thin, single pane glass in windows and glass walls. Moreover, exterior wooden wall insulation was an option that buyers sometimes did not elect. In other words, Eichler homes heated their neighborhoods on cool days.

To rectify the energy efficiency problems, subsequent Eichler Home owners have insulated roofs and exterior walls, as well as replaced glass walls with other materials such as energy-efficient double pane glass or acrylics. To address a natural failure of the floor's radiant heating system, some owners have resorted to other heating approaches, including retrofitting new gypsum flooring with heated water pipes when they wanted to preserve the traditional Eichler ambiance.

In short, Eichler home energy efficiency problems have usually proven solvable. It's just a matter of money – maybe more than the house originally cost in present value dollars.

And, sadly, so it is with many data centers today…

Your Data Center

Many people, previously unaware of how energy inefficient original Eichler homes were, often gently chuckle when they learn the facts. However, what these people are possibly unaware of is that an analogous situation exists in most enterprise data centers. And, there are few chuckles to be had there because, if it does exist, it threatens the owning enterprise's growth and vitality.

You see, electrical inefficiency unnecessarily consumes electric energy that society now considers a shared, finite resource. In an age of increasing transparency and Internet blogging, the public's umbrage can be widespread, varied, harsh, and swift. Here, it is useful to recall that Ida M. Tarbell's investigative journalism for *McClure's Magazine* over a century ago eventually led to the government-dictated Standard Oil Company dismemberment, much to the chagrin of John D. Rockefeller – arguably history's wealthiest American.

Here, it is useful to note that the Congress's December 20, 2006 Public Law 109–431[9] authorized the U.S. Environmental Protection Agency *To study and promote the use of efficient computer servers in the United States* under its Energy Star Program. On April 23, 2007, the EPA released its Public Review Draft which is informative but cannot be cited at this time[10] because of its draft status. Finally, consider that various legislation discussions suggest banning low-efficiency incandescent lighting. It now seems certain that governmental oversight is inevitable.

From a civil litigation perspective, consider that the United States Supreme Court has now ruled that carbon dioxide and other greenhouse gases are air pollutants under the Clean Air Act, and that the current U.S. administration has violated law by refusing to limit carbon dioxide emissions.

It follows that adversarial litigants may now assert that when enterprises fail to rectify data center power waste that needlessly generates carbon dioxide; they are indistinguishable from committing a corporate governance and social responsibility dereliction of duty. It's really that serious.

The good news is that it is easy to initiate changes to eliminate data center power waste though admittedly harder to implement the changes. However, the solutions can easily pay for themselves by reclaiming underutilized data center resources. This can delay or subsidize future capital expenditures while eliminating considerable energy OPEX costs.

We begin by discussing data.

[9] http://energystar.gov/ia/products/downloads/Public_Law109-431.pdf
[10] http://www.energystar.gov/ia/products/downloads/EPA_Server_Report_Public_Draft_23Apr.pdf

Enterprise Data

Data is a deceptively simple-sounding, two-syllable word. About all that some people know about data is that enterprises derive vital information from data they must safely store in ways that permit only authorized users to read and perhaps change it. Other than that, it's a puzzlement - as it should be.

But data needs careful protection because it is stored using mechanical devices such as hard disk drives that have a tendency to fail at inconvenient times. As recent events have shown, preventable data loss can be extraordinarily inconvenient when associated with litigation discovery processes. Courts have historically adopted a predisposition against *oops* as an acceptable data loss explanation.

Finally, authorized users needs to access data for their purposes according to Service Level Agreements (SLAs) - pre-agreed upon conventions that establish response time levels between user groups and the IT organization.

Barricading Data Against Loss

Because mechanical devices fail and thereby place data at risk, the IT industry uses various approaches that allow data storage subsystems to provide a substantial immunity to the failure of one or more disks. You've likely heard of it – it's called *Redundant Array of Independent Disks* or *RAID* for short.

There are many RAID variants – RAID 0, RAID 1, RAID 1+0, RAID 5, and RAID 6 being the most well known. In the event you are unfamiliar with RAID, Appendix B contains a brief RAID overview.

