



The Storage Application: A Fourth Generation Storage Device

A StoreAge White Paper

The Virtualization Architects

The Storage Application: A Fourth Generation Storage Device

A StoreAge White Paper

Copyright © StoreAge Networking Technologies Ltd. 2001.

All rights reserved.

No part of this publication may be reproduced in any material form (including photocopying, or storing it in any medium by electronic means whether or not transiently or incidentally to some other use of this publication) without the prior written permission of the copyright owner, or under the terms of a license issued by the copyright owner.

The information contained in this document is subject to change without notice. StoreAge is neither responsible for nor liable to anyone in connection with this document.

StoreAge, SVM and MultiView are trademarks of StoreAge Networking Technologies Ltd. All other marks mentioned in this material are the property of their respective owners.

Contact Information:

StoreAge Networking Technologies Ltd.
Gutwirth Science Park, Technion City
Haifa 32000, ISRAEL
Tel: +972-4-832-9030 Fax: +972-4-832-9040
<http://www.Store-Age.com>

Solution Centre Limited
Vickers House
Priestley Road
Basingstoke, RG24 9NP.
Tel: 01256 818600
Fax: 01256 819600
E-Mail: sales@solutioncentre.co.uk
www.solutioncentre.co.uk

<http://www.store-age.com>

Introduction

Storage devices have developed over basically three generations:

In the **first** generation, file systems communicated directly with individual physical devices (such as floppy or hard disks).

In the **second** generation, arrays were created and the file systems or applications communicated with RAID controllers instead of individual disks.

In the **third** generation, a virtual abstraction of RAID was introduced and the file systems communicated with Virtual Volumes instead of directly with RAIDs.

This paper suggests that a **fourth** generation is emerging, whereby file systems will communicate with Storage Applications that will react as storage devices. Fourth generation storage devices are software elements rather than physical or even virtual devices. They are not dependent on the specific file structure of the data. Fourth generation storage devices enable the Operating Systems and the File Systems to communicate with storage in a software-to-software manner rather than software-to-device manner, with all the advantages that such a mode of communication provides.

Why a Fourth Generation Storage Device

The need to properly manage the enterprise storage resources is becoming more and more apparent to all organizations. Users are already willing to pay significant premium for what they perceive as a manageable storage system.

Vendors of various hardware components in a Storage Area Network (HBAs, switches, RAID controllers) know they must provide ways to manage and control them. Until now, there have been few standards for this and the management tools and methods have been specific to each device and vendor. This includes tasks such as creating LUNs and their RAID characteristics and sizes, defining switch zoning, specifying LUN access control, etc.

The addition of Virtualization does not completely eliminate the need to configure all of these components, but provides a layer of management that is uniform and independent of various vendors. Thus LUNs created on various storage devices can be assigned to Storage Pools and allocated as needed to server and applications. Beyond basic volume management, the general capability to create point-in-time snapshot images of virtual volumes using “copy-on-write” technology is also essential

This paper suggests that the storage virtualization component is a key platform for creating high-value storage components that can be used to increase the return on the storage area network investment. In short, the value of these fourth generation devices is greater than the sum of the parts that make it up.

The StoreAge Solution

StoreAge Networking Technologies Ltd. developed the SVM (Storage Virtualization Manager) that provides high-performance virtual volume management via its out-of-the-data path architecture. SVM also creates a new platform on which Storage Applications can be executed. Storage Applications are software modules that can take Volumes, manipulate them and produce new Volumes. The result is a fourth generation storage device, whereby the OS and file system write or read data to a storage application instead of to a traditional device.

Examples of such storage applications are:

Volume Replication

Volume Remote Mirror

Volume Snapshot

Every Storage Application is an independent software product that can be added as needed. Each has a Web based GUI, for configuration and setting of parameters, or well defined Application Programming Interface (API) to be used from 3rd party applications. In the StoreAge solution, storage applications are run in the SVM Agent and managed by SVM Appliance. StoreAge provides SAN API and Command Line Interface (CLI)s for implementation of the applications on various servers in the SAN. By viewing a storage application as a device, this application becomes transparent to and independent of the operating system. This method enables the creation of more intelligent devices. In fact, storage applications can be stacked, that is, a storage application can write/read data to/from a device that by itself is a storage application, which produces a totally new kind of storage devices with exciting possibilities.

Example:

You can create a “Snapshot” of a specific volume, and begin to read and write it as a virtual volume independent of the volume on which it was based. Later, this virtual volume can itself have a snapshot taken. This process can go on indefinitely, with virtual volumes created over themselves. In this way you have only one “real” volume that contains the entire original set of data, and a number of snapshot volumes, based on the original one, where each snapshot volume requiring space to hold changes made since the last snapshot was created.

An example of how this can be used:

Assume we have volume “X:”

After performing a “snapshot” operation, we can create virtual volume “Y”.

This volume can be assigned to a group testing a new version of a software package, in order to verify its operation against “live” data

If this testing group wants to create 2 copies of the volume in order to work in two parallel directions, another snap shot of the virtual volume can be created.

Another example:

Let's assume that two different storage applications coexist within the same environment:

- ?? Remote Mirroring, in which a volume's contents are maintained at two different geographic locations.
- ?? Snapshot, described above, where a new volume is created whose initial contents are the same as the volume upon which it was based, but then has its own lifecycle.

Stacking these applications in a different order will produce different results! In one case a Volume can be first Remote Mirrored, and then a Snapshot can be created from the Remote Mirrored Volume.

Creating a Snapshot of a volume and then creating a Remote Mirror of the Snapshot Volume will obtain a different result.

This phenomenon can produce numerous combinations, giving the user the capability to almost "customize" the storage device according to specific needs.

Because the SVM Architecture is "out of the data path" with the virtualization agents distributed on the various servers, overall system performance for the "stacked" applications of virtualization technology is much greater than if all operations had to go through a central, "in the data path" virtualization appliance.

Conclusion

Fourth generation storage devices, which embed some of the important data management functions, will greatly improve overall storage system manageability at significant TCO savings to enterprises. StoreAge's Storage Applications are actual implementation of such devices within an actual SAN environment.