

Be SURE about Storage Capacity Management

*Written by
Quest Software, Inc.*



White Paper

**© 2009 Quest Software, Inc.
ALL RIGHTS RESERVED.**

This document contains proprietary information, protected by copyright. No part of this document may be reproduced or transmitted for any purpose other than the reader's personal use without the written permission of Quest Software, Inc.

WARRANTY

The information contained in this document is subject to change without notice. Quest Software makes no warranty of any kind with respect to this information. QUEST SOFTWARE SPECIFICALLY DISCLAIMS THE IMPLIED WARRANTY OF THE MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Quest Software shall not be liable for any direct, indirect, incidental, consequential, or other damage alleged in connection with the furnishing or use of this information.

TRADEMARKS

Quest, Quest Software, and the Quest Software logo are trademarks and registered trademarks of Quest Software, Inc. in the United States of America and other countries. Other trademarks and registered trademarks used in this document are property of their respective owners.

World Headquarters
5 Polaris Way
Aliso Viejo, CA 92656
www.quest.com
e-mail: info@quest.com

Please refer to our Web site (www.quest.com) for regional and international office information.

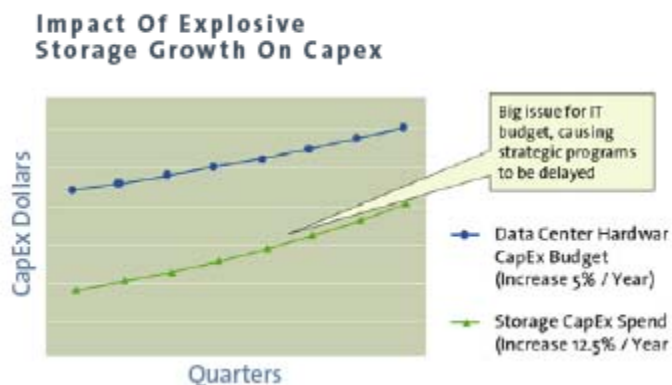
Updated—January 16, 2009

CONTENTS

- THE NEED FOR SYSTEMATIC STORAGE CAPACITY MANAGEMENT 1**
- STORAGE UTILIZATION INEFFICIENCIES 3**
- TODAY’S COMMON PRACTICES FOR STORAGE CAPACITY MANAGEMENT
ARE WASTEFUL 5**
 - ALLOCATION-BASED STORAGE CAPACITY MANAGEMENT 5
 - SERVER-BASED STORAGE CAPACITY MANAGEMENT 5
 - WHY CURRENT STORAGE CAPACITY MANAGEMENT PROCESSES FAIL..... 6
 - WHY TODAY’S SRM TOOLS FAIL AT CAPACITY MANAGEMENT 6
- CONSEQUENCES OF POOR PLANNING 8**
- SURE STORAGE CAPACITY MANAGEMENT SUCCEEDS 10**
- OPERATIONAL BENEFITS OF SURE STORAGE CAPACITY MANAGEMENT 12**
- FINANCIAL BENEFITS OF SURE STORAGE CAPACITY MANAGEMENT 13**
- ABOUT QUEST SOFTWARE, INC. 14**
 - CONTACTING QUEST SOFTWARE..... 14

THE NEED FOR SYSTEMATIC STORAGE CAPACITY MANAGEMENT

Most companies have a capacity management process in place to manage service levels and to control costs for their networks and servers, yet many overlook the fastest growing component of the IT infrastructure—storage. While storage may appear to be highly affordable, in reality the opposite is true. Even as storage unit prices continue to decline 25% annually, capacity continues to rise 50% or more, due to exponential growth of storage consumption by applications, compliance