For our purposes here, RAID technical details are not important. Each RAID variant has its own advantages, achieved by trading different amounts of storage space for different levels of performance and immunity to failure. But the important point is this:

> *When any RAID scheme provides protection against disk failure, there is a difference between the total amount of data an enterprise has and how much disk space the data occupies.*

RAID protection *always* requires more space than unprotected data requires.

- With simple RAID 1 support (mirroring), the required space is 200 percent (twice) what unprotected data requires.

- With RAID 5 support that protects the capacity of four disks using the capacity of a fifth requires 125% (5/4) what unprotected data requires.

Again, Appendix B has RAID technology information including a discussion on how to calculate effective storage utilization for various RAID configurations.

Harnessing Enterprise Data Growth

A primary challenge with enterprise data is that is continues to grow. This is partly a consequence of the increasing scope of the IT mission as well as regulatory compliance statutes that require electronic communication retention. It's also a consequence of data naturally increasing with record and file size increases (emails, attachments, richer documents, larger spreadsheets, etc).

One potential opportunity for substantial power savings involves the amount of inactive data existing on systems. Indeed, Symantec often finds that an enterprise has often provisioned more than five times the storage their active data requires.

As another example, email usually exhibits explosive growth in any enterprise. Moreover, because email and other electronic systems proliferate identical copies of files, it follows that one way to economize storage consumption is to use de-duplication technologies. Here, programs such as Symantec's Enterprise Vault[11] email and *Veritas NetBackup PureDisk*[12] packages can make significant contributions in this arena.

One common estimate is that enterprise data grows at a 70% compound annual growth rate (CAGR) or more. This discussion uses a 70% CAGR because that often proves conservative. However, some enterprises experience a 130% growth rate and your enterprise's growth rate likely experiences a different rate yet.

Figure 1 shows how serious a problem a conservative 70% growth rate can be by using carefully constructed triangles. Just glancing at the figure suggests that without proper data storage capacity planning and provisioning, in a few years, your enterprise's data is probably not going to fit within your storage subsystems.[13] If so, there is little time to waste, *carpe diem*.

Finally, comprehensive data center data storage assessments invariably reveal *substantial* unutilized, unallocated, hence collectively wasted, storage capacity. An excellent tool to help perform such discovery efforts is Symantec's *Veritas CommandCentral™ Storage*[14] (*CC Storage*), a tool that identifies waste and file usage patterns.

Appendix C contains a brief overview and example CC Storage screen captures.

[11] http://www.symantec.com/enterprise/products/overview.jsp?pcid=1018&pvid=322_1
[12] http://www.symantec.com/enterprise/products/overview.jsp?pcid=1018&pvid=1381_1

[13] Note that a 50% CAGR in storage capacity partially offsets the severity of the continuing data growth challenge. Thus, an adjusted *storage device capacity* growth rate would be approximately (1.7/1.5) or 13% in this discussion.

[14] http://www.symantec.com/enterprise/products/overview.jsp?pcid=1020&pvid=19_1

