

Realtime  
publishers

The Essentials Series: Mid-Market  
Storage Management and Virtualization

# Optimizing Storage Performance

*sponsored by*



Enterprise Computing Solutions  
MIDMARKET GROUP

*Grow. With the leader.*

by Dan Sullivan

---

Optimizing Storage Performance.....	1
Improving I/O Performance.....	1
Improving I/O Performance Through Load Balancing.....	1
Improving I/O Performance with Caching.....	2
Improving Storage Management Operations.....	3
Improving Business Reporting.....	4
Improving Application-Specific Storage Management.....	5
Summary.....	5

---

## Copyright Statement

© 2008 Realtime Publishers, Inc. All rights reserved. This site contains materials that have been created, developed, or commissioned by, and published with the permission of, Realtime Publishers, Inc. (the “Materials”) and this site and any such Materials are protected by international copyright and trademark laws.

THE MATERIALS ARE PROVIDED “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE AND NON-INFRINGEMENT. The Materials are subject to change without notice and do not represent a commitment on the part of Realtime Publishers, Inc or its web site sponsors. In no event shall Realtime Publishers, Inc. or its web site sponsors be held liable for technical or editorial errors or omissions contained in the Materials, including without limitation, for any direct, indirect, incidental, special, exemplary or consequential damages whatsoever resulting from the use of any information contained in the Materials.

The Materials (including but not limited to the text, images, audio, and/or video) may not be copied, reproduced, republished, uploaded, posted, transmitted, or distributed in any way, in whole or in part, except that one copy may be downloaded for your personal, non-commercial use on a single computer. In connection with such use, you may not modify or obscure any copyright or other proprietary notice.

The Materials may contain trademarks, services marks and logos that are the property of third parties. You are not permitted to use these trademarks, services marks or logos without prior written consent of such third parties.

Realtime Publishers and the Realtime Publishers logo are registered in the US Patent & Trademark Office. All other product or service names are the property of their respective owners.

If you have any questions about these terms, or if you would like information about licensing materials from Realtime Publishers, please contact us via e-mail at [info@realtimepublishers.com](mailto:info@realtimepublishers.com).

---

## Optimizing Storage Performance

Businesses can optimize storage management in several ways, all of which are enabled or improved by storage virtualization. Some of these areas focus on technical implementation and configuration issues and others are more business-process oriented. This article will examine how businesses can optimize their storage operations by

- Improving I/O performance
- Improving storage system manageability
- Improving reporting and chargeback services
- Addressing specialized needs of applications

The most efficient implementations leverage the benefits of all these focus areas.

### Improving I/O Performance

The first step to improving I/O performance is to utilize virtualized storage management, which can improve I/O performance in two ways. I/O performance can be improved through load balancing and with caching.


#### *Improving I/O Performance Through Load Balancing*

In a virtualized environment, storage managers will be able to allocate storage across disk arrays to reduce bottlenecks even without changes to storage hardware. Consider, for example, the following scenario. Suppose a business currently owns three disk arrays: one for an enterprise resource planning (ERP) system, one for network shared drives and a variety of application servers, and the third for a data warehouse. The ERP system is relatively mature and has fairly constant storage requirements. Utilization in the storage array for network shared drives and application servers vary widely over time. The demands from end users for file storage fluctuate in relatively unpredictable patterns. Software developers and systems administrators frequently create and decommission virtual servers and related storage. The data warehouse storage array is underutilized because that project is new and will not reach storage capacity for at least 2 years.

Without virtualization, the load on the data warehouse server is much less than on the other servers. The ERP system and application servers, however, are performing nearly constant I/O operations on the other storage arrays. In a virtualized environment, a storage manager could distribute the ERP, file storage, application server, and data warehouse storage across multiple storage arrays evening out the load. Without adding any storage devices or upgrading controllers, the I/O performance over all applications is improved.

## Improving I/O Performance with Caching

A second way to improve performance is through caching. The virtual storage manager can cache data as it is read or updated and make it available to other read operations. Retrieving from cache can be much faster than retrieving from disk (see Figure 1).



**Figure 1: Caching allows data to be stored and retrieved from high-speed memory faster than repeated fetches from disk.**

In addition to improving performance, storage virtualization improves availability of data. For example, a database may be replicated and, in the case of failure, the replicated copy can be used by updating the logical to physical mapping information in the virtualized storage manager. Changes to the configuration of a virtual storage array can also change with minimal disruption; failed disks can be removed and new hardware added with minimal impact on applications. The virtualized storage manager can implement high-level services, such as replication and snapshot management, so all storage devices in the system can benefit from these services in a consistent manner. Another form of optimization comes from consolidating management operations.