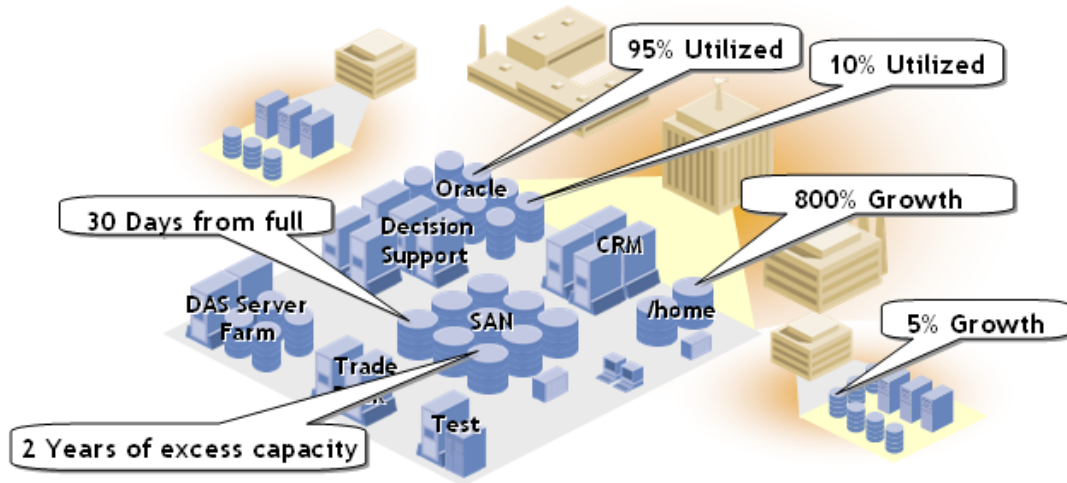
requirements, and data protection strategies. The result is alarming storage spending increases of 10% to 15% every year. In most businesses, storage costs are rising much faster than IT budgets and have become a financial burden that is impacting the rollout of major company initiatives¹.

Despite these rising costs, storage is generally significantly under-utilized. Based on many petabytes of production data analyzed by Quest, expensive storage assets experience a wasteful average utilization of approximately 30%, and most companies have 15-24 months of excess capacity on the data center floor.

Utilization is so unacceptably low because storage usage has been difficult to project due to the rapid creation and change rate of data, and the size and complexity of storage infrastructures. Out of the tens of thousands of individual volumes, file systems, files, and databases that consume storage infrastructure, some are highly utilized, while others have plenty of extra space. Some logical elements experience explosive growth, while others remain flat. Efficiently managing the sheer number and heterogeneity of these different storage elements with existing tools has proven to be impossible. This complexity has left administrators with no choice but to overprovision storage to avoid shortages.

¹ Quest 2006 survey of 140 storage professionals.

Be SURE about Storage Capacity Management

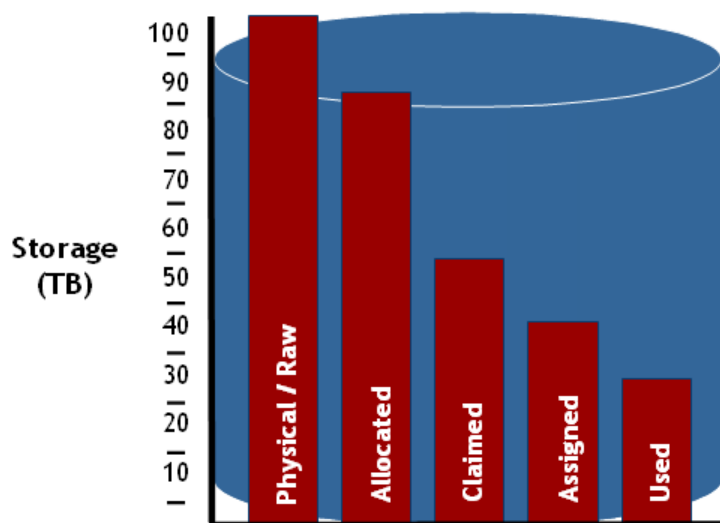


Systematic, Usage-based Regimen for Enterprises™ (SURE™) is a storage capacity management system that determines current and predicts future storage array usage (the amount of storage that contains application or end user data). SURE can significantly raise this utilization to 75% or more, drastically reducing storage capital expenditures and further reducing operational expenditures such as power, cooling, and data center floor space—without the jeopardy of an out-of-storage condition.

STORAGE UTILIZATION INEFFICIENCIES

In order to increase utilization, contributors to storage inefficiencies inherent in the provisioning process must be identified. Before exploring storage inefficiencies, let's establish common definitions.

- Raw Storage—the total of all physical storage installed on the array.
- Configured Storage—storage that has been formatted with a RAID protection level and assigned by the administrator to a LUN.
- Allocated Storage—configured storage that is mapped to a host server.
- Claimed Storage—allocated storage that the server has recognized and picked up.
- Assigned Storage—claimed storage that the server presents to users and applications as file systems and volumes.
- Used Storage—the amount of actual data that applications and/or users have written to configured storage.
- Utilization—the percentage of raw storage that has been used—used storage / raw storage, expressed as a percentage.