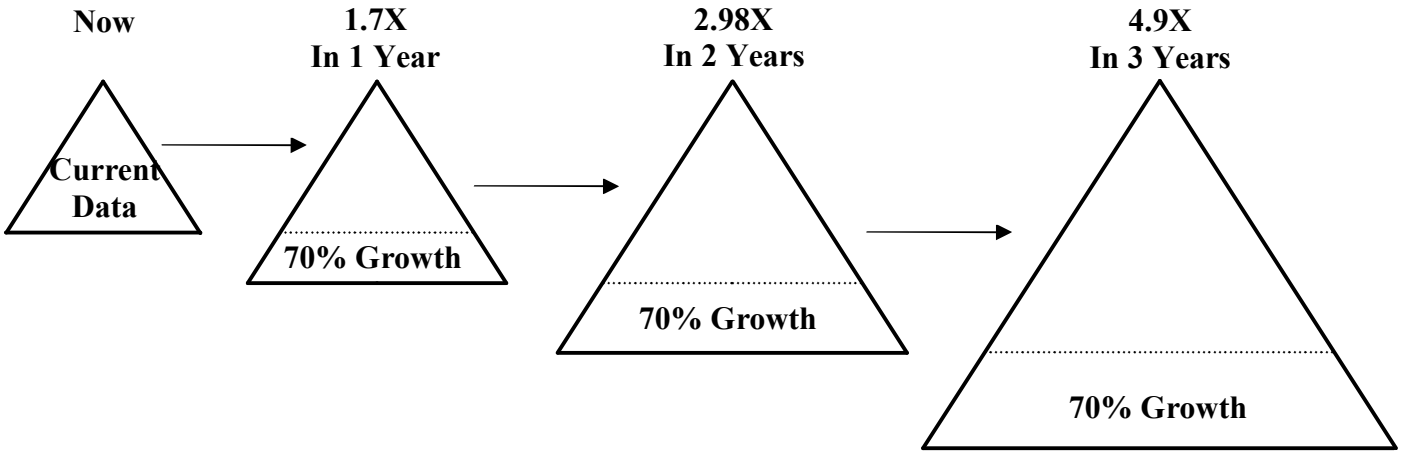

Figure 1 – The effect of a 70% data growth rate

Part I Conclusion

Enterprise data centers are under siege from a variety of challenges. Power and cooling challenges are among the most serious and promise to be chronic, long-term problems. While largely hardware caused, software plays a critical role in mitigating these challenges. Subsequent parts to this paper outline strategies and solutions that can substantially decrease data center power and cooling costs.

Appendix B – RAID Overview

There are many forms and variants of RAID. Here are some of them:

While providing no protection against disk failure, **RAID 0** spreads data uniformly across a number of disks. This is useful when many users need simultaneous access to a data such as a shared database. Here, spreading the data out across many disks allows multiple overlapped concurrent accesses though the probability of data loss increases with each additional disk used.

RAID 1 stores data in the necessary disk space and simultaneously makes one or more copies of the data on other disks. So, the multiple copies occupy much more disk capacity as the actually data values require. Clearly, there is a difference between the total amount of data an enterprise has and how much disk space the data requires.

RAID 0+1 combines RAID 0 and RAID 1. A data set resides on multiple disks and a second, independent set of disks contains mirror copies of the data. The mirroring provides protection against the loss of disks, and using twice as many disks allow twice as many concurrent accesses, greatly accelerating performance.

RAID 5 similarly spreads data across a number of disks, say four, and then uses the data to create a checksum value that can reconstruct lost data if one of the four disks fails. The storage subsystem can store the value on a fifth disk drive but usually uses the five drives together, spreading the data and checksum information across the set of disks.

Like RAID 5, **RAID 6** spreads data across a number of disks, say four, and then uses the data to create two or more checksum values that can reconstruct lost data if more than one the four data disks fails. The storage subsystem can store the value on a fifth and sixth disk drive but usually uses the drives together, spreading the data and checksum information across the set of disks.

Clearly, there is a difference between the total amount of data an enterprise has and how much disk space the data requires.

Calculating various RAID Capacity Factors

The formula to compute a RAID set capacity factor is:

(Total Drives in the RAID set) / # of Data Drives

RAID 1: RAID 1 replicates data. A 2-way mirror has 2X the total number of drives as compared to data drives. A 3-way mirror has 3X the number of drives. These

yield 2X and 3X capacity factors respectively. This scheme also applies to RAID 0+1 or 1+0 that combines striping and mirroring.

RAID 5: RAID 5 is usually a number of data drives plus one additional parity drive that completes the RAID set. A RAID set with 5 data drives and one parity drive has 6 drives total and 5 data drives worth of capacity

RAID 6: RAID 6 has 2 parity drives per RAID set. So a RAID set with 5 data drives would have 7 total drives / 5 data drives or a RAID capacity factor of 7/5.

I/O Performance Measurements:

Many factors affect I/O performance. Two primary factors are the total number of disk drives and individual disk drive performance. A future revision of this white paper will address this topic in more detail.

Appendix C – CC Storage Analysis Tool

CC Storage is a Storage Area Network (SAN) hardware management and monitoring application that discovers SAN resources - switches, Host Bus Adapters (HBAs), disk arrays, host servers, and applications.

CC Storage also has the ability to monitor file system activity and determine which files are active and inactive, hence eligible for archiving or migration to more energy efficient storage devices (Figures C-1 through C-4). In addition to understanding where the inactive files are, CC Storage can also help to reconfigure the SAN and help move the data to more appropriate storage. Data administrators can identify inactive data on Tier 1 storage and help to migrate it to the Tier 2 storage using SAN reconfiguration by:

- Opening a connection between the source and the destination
- Copying the data to Tier 2
- Closing the connection.
-

Thus, CC Storage is a powerful primary storage subsystem management tool that enable enterprises to reduce power and cooling costs significantly through increased storage utilization.

