

Symantec Technology Network

Using Software to Reduce Hardware Requirements

- by -

Bruce Naegel

Symantec Sr. Product Manager, Data Protection Group,

April 15, 2008



Table of Contents

| | |
|--|---|
| Table of Contents | 2 |
| Overview | 3 |
| Clustering | 3 |
| Storage Tiering | 4 |
| Data Deduplication | 4 |
| Virtualization | 4 |
| Storage Resource Management | 5 |
| About the Author | 7 |
| About Symantec Technology Network (STN) | 7 |

Overview

Maximizing the value of existing hardware investments has never been a bad idea. Perhaps never before has it also been such a good idea.

Today's IT organizations are up against a wall. It's not just that data centers are running out of space, or that power and energy costs are escalating, or that corporate social responsibility initiatives demand greater efficiencies throughout IT. It's also that few IT organizations can afford to simply rip and replace their data center infrastructures to address space, environmental, corporate responsibility, or any other issue. After all, if the capital costs of doing so were not enough to prohibit such dramatic action, then the potential risk to business continuity likely would be.

As a result, a growing number of IT organizations are looking at a more affordable yet effective solution for reducing their hardware requirements: software. As it turns out, a handful of technologies that IT administrators have relied on for years in their server and storage environments are also proving to be compelling tools for getting more out their hardware:

- Clustering
- Storage tiering
- Data deduplication
- Virtualization
- Storage resource management

These are components of a cost-effective strategy for maximizing IT investments while also business continuity.

Clustering

Reducing application downtime is typically the objective behind the deployment of clustering technologies. But application protection is only one benefit, and not all clustering tools offer the same capabilities. Most clustering solutions require identical hardware and operating system environments; yet, this can be an expensive and restrictive solution since organizations must also ensure platform parity in a cluster.

However, clustering solutions are now available that can support a variety of operating systems—from UNIX to Windows, Linux, and virtual platforms—as well as a wide range of heterogeneous hardware configurations. With these tools, IT can maximize resource utilization by consolidating workloads running from underutilized hardware onto a smaller number of machines.

In addition, while many clustering products and deployments are pairs of active-passive nodes in a high availability cluster, solutions have emerged that support N+1 clusters. When paired with workload management, resilience, and advanced failover logic, using these solutions in an N+1 configuration preserves availability while reducing power draw, increasing utilization, and reducing the servers required to support the applications.

Storage Tiering

The key to deriving value from multiple tiers of storage is to place each file on the appropriate type of storage device based on a number of criteria, including its availability and performance requirements as well as protection and other business needs. Today, enterprises can exploit different storage components, configurations, and capabilities to create tiers of storage that match cost with data value and I/O performance with data access needs, while meeting business application requirements for availability and security.

An ancillary benefit of storage tiering is that many organizations discover that clearly matching data with appropriate storage also reduces the amount of required physical storage hardware. Typically, a relatively small portion of an organization's files are considered critical; these files can be placed on lower capacity, higher performing devices. The bulk of the organization's data, which is considered non-mission critical, can then be placed on higher capacity, lower performing devices. Since a 1TB lower performance drive takes up the same amount of physical space as a high performance 70GB drive, fewer devices overall are required for storing data.

Furthermore, organizations with dynamic storage tiering capabilities enjoy additional efficiencies. Dynamic storage tiering enables IT to utilize the appropriate tier of storage for their applications without impacting the availability of an application to the business.

Information is non-disruptively moved to less expensive storage based on the changing value of the information. IT defines the policies that move the files, and the policies shift the files dynamically to higher or lower tiers of storage based on the defined criteria, eliminating administrative overhead.

Data Deduplication

Perhaps one of the most overlooked strategies for reducing hardware requirements is to decrease the overhead associated with holding multiple data copies by identifying common data and reducing copies to a single entity. Often referred to as data deduplication or single-instance storage, this technology can have a dramatic impact—in some cases, reducing data by 98 percent—on the amount of disk storage required for archiving purposes as well as the number of tapes and tape drives required for backup purposes.

Deduplication is commonly used in messaging applications. For example, if an attachment was sent to 100 people, it would be stored only once, thus saving the disk space that would have been required to store the other 99 copies.

Virtualization

Server virtualization allows for better processing load distribution across server hardware, thereby optimizing this hardware. Having the right software infrastructure for server virtualization enables virtualization to be extended into enterprise level applications. In acknowledgement of the benefits that can be realized through server consolidation, IT managers are also exploring storage virtualization.

Storage virtualization helps improve the hardware utilization challenge by enabling administrators to pool all storage into logical groups that can be reallocated quickly or in real-time based on demand. The best virtualization software can do this across any storage array from a variety of vendors, running under a variety of operating systems, and from a single management interface.

When storage resources are virtualized, they appear as a single resource. Data can be moved transparently across vendors and operating systems to utilize available capacity. Storage management tools also enable IT to classify data by age or type so that less valuable or less current data can be moved automatically to higher capacity, less costly storage.

Storage Resource Management

The majority of storage devices contain a large percentage of data that is of little or no immediate business value. But simply reviewing overall disk capacity usage does not give IT administrators this information. Indeed, a disk array may very well be full, making it appear as though storage is at 70 or 80 percent of capacity, when in reality the majority of files on the disk are old, rarely accessed, or non-business-related.

The challenge, then, is to better manage where and how critical and non-critical information is stored. This requires visibility into both current and historical disk usage. IT must be able to accurately track the availability and usage of resources and make the necessary adjustments based on such information to maximize their storage assets.

Storage resource management tools with discovery, visualization, and reporting capabilities enable IT to reclaim lost or wasted storage and more accurately forecast current and future storage capacity.

By using familiar technologies that have proven effective in their server and storage environments, IT can also realize fundamental cost efficiencies through improved resource utilization. Through workload consolidation, storage tiering, the elimination of duplicate data, server and storage resource virtualization, and the reclamation of lost or wasted storage, IT organizations can improve existing data center efficiency and maximize their current hardware infrastructure.

About the Author

Bruce Naegel is a Senior Product Manager in Symantec's Data Protection Group located in Mountain View, CA. He is responsible for Symantec's Shared Infrastructure Group Security (SIGSEC) software infrastructure and has over 20 years of storage product management experience. Bruce has published articles on a variety of subjects including performance optimization and green computing. He has led industry standards committees and presented at various industry forums including Storage Networking Industry Association (SNIA), AFCOM® and Computer Measurement Group (CMG). He has a BS in Biomedical Engineering from Rensselaer Polytechnic Institute in Troy, NY.

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>

W. David Schwaderer is STN's Editor-In-Chief. He has authored ten technical books, and six commercial software programs. His eleventh book titled *Innovation Survival* is nearing completion. He 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 and in Silicon Valley on the subject of innovation.

Russ G. Gregg is STN's Video Content Producer. He is presently focusing on producing state-of-the-art digital video for the Web and small form-factor, portable video devices such as Apple Inc.'s iPhone™ and Video iPod®. Russ has a Bachelor of Arts in English Literature, Texas from A&M University.



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.

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

Symantec and the Symantec logo are U.S. registered trademarks of Symantec Corporation. Other brands and products are trademarks of their respective holder/s. Any technical information that is made available by Symantec Corporation is the copyrighted work of Symantec Corporation and is owned by Symantec Corporation. NO WARRANTY. The technical information is being delivered to you as-is and Symantec Corporation makes no warranty as to its accuracy or use. Any use of the technical documentation or the information contained herein is at the risk of the user.
Copyright © 2007 Symantec Corporation. All rights reserved.