Sources: James Damoulakis, CTO, Glasshouse, Storage Magazine, July 2006
Quest storageanalysis of petabytes of customer data

Be SURE about Storage Capacity Management

Significant inefficiencies are introduced in each step of the process which contributes to poor utilization of expensive storage assets. In mid to large sized storage environments the following observations are typical:

- All raw storage is rarely allocated because some capacity is set aside for emergencies or unanticipated needs. The goal should be to configure and allocate 90%-95% of raw storage.
- All allocated storage is not claimed by hosts and all claimed storage is not assigned to file systems or volumes. These two inefficiencies make up “dark storage” and often contribute significantly to storage waste. This is usually caused by handoff issues between the storage administrator and the server/system administrator. The goal should be to eliminate all dark storage.
- Utilization of the file systems and volumes on application servers (assigned storage utilization) is typically 45%-60%. However, with a SURE storage capacity planning solution in place, 75%-85% utilization is typically achieved.

As previously mentioned, typically only about 30% of raw storage is used and most companies have 15-24 months of excess capacity on the data center floor. Utilizing this unused storage before purchasing new storage will significantly increase storage utilization, decrease storage capital spending, and increase storage administrator productivity.

TODAY'S COMMON PRACTICES FOR STORAGE CAPACITY MANAGEMENT ARE WASTEFUL

All enterprises, large or small, are dealing with capacity planning issues and are solving them in one way or another. As storage consumption continues to grow, its dynamic nature has storage administrators struggling to predict how much of what resource will be needed to support enterprise applications. Common practices used today are:

Allocation-Based Storage Capacity Management

The most popular storage capacity management process in place today is allocation-based storage capacity management. In this process, the storage administration team allocates storage to consumers based on how much they ask for—without knowing if the requests are too much, too little, or will meet storage consumption demands for weeks, months, quarters, or in fact, years into the future. This process assumes that allocated LUNs will be 100% claimed, assigned, and used (which has been shown to be highly inaccurate). When the un-allocated pool of storage starts to dwindle, additional storage is purchased to replenish the pool to fulfill future requests.

This practice is most commonly employed because today's Storage Resource Management (SRM) tools only provide storage array information that shows how much storage has been allocated and how much is kept in reserve. Since this is the only information many storage teams have access to, it has become the most prevalent storage capacity management process.

Server-Based Storage Capacity Management

Some storage teams have spent significant time, money, and sanity deploying and maintaining SRM agents on each of their application servers to collect information on the application server volumes in their environment. From this information, storage teams can determine how much data is written to each application server volume and then makes an assumption that the utilization of the storage systems is equal to the "utilization" of the application server volumes (which has also been shown to be highly inaccurate).

Why Current Storage Capacity Management Processes Fail

Neither process provides the fundamental data required for any successful capacity management process, which requires the actual usage and growth rates of all the storage consumers of the storage systems.

Basic assumptions built into these processes are inaccurate in real production environments:

- **Allocated storage is 100% used**—As shown in Figure 1, used storage is a small percentage of allocated storage. Analysis of real production environments conducted by Quest found that overall utilization (used / raw storage) is 30% and utilization of allocated storage (used / allocated storage) is only slightly better at 36%. With millions of dollars of storage hardware, it is extremely inaccurate and wasteful to assume that 100% allocated storage is 100% used.
- **Application server utilization is equal to storage system utilization**—As shown in Figure 1, assigned storage is much smaller than raw storage, so the application server “utilization” (used / assigned storage) is going to be much higher than real storage system utilization (used / raw storage). Quest has found that storage system utilization is typically less than ½ application server “utilization”.

Why Today’s SRM Tools Fail at Capacity Management

Today’s SRM tools discover storage resources, collect data, and provide various levels of reports. The capacity management component of SRM attempts to integrate the discovery and data collection process with a data repository and analysis. However, they fall far short of providing a complete capacity management solution. Since SRM primarily focuses on automating manually-recurring discovery and provisioning activities, capacity management is glaringly absent from these tools. SRM shortcomings include:

1. **Application server data collection requires agents**—High deployment and maintenance costs have greatly hindered the use and effectiveness of agent-based IT management solutions. Countless real world horror stories abound.
2. **Do not determine array usage**—Capacity management on any system simply cannot be completed without knowing the actual usage of the system.