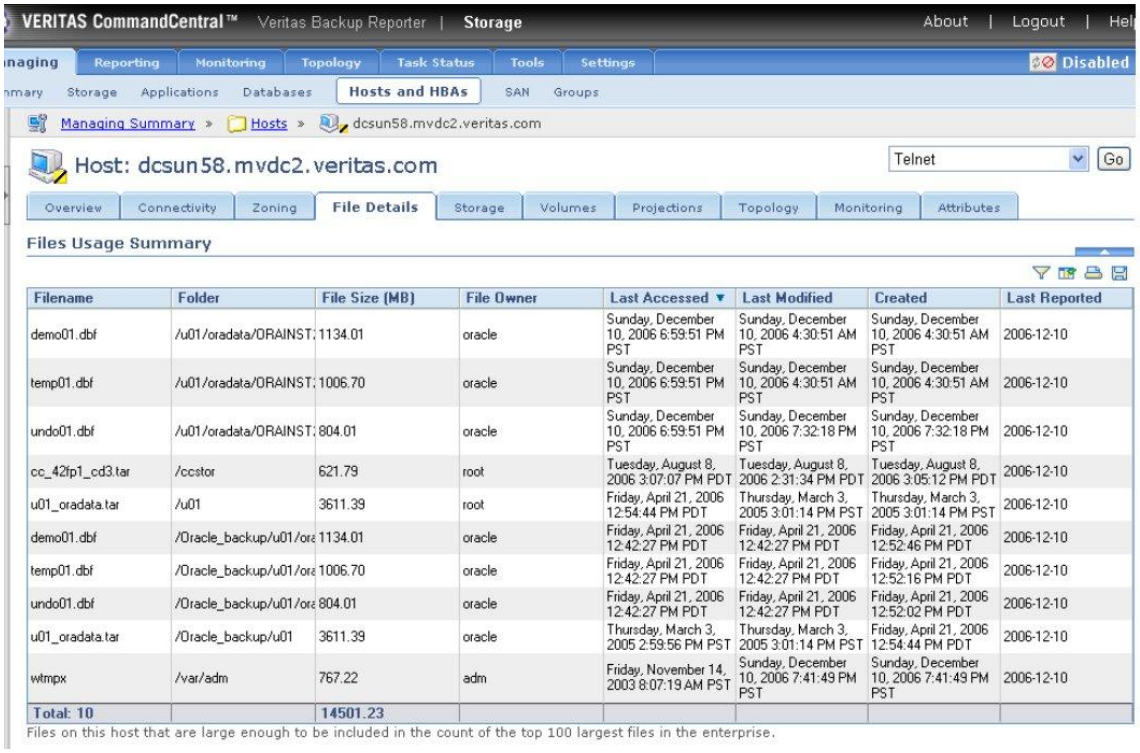

Filename	Folder	File Size (MB)	File Owner	Last Accessed ▼	Last Modified	Created	Last Reported
demo01.dbf	/u01/oradata/ORAINST:	1134.01	oracle	Sunday, December 10, 2006 6:59:51 PM PST	Sunday, December 10, 2006 4:30:51 AM PST	Sunday, December 10, 2006 4:30:51 AM PST	2006-12-10
temp01.dbf	/u01/oradata/ORAINST:	1006.70	oracle	Sunday, December 10, 2006 6:59:51 PM PST	Sunday, December 10, 2006 4:30:51 AM PST	Sunday, December 10, 2006 4:30:51 AM PST	2006-12-10
undo01.dbf	/u01/oradata/ORAINST:	804.01	oracle	Sunday, December 10, 2006 6:59:51 PM PST	Sunday, December 10, 2006 7:32:18 PM PST	Sunday, December 10, 2006 7:32:18 PM PST	2006-12-10
cc_42fp1_cd3.tar	/ccstor	621.79	root	Tuesday, August 8, 2006 3:07:07 PM PDT	Tuesday, August 8, 2006 2:31:34 PM PDT	Tuesday, August 8, 2006 3:05:12 PM PDT	2006-12-10
u01_oradata.tar	/u01	3611.39	root	Friday, April 21, 2006 12:54:44 PM PDT	Thursday, March 3, 2005 3:01:14 PM PST	Thursday, March 3, 2005 3:01:14 PM PST	2006-12-10
demo01.dbf	/Oracle_backup/u01/ora	1134.01	oracle	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:52:46 PM PDT	2006-12-10
temp01.dbf	/Oracle_backup/u01/ora	1006.70	oracle	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:52:16 PM PDT	2006-12-10
undo01.dbf	/Oracle_backup/u01/ora	804.01	oracle	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:42:27 PM PDT	Friday, April 21, 2006 12:52:02 PM PDT	2006-12-10
u01_oradata.tar	/Oracle_backup/u01	3611.39	oracle	Thursday, March 3, 2005 2:59:56 PM PST	Thursday, March 3, 2005 3:01:14 PM PST	Friday, April 21, 2006 12:54:44 PM PDT	2006-12-10
wtmpx	/var/adm	767.22	adm	Friday, November 14, 2003 8:07:19 AM PST	Sunday, December 10, 2006 7:41:49 PM PST	Sunday, December 10, 2006 7:41:49 PM PST	2006-12-10
Total: 10		14501.23					

Files on this host that are large enough to be included in the count of the top 100 largest files in the enterprise.

