

## Market Perspective

MARCH 21, 2005

### Storage Virtualization—Why Architecture Matters

By Charles King

Like the weather, storage virtualization is something everybody talks about, but unlike the weather, many vendors are actually trying to do something about it. A simple reading of the news suggests that this has been the case since 2003, but the topic is likely to heat up considerably this year as increasing numbers of virtualization solutions come to market. The fact is that storage virtualization is a subject whose underlying complexity is often masked by simplistic marketing terminology. While some vendors discuss storage virtualization as if it were a simple tactical solution, we believe it is more accurate to consider it an enabling technology that can address a range of incrementally evolving business challenges.

A close examination of solutions currently available or planned for release in the first half of 2005 reveals that vendors are taking strikingly different architectural approaches to storage virtualization, with solutions that incorporate everything from in-band appliances to array-based processes to intelligent network switches. Some would say that considering these architectures' relative merits is a waste of time as storage customers are primarily concerned with the practical application of virtualization solutions for immediate storage problems. We disagree with this contention, since every tactical purchase incorporates long-term strategic implications. If businesses are to make informed and intelligent purchasing decisions concerning storage virtualization, it is critical for them to understand underlying architectural issues and how those features map against their immediate and future business requirements.

#### What Is Storage Virtualization?

Despite sometimes blinding quantities of hype, virtualization does not provide a cure for the common cold or the recipe for Kentucky Fried Chicken, but it can solve some and mitigate other key issues facing storage customers. In essence, virtualization enables the creation of logical (virtual) representations of physical IT resources such as memory, networks, servers, and storage -- IT doppelgangers, if you will, that perform as if they were actual resources. In virtualized storage environments, applications can see and interact with these logical components, which are independent from but able to interact with their physical counterparts including SANs, disk arrays, tape components, and other storage media. It all sounds simple enough, so why are many storage vendors and industry associations making such a big deal over virtualization? Quite simply, by freeing applications from physical resources, virtualization offers a simplified and more effective means for centrally accessing and managing infrastructure components and stored information.

#### Potential Benefits/ Probable Challenges

Virtualization also promises to enhance overall platform independence, along with

system flexibility and utilization. Maximizing system flexibility and utilization are critical to ensuring that a storage investment is delivering the benefits and dividends its owners planned on. Overall, virtualization is not simply a fancy way of viewing and interacting with data, but can tangibly improve the performance of enterprise storage and the value of business information.

The problem is that virtualization is still a long way from being an off the shelf solution for every storage woe. Commercial virtualization solutions are not created equal. Many currently available products offer significantly more functionality than did previous versions, and future releases will incorporate additional enhancements. In other words, storage customers should ensure not just that a solution's current capabilities meet company needs, but must also consider how the vendor's plans map against its own future business requirements.

In general, we believe that to be truly effective, storage virtualization solutions must successfully meet a number of intrinsic challenges. First, they need to support customers' existing and critical storage processes, not just a few targeted functions. To perform effectively, they also need to integrate successfully with existing storage management applications and tools. A virtualization solution that requires customers to deploy new or proprietary applications or tools should be considered very carefully, if at all. Additionally, while virtualized solutions promise to ease overall storage complexity, they will require careful planning and appropriate resources. Finally, workable virtualization solutions must scale. What is the point, after all, of spending good money on a product that meets your current needs but eventually leaves you high and dry?

### **Mapping the Ideal Virtualization Solution**

While perfect IT solutions tend to inhabit the same celestial regions as Santa Claus, the Tooth Fairy, and the 2004 Boston Red Sox, they are still well worth considering. So what would the ideal storage virtualization solution look like? Like all effective technologies, virtualization solutions should help solve problems without creating new ones. More specifically, ideal virtualization solutions should possess the following characteristics:

- Non-disruptive – Such solutions should be easy to integrate into existing storage infrastructures, causing minimal pain, headaches, or heartbreak
- No effect on data integrity – Enhancing information value is not possible if its integrity is compromised
- Allow access – To array processes and functions. Solutions that require customers to deploy proprietary applications or specialized interfaces should be considered carefully or avoided entirely
- Highly scalable/extensible – To allow for future expansion of SAN or other storage infrastructure components
- Little/no added latency – No one needs a solution that benefits one element of storage performance at the expense of another
- Play well with others – Existing management tools and applications, that is. The best solutions require few if any new or proprietary additions
- Open – Supports multiple options and does not force customers into single vendor

relationships

### **Virtualization Architectures – Approaches and Considerations**

How does this ideal list stack up against solutions based on common virtualization architectures? To begin, most commercial solutions are based on one of three architectural models: appliance-enabled, array-enabled, or network-enabled.

In practice, appliance-enabled solutions feature a midrange storage controller/server placed in the data path to perform a variety of virtualization functions. However, the controller's position in the data path tends to add latency, though some vendors enhance controller performance as a means of addressing this issue. In addition, some appliance-enabled solutions utilize cache for copy processes. While this approach can enhance performance, holding data in both the appliance and the array can create data integrity issues. While it is possible to address many of the technical and performance shortcomings of appliance-enabled solutions, some management and functionality limitations are more troubling. All too often, such solutions only support network replication and provide no access to array replication processes. In addition, many solutions require customers to use separate or proprietary management tools. While this issue might be a minor concern for customers already utilizing the same vendors other products, we believe it detracts from these solutions' usability among other businesses.

Array-enabled solutions share many similarities with appliance-enabled solutions. A specialized array that performs virtualization processes is placed in the data path, potentially causing many of the same latency problems. Some array-enabled solutions also perform copy processes in cache, causing potential data integrity issues. With such solutions, scalability is limited to the capacity of the specialized array (which customers will use for both virtualization and conventional array processes), an issue which limits the usefulness of such solutions for some businesses. Like appliance-enabled solutions, array-based products also possess significant functionality and management challenges. Many such solutions support only their own replication processes and provide little or no access to functionality in virtualized arrays. In addition, they often require customers to use separate management tools for SANs and virtual arrays, increasing overall complexity. While we believe array-based solutions will find some success among their vendors' existing customers, we doubt their potential outside this circle.

Network-enabled solutions offer software-based processes compatible with and deployed on intelligent switches and directors from many popular SAN vendors. These solutions perform at full wire speed and suffer few if any of the latency problems common to appliance- or array-enabled solutions. In addition, since network-enabled solutions are inherent parts of the SAN, they scale as the SAN scales, providing considerably more flexibility for future growth than other architectures. Network-enabled virtualization solutions also offer customers wider choices for replication, as they often support both network replication processes and provide access to array replication. Finally, since such solutions tend to incorporate management tools and options that are integrated with existing storage and network management applications, they are far less likely to limit users as many single vendor products do.

### **Mission Accomplished?**

After months of hype and discussion, increasing numbers of commercial storage virtualization solutions are now or will soon become available. However, not all of these solutions are created equal. Along with price, performance, and feature issues, businesses shopping for storage virtualization solutions will likely need to sort out a bewildering array of claims and counter claims concerning these products. We believe that most virtualization solutions are fully capable of delivering short-term benefits but that many also offer customers long-term challenges related to architectural design issues. As a result, it is critically important for businesses considering such solutions to understand clearly these architectural differences.

Storage virtualization solutions based on network-enabled architecture offer an elegant approach to supporting functionalities and processes that are critical to most business users. In general, such solutions tend to be largely non-disruptive and easily integrated with existing infrastructures and applications, and provide access to network and array replication processes. In addition, they offer high scalability, minimal latency, and little or no impact on data integrity. Overall, we believe that enterprises seriously considering the purchase of storage virtualization solutions would do well to investigate those based on network-enabled architectures.

© 2005 Pund-IT, Inc. All rights reserved.

Contact:  
Pund-IT, Inc.  
Phone: 510-909-0750  
E-mail: [charles@pund-it.com](mailto:charles@pund-it.com)  
Web: [www.pund-it.com](http://www.pund-it.com)