3. **Usage forecasts provide no value**—With a painstaking manual process, some SRM provide rudimentary usage forecasts for host volumes only if and when the user asks for it. Rudimentary means that it just provides a straight line “least sum of the squares” algorithm. Not only are these forecasts not accurate, but more importantly, they are not saved in a database so that they can further analyzed. Examples of basic analysis that is far beyond the capabilities of SRM software, even if operated by expert administrators, are:
 - Identify volumes and LUNs which are forecasted to run out of storage (100% utilization) in the next 90 days?
 - Identify volumes and LUNs which are forecasted to stay below a low threshold in the next 180 days and have predicted free storage greater than, say, 200 GB? (These volumes and LUNs would be good candidates for storage reclamation.) SRM tools provide no usage forecasts for storage systems at all.
4. **Do not provide automated storage capacity plans**—A primary benefit of storage capacity management is to determine how much capacity to purchase, when to purchase it, and where to deploy it. With SRM, consumption and supply data needs to be collected, correlated, and dumped into a spreadsheet and extensive calculations must be manually performed by programming experts.
5. **Generating customized reports is difficult**—Although SRM tools come with pre-packaged reports, most IT organizations require customized reports to meet their individual needs. For most SRM tools, creating customized reports is a long and difficult process and typically requires extensive knowledge of SQL or scripting.
6. **User defined groups are not supported**—Administrators require the capability to establish custom groups which allow them to organize storage elements associated with users, business units, applications, storage tiers etc. into logical views. Most SRM tools do not provide this capability.

While they may help identify current problem areas that need attention, SRM tools do a very poor job of supporting future planning and forecasting. In addition, most SRM tools require that thousands of software agents be deployed and maintained, which translates to deployment schedules measured in quarters and in many organizations, years.

CONSEQUENCES OF POOR PLANNING

Many enterprises underestimate the consequences of poor planning. Poor planning often leads to the following issues:

- **Emergency Fire Drills**—Emergency out-of-capacity conditions are typically extremely visible to end users, are disruptive to normal business flow, can negatively impact overall company profitability, and can damage the company's and the storage team's reputations. Alarms resulting from thresholds that have been exceeded are often too late to avoid emergency out-of-hours calls to provision more storage.
- **Overspending and Underutilization**—Avoiding the dreaded fire drill usually leads to over-spending and underutilization to ensure there is always plenty of excess storage capacity on hand. The more risk averse the capacity planner is, the more excess capacity is made available to safeguard against capacity shortages. While it may be completely reasonable to have 3-6 months worth of excess capacity on hand, having 1 to 2 years—or even 9 months—usually does not make sound business sense. By the time some of this capacity is ready to be used, it will already be obsolete. While it's prohibitively risky to run at 95% utilization or just 1 month of excess capacity, 75—85% utilization and 3-6 months of excess capacity is perfectly achievable through the use of the SURE storage capacity management process.
- **Purchasing Storage Too Far in Advance**—Purchasing storage too far in advance of when it is actually needed adversely impacts the organization in several ways. First and foremost, it fails to take advantage of falling storage prices. With storage prices declining by an average of almost 2% per month (about 25% per year), delaying storage purchases by 9 months leads to an automatic built-in discount of about 17%. Additionally, storage hardware maintenance contracts generally don't discriminate between used and excess capacity—the result is that customers pay maintenance on all the physical storage on-site. If the excess capacity offered at a special 20% discount to buy now is not needed for another year, it is not a good deal.

Another negative impact of forward buying is failing to acquire the latest technology at the right time. While the administrator may be buying the latest technology now, it is going to be technologically dated by the time it gets used a year later. In addition, you may well be eating into your useful reliability lifetime by having disks in place, powered up, and spinning. With storage technology having an average shelf life of three years before having to be upgraded (thanks in part to high maintenance renewal rates after three years), purchase timing is a critical factor. If the storage resource does not get used until a year or more after it is purchased, its associated benefit gets trimmed from three years to two. Factoring in the cost of capital and the opportunity cost of that same capital over the year deferral, along with falling storage prices, suddenly even a 50-60% special discount to buy more now may not be such a great deal.

- **Excessive Power, Cooling and Data Center Real Estate Costs**—Having more hardware than you need also wastes expensive energy for power and cooling and takes up precious data center real estate. Studies have shown that over a 3-year life cycle, these energy costs are typically around 12% of the original equipment costs².

² Enterprise Strategy Group (ESG), 2006.

SURE STORAGE CAPACITY MANAGEMENT SUCCEEDS

Dramatically improving utilization and cutting storage CapEx requires a comparatively small investment in a SURE (Systematic, Usage-based Regimen for Enterprises) storage capacity management system. What makes SURE storage capacity management best in class is its ability to determine current and predict future storage array usage (the amount of storage that contains application or end user data) through its agent-less server, application and storage system data collection, association, and predictive analytics processes. These processes are critical to effectively managing storage capacity and provide compelling benefits which include:

- Automatically developing optimized storage capacity plans allowing application data to grow into existing capacity without purchasing additional storage.
- Clearly detailing the storage environment, current usage, usage growth rates, and forecasted usage.
- Identifying dark storage that can be reclaimed for applications in need of additional capacity.
- Pinpointing over-provisioned areas so wasted storage can be reclaimed.
- Providing early warning of potential over-utilization issues before they become out-of-storage emergencies.