Figure C-1 – CC Storage File System Analysis Screen Shot

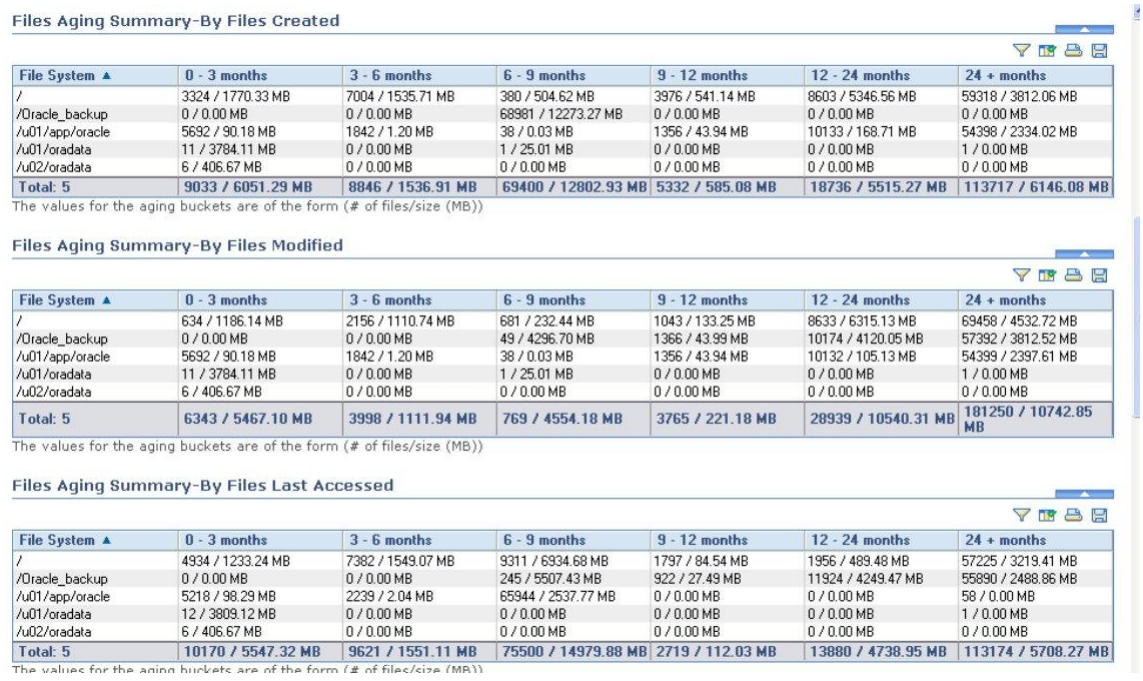

Files Aging Summary-By Files Created

File System ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
/	3324 / 1770.33 MB	7004 / 1535.71 MB	380 / 504.62 MB	3976 / 541.14 MB	8603 / 5346.56 MB	59318 / 3812.06 MB
/Oracle_backup	0 / 0.00 MB	0 / 0.00 MB	68981 / 12273.27 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB
/u01/app/oracle	5692 / 90.18 MB	1842 / 1.20 MB	38 / 0.03 MB	1356 / 43.94 MB	10133 / 168.71 MB	54398 / 2334.02 MB
/u01/oradata	11 / 3784.11 MB	0 / 0.00 MB	1 / 25.01 MB	0 / 0.00 MB	0 / 0.00 MB	1 / 0.00 MB
/u02/oradata	6 / 406.67 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB
Total: 5	9033 / 6051.29 MB	8846 / 1536.91 MB	69400 / 12802.93 MB	5332 / 585.08 MB	18736 / 5515.27 MB	113717 / 6146.08 MB

The values for the aging buckets are of the form (# of files/size (MB))

Files Aging Summary-By Files Modified

File System ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
/	634 / 1186.14 MB	2156 / 1110.74 MB	681 / 232.44 MB	1043 / 133.25 MB	8633 / 6315.13 MB	69458 / 4532.72 MB
/Oracle_backup	0 / 0.00 MB	0 / 0.00 MB	49 / 4296.70 MB	1366 / 43.99 MB	10174 / 4120.05 MB	57392 / 3812.52 MB
/u01/app/oracle	5692 / 90.18 MB	1842 / 1.20 MB	38 / 0.03 MB	1356 / 43.94 MB	10132 / 105.13 MB	54399 / 2397.61 MB
/u01/oradata	11 / 3784.11 MB	0 / 0.00 MB	1 / 25.01 MB	0 / 0.00 MB	0 / 0.00 MB	1 / 0.00 MB
/u02/oradata	6 / 406.67 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB
Total: 5	6343 / 5467.10 MB	3998 / 1111.94 MB	769 / 4554.18 MB	3765 / 221.18 MB	28939 / 10540.31 MB	181250 / 10742.85 MB

The values for the aging buckets are of the form (# of files/size (MB))

Files Aging Summary-By Files Last Accessed

File System ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
/	4934 / 1233.24 MB	7382 / 1549.07 MB	9311 / 6934.68 MB	1797 / 84.54 MB	1956 / 489.48 MB	57225 / 3219.41 MB