---

## Improving Storage Management Operations

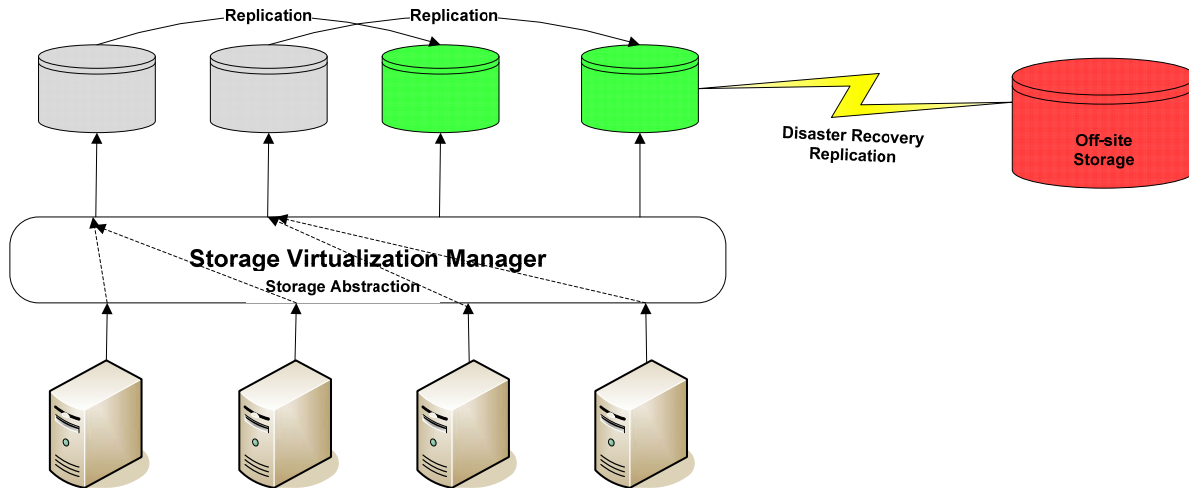
Virtualized storage management offers several ways to optimize management operations. The first is *thin provisioning*. This is the ability to quickly configure a logical unit of storage and assign it to a server without actually taking an equivalent amount of physical storage out of storage pool. Instead, the application that needs that space gets the physical space allocated as needed.

In the old, non-virtualized model of storage allocation, if a database, for instance, needed up to 1 terabyte of storage, then 1 terabyte was allocated and no other application could use that storage. This is known as *fat provisioning*. A drawback of this approach is that much of the allocated space may be unused for much of the time. With thin provisioning, multiple applications can be assigned large amounts of storage that is allocated as needed. Applications are unlikely to all need the maximum amount of storage at the same time, so it is safe to “promise” more logical storage than there is physical storage. This is somewhat analogous to a banking system in which the bank keeps less cash on hand than is actually in customer accounts; as long as there is not a run on the bank, there is no problem providing cash to customers when it is need. The same model applies with thin storage allocation. The advantage of thin provisioning is that it reduces the wasted resources that occur when you allocate physical storage based on peak demand for all applications.

User management can be improved with virtualized storage management as well. Access controls and security policies can be standardized across storage platforms and enforced from a single control point.

I/O redirection enables more manageable disaster recovery, data replication, and data migration. For example, data replication is often used in high-availability systems to ensure a current copy of data is available should the application’s primary storage system fail; the I/O redirection metadata in the storage virtualization manager needs to be updated only to direct reads and writes to the replicated version. That replicated data is updated as the primary source changes so that it is always current.

This setup can also help with disaster recovery. The replicated data can be copied to an offsite storage system, in the case of catastrophic failure in the primary site, but without putting additional load on the primary data source (see Figure 2). An area that extends from storage management operations into business operations is business reporting on storage services.



**Figure 2: High-availability and disaster recovery services can be managed for all applications through the storage virtualization manager.**

## Improving Business Reporting

It may seem like a distraction for IT staff focused on technical issues, but business reporting is an essential part of a well-run business. When it comes to chargebacks to end user departments, it is important to have accurate and timely information. A virtual storage management solution can help in this area with at least two types of reporting.

With regards to license management, storage management systems can report on the size of data under management by particular applications and the number of users who have access to that data. This information can help ensure that departments are charged properly and that the business does not over-subscribe to use-based licensing.

Another advantage is that a comprehensive and unified storage reporting system can provide fine-grained chargeback reporting. A single department, for example, might have multiple projects or initiatives under way, and each might be charged to a different line item. For these customers, it is important to have detailed reporting by project rather than having all charges lumped together for the entire department. In addition to these generalized optimizations in IT operations and business reporting, there may be application-specific optimizations enabled by storage virtualization.

---

## Improving Application-Specific Storage Management

Enterprise applications, such as relational databases and email systems, have idiosyncratic storage requirements that can be efficiently addressed by virtualized storage management. In the case of database applications, there are a few types of storage required. Data files are used to store tables and indexes that comprise the bulk of database storage. This is a basic form of long-term storage. In addition, databases require temporary storage for ad hoc operations, such as sorting; they also depend upon archive information (sometimes referred to as “redo logs”) to recover in case of a database failure. Typically, a database will require a moderate amount of temporary storage which, ideally, provides the highest performance possible because interactive queries often make use of this type of storage. Long-term data stores must provide acceptable performance and reasonable cost because large volumes of data are stored there. Archive logs should provide acceptable write performance, but read performance does not have to be as fast as that used for sorting and other operations tied to interactive response time. Given these requirements, a tiered-storage model could be used to balance the need for high-performance, reliability and cost effectiveness.

Email can quickly become a major consumer of storage services, so managing effectively is often a top priority. With a virtualized storage environment, systems administrators may be able to configure email archives so that attachments and older emails are moved to lower-cost and possibly lower-performance storage. Other enterprise-scale applications and even departmental-level applications may have particular requirements that can be met more effectively and in a more cost-effective manner with a tiered-storage model.

### Summary

Inefficiencies can easily creep into a storage management program. I/O performance can suffer because of poor distribution of data, the need to allocate storage for peak demand, and the lack of high-performance caching. Staff time can be unnecessarily taken up with manual storage allocation procedures and cumbersome reporting tasks, both of which could be automated in a virtualized storage management environment. Specialized applications—such as databases, email systems, and ERP systems—are large enough and complex enough to warrant specialized storage management plans. Fortunately, these can be accommodated in a virtual storage environment. The collective benefits of storage virtualization clearly present a number of opportunities to optimize storage performance.