As a result, SURE storage capacity management system can decrease storage capital expenditures by as much as 50%, and drastically reduce operational expenditures such as power, cooling and data center floor space.

SURE storage capacity management consists of the following integrated steps:

- **Collect** storage related meta information from storage subsystems (storage inventory) and application hosts (storage usage). Best in class processes do this without the use of installing software agents on application servers, understand data protection schemes, and account for clusters and virtualization.
- **Associate** host volumes with storage subsystem logical devices, including consideration for RAID protection levels, snapshots, and replication, at both host and storage subsystems. This enables the determination of true usage of each logical device (LUN) and the overall array.
- **Baseline** current utilization and develop future utilization targets.
- **Forecast** future usage and utilization for each volume, database, file system, and logical device after a number of collections have been completed.

- **Analyze** current and projected storage consumption and supply details.
- **Plan** for future provisioning and purchases by developing a storage capacity plan to optimize what to buy, when to buy it, and where to deploy it.
- **Provision** storage per the plan, using hardware vendor or user generated tools.
- **Enterprise-Wide Monitoring** of storage capacity on a daily basis ensures that changes in storage consumption don't cause emergencies.



OPERATIONAL BENEFITS OF SURE STORAGE CAPACITY MANAGEMENT

A SURE storage capacity management system enables the storage team to accomplish their tasks much more efficiently, and scale to support much larger environments. Benefits include:

- Improved administrator productivity by eliminating unnecessary provisioning tasks.
- Clear, accurate, and objective cost justification for future storage purchases.
- Excess capacity requires excess administration—reduced storage capacity requires fewer administrators to manage it.
- Fewer unnecessary calls from end users for storage that they do not need.
- Proactive planning of provisioning tasks, enabling better administer work load balancing.
- Fewer emergency, out-of-hours provisioning tasks to solve out-of-storage problems.
- Fewer re-provisioning tasks due to inaccurate sizing estimates.

FINANCIAL BENEFITS OF SURE STORAGE CAPACITY MANAGEMENT

SURE storage capacity management takes dead aim at poor storage utilization and resulting over-spending. The table below shows a straightforward example of how economically compelling a SURE storage capacity management solution can be.

Table 1. Storage and Power Requirements for 1 PB of Data

STORAGE UTILIZATION	REQUIRED STORAGE	STORAGE CAPEX ³	POWER & COOLING COSTS / 3 YEARS ⁴
40%	2.5 PB	\$50M	\$6M
80%	1.25 PB	\$25M	\$3M
Savings	1.25 PB	\$25M	\$3M

Let's say a company has 1 PB of data. If they are at 40% utilization (slightly above industry average, but simplifies the math), they would have purchased 2.5 PB of storage ($1 \text{ PB} / 40\% = 2.5 \text{ PB}$ or $1 \text{ PB} / 2.5 \text{ PB} = 40\%$). At \$20,000 per TB, 2.5 PB costs approximately \$50M. With energy costs for power and cooling at about 12% of purchase price over a 3 year period, that costs \$6M over three years.

Using a SURE storage capacity management system, the company increases their utilization to 80%. They would have only purchased 1.25 PB of storage ($1 \text{ PB} / 80\% = 1.25 \text{ PB}$ or $1 \text{ PB} / 1.25 = 80\%$). Now, storage CapEx and energy costs are halved, saving \$25M in CapEx and \$3M in energy over 3 years.

In most companies, the financial benefits of implementing a SURE storage capacity management solution will far exceed the economic benefit of all other cost containment projects in it.

³ Assuming storage costs of \$20,000 per TB.

⁴ Enterprise Strategy Group, 2007.

ABOUT QUEST SOFTWARE, INC.

Quest Software, Inc., a leading enterprise systems management vendor, delivers innovative products that help organizations get more performance and productivity from their applications, databases, Windows infrastructure and virtual environments. Quest also provides customers with client management through its ScriptLogic subsidiary and server virtualization management through its Vizioncore subsidiary. Through a deep expertise in IT operations and a continued focus on what works best, Quest helps more than 100,000 customers worldwide meet higher expectations for enterprise IT. Visit www.quest.com for more information.

Contacting Quest Software

Phone: 949.754.8000 (United States and Canada)
Email: info@quest.com
Mail: Quest Software, Inc.
World Headquarters
5 Polaris Way
Aliso Viejo, CA 92656
USA
Web site: www.quest.com

Please refer to our Web site for regional and international office information.