/Oracle_backup	0 / 0.00 MB	0 / 0.00 MB	245 / 5507.43 MB	922 / 27.49 MB	11924 / 4249.47 MB	55890 / 2488.86 MB
/u01/app/oracle	5218 / 98.29 MB	2239 / 2.04 MB	65944 / 2537.77 MB	0 / 0.00 MB	0 / 0.00 MB	58 / 0.00 MB
/u01/oradata	12 / 3809.12 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	1 / 0.00 MB
/u02/oradata	6 / 406.67 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB
Total: 5	10170 / 5547.32 MB	9621 / 1551.11 MB	75500 / 14979.88 MB	2719 / 112.03 MB	13880 / 4738.95 MB	113174 / 5708.27 MB

The values for the aging buckets are of the form (# of files/size (MB))

Figure C-2 – CC Storage File Aging Summary Report

File Types Usage Summary

File Type ▲	# of Files	Size (MB)	Average File Size (MB)	Last Reported
tar	22	8905.45	404.79	2006-12-10
dbf	27	7290.80	270.03	2006-12-10
	17361	3213.91	0.19	2006-12-10
sql	4410	2301.11	0.52	2006-12-10
so	4663	1998.61	0.43	2006-12-10
jar	1452	1307.99	0.90	2006-12-10
log	905	1036.77	1.15	2006-12-10
dfj	4	604.54	151.13	2006-12-10
a	449	588.14	1.31	2006-12-10
ora	33	400.16	12.13	2006-12-10
Total: 1377				

1 2 3 4 5 6 7 8 9 10

File Types Aging Summary-By File Types Created

File Type ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
	686 / 884.52 MB	552 / 130.63 MB	1118 / 525.23 MB	1718 / 114.83 MB	1220 / 315.40 MB	12067 / 1243.28 MB
0	6 / 13.53 MB	8 / 4.42 MB	7 / 55.64 MB	0 / 0.00 MB	7 / 65.42 MB	49 / 92.18 MB
001	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	6 / 0.02 MB
002	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
003	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
004	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.04 MB
005	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.01 MB
00503	0 / 0.00 MB	0 / 0.00 MB	1 / 1.17 MB	0 / 0.00 MB	0 / 0.00 MB	2 / 1.18 MB
006	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.03 MB
007	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.01 MB
Total: 1377						

1 2 3 4 5 6 7 8 9 10

The values for the aging buckets are of the form (# of files/size (MB))

Figure C-3 – CC Storage File Types Usage Summary Report

File Types Aging Summary-By File Types Modified

File Type ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
	113 / 788.26 MB	188 / 136.22 MB	111 / 67.41 MB	485 / 22.96 MB	1370 / 361.15 MB	15094 / 1837.90 MB
0	2 / 8.98 MB	8 / 1.90 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 31.01 MB	62 / 189.30 MB
001	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	6 / 0.02 MB
002	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
003	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
004	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.04 MB
005	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.01 MB
00503	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	3 / 2.35 MB
006	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.03 MB
007	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.01 MB
Total: 1377						

1 2 3 4 5 6 7 8 9 10

The values for the aging buckets are of the form (# of files/size (MB))

File Types Aging Summary-By File Types Last Accessed

File Type ▲	0 - 3 months	3 - 6 months	6 - 9 months	9 - 12 months	12 - 24 months	24 + months
	1680 / 291.18 MB	878 / 139.19 MB	3011 / 643.12 MB	709 / 22.41 MB	834 / 183.57 MB	10249 / 1934.42 MB
0	14 / 47.53 MB	4 / 4.42 MB	10 / 75.19 MB	1 / 0.17 MB	3 / 32.10 MB	45 / 71.78 MB
001	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	6 / 0.02 MB
002	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
003	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.02 MB
004	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.04 MB
005	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	5 / 0.01 MB
00503	0 / 0.00 MB	0 / 0.00 MB	1 / 1.17 MB	0 / 0.00 MB	0 / 0.00 MB	2 / 1.18 MB
006	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.03 MB
007	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	0 / 0.00 MB	4 / 0.01 MB
Total: 1377						

1 2 3 4 5 6 7 8 9 10

The values for the aging buckets are of the form (# of files/size (MB))

Figure C-4 - CC Storage File Types Aging Summary – By File Types Modified Report

About Symantec Technology Network (STN)

Symantec Technology Network (STN) is Symantec Corporation's technical information generation and dissemination organization. It distributes a free monthly technical newsletter that discusses timely technology events to 120,000 email subscribers across the globe. STN also publishes technical data storage and security articles each month for large enterprise and Small and Medium Business (SMB) readers, as well as hosts a variety of blogs and product discussion forums discussing Symantec product tips and insights. To subscribe to STN's Technical Newsletter and review other STN materials, please visit STN at:

http://www.symantec.com/enterprise/stn/index.jsp

About the Authors

Bruce Naegel

Bruce Naegel is a Senior Product Manager in Symantec's Data Center Management Group. He has been in various positions in Storage Product Management for over 20 years, spanning companies from 10 employees to 100,000. During that tenure, he participated in a number of award winning product releases. He has authored a number of published articles, delivered numerous mass storage presentations and been active in various Mass Storage Software standards bodies. He also has an active interest in kaleidoscopic photography.

John Colgrove

John Colgrove is a Symantec Fellow and Vice President.

W. David Schwaderer

W. David Schwaderer is the Symantec Technology Network (STN) Editor-In-Chief.

He has authored six commercial software programs and ten technical books including ***Data Lifecycles—Managing Data for Strategic Advantage*** published by John Wiley & Sons Ltd. His soon-to-be-published eleventh book is titled ***Innovation Survival—Concept, Courage and Change***. He is also working on a twelfth book titled ***Data Center of the Future*** that examines the tectonic forces reshaping today's enterprise data centers.

David has a Masters Degree in Applied Mathematics from the California Institute of Technology and an MBA from the University of Southern California. He lectures at Stanford on the subject of innovation.

About Symantec

Symantec is the world leader in providing solutions to help individuals and enterprises assure the security, availability, and integrity of their information. Headquartered in Cupertino, Calif., Symantec has operations in more than 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 800 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
408 517 8000
800 721 3934
www.symantec.com