

The Next Phase of SANs and IT Consolidation

April 2006

Copyright © 2006 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, EMC ControlCenter, AlphaStor, ApplicationXtender, Captiva, Catalog Solution, Celerra, CentraStar, CLARAlert, CLARiiON, ClientPak, Connectrix, Co-StandbyServer, Dantz, Direct Matrix Architecture, DiskXtender, DiskXtender 2000, Documentum, EmailXaminer, EmailXtender, EmailXtract, eRoom, FLARE, HighRoad, Navisphere, OpenScale, PowerPath, Rainfinity, RepliStor, ResourcePak, Retrospect, Smarts, SnapShotServer, SnapView/IP, SRDF, Symmetrix, TimeFinder, VisualSAN, VSAM-Assist, WebXtender, where information lives, Xtender, and Xtender Solutions are registered trademarks and EMC Developers Program, EMC OnCourse, EMC Proven, EMC Snap, EMC Storage Administrator, Acartus, Access Logix, ArchiveXtender, Authentic Problems, Automated Resource Manager, AutoStart, AutoSwap, AVALONidm, C-Clip, Celerra Replicator, Centera, CLARevent, Codebook Correlation Technology, Common Information Model, CopyCross, CopyPoint, DatabaseXtender, Direct Matrix, EDM, E-Lab, Engenuity, FarPoint, Global File Virtualization, Graphic Visualization, InfoMover, Invista, MirrorView, NetWin, NetWorker, OnAlert, Powerlink, PowerSnap, RecoverPoint, RepliCare, SafeLine, SAN Advisor, SAN Copy, SAN Manager, SDMS, SnapImage, SnapSure, SnapView, StorageScope, SupportMate, SymmAPI, SymmEnabler, Symmetrix DMX, UltraPoint, UltraScale, Viewlets, and VisualSRM are trademarks of EMC Corporation. VMware is a registered trademark of VMware, Inc. All other trademarks used herein are the property of their respective owners.

Part Number H2156

Table of Contents

The Next Phase of SANs	4
Non-disruptive operations	4
Efficient data movement	4
The next step for enterprise SANs.....	4
An EMC and Brocade Solution	4
Intelligence in the Fabric.....	6
Which storage services will leverage network intelligence?	6
What are the benefits of placing this intelligence in the network?	8
The Platform: Intelligent Switch	10
Brocade XPath Technology	10
A More Effective Model for Fabric Applications	10
XPath Design Goals.....	11
XPath Technology Components	11
XPath Partitioned Processing.....	11
Split Control and Data Path Processing	12
Distributed Data Path Processing	12
XPath Storage Processors	12
The Solution: Invista.....	13
Deployment configurations	14
Conclusion	15

The Next Phase of SANs

In the past, the forces of consolidation and decentralization were significant drivers of IT architectures. Today, the need for consolidation is greater than ever. For instance, increasingly rapid proliferation of distributed data and systems has strained the resources of staff-limited IT organizations driving a renewed focus on the consolidation, centralization, standardization, and automation of IT infrastructures.

Storage Area Networks (SANs) have been a key enabler for an ever-increasing IT consolidation trend. SAN fabric connectivity has facilitated the consolidation of storage and server resources, helping organizations to reap total cost of ownership (TCO) benefits from storage resource sharing, increased data protection, higher availability, and greater operational efficiency. The maturing consolidation of network and resource management has made SANs easier to deploy and manage as they grow to encompass larger parts of enterprise IT infrastructures. These maturing capabilities have resulted in SAN-attached storage being deployed from the most basic levels of the enterprise to mission-critical environments.

The next phase of SAN maturity will include the consolidation of the appropriate storage management functionality in the form of SAN fabric-based software applications. Because they will be optimized for deployment within SAN fabrics, these new fabric applications will be able to provide even greater levels of storage availability, data mobility, and data management.

Non-disruptive operations

SAN-based storage has matured to the point where it is regularly deployed in mission-critical environments where application downtime is unacceptable. Certain mandatory storage operations on a SAN are by their very nature disruptive and require host or application downtime. These operations range from basic administrative tasks such as changing storage volume attributes or protection levels to significant periodic tasks such as data migrations.

IT organizations invest significant funds and resources to manage tasks associated with volume management and data migrations: planning, scheduling, provisioning, data movement, application reconfiguration, and retiring the source. Scheduling planned downtime is often the number one pain point, given the need to manage and minimize the impact to the customer and the business.

Fabric-based applications like EMC® Invista™ can accelerate the deployment of new, more functional, higher capacity device technologies into existing infrastructures by making it quicker and easier to non-disruptively move existing data onto the new devices.

Efficient data movement

Data movement capabilities are becoming increasingly important to networked storage implementations. As organizations begin to employ tiered-storage infrastructures and information lifecycle management (ILM) strategies, they need a simple and non-disruptive method for moving data from tier to tier as certain requirements change.

The next step for enterprise SANs

Reducing downtime by enabling non-disruptive storage operations and efficient data movement in support of ILM is an essential “next step” for enterprise SANs. EMC Invista and the EMC Connectrix® AP-7420B Intelligent Switch provide a powerful new technology platform that together makes, network-based storage virtualization hosted on an intelligent switch an exciting new reality.

An EMC and Brocade Solution

EMC and Brocade have teamed together to create a solution that addresses these issues. To help solve these data management challenges

- Brocade has developed the SilkWorm Fabric Application Platform, the base for the EMC Connectrix AP-7420B Intelligent Switch, as a platform for a new class of fabric intelligence based on industry-unique Brocade XPath Technology. The EMC Connectrix AP-7420B Intelligent Switch is specifically designed to support fabric

applications at wire-rate performance while seamlessly integrating with existing SAN infrastructures. As a result, the Connectrix AP-7420B provides a much simpler and more scalable solution for hosting fabric applications compared to current alternatives, such as fabric-attached appliances based on standard servers.

- EMC Invista is the industry's first high-performance networked based storage virtualization solution that leverages intelligent SAN switches. This new technology places virtualization intelligence into the existing SAN infrastructure while allowing for the highest levels of scalability and availability. Invista's functionality addresses the ever-expanding requirement for non-disruptive operations; keeping applications available, reducing administrative overhead, increasing the effectiveness and utility of current storage resources and, reducing costs by maximizing current storage investment and minimizing management effort.

As industry leaders in their respective fields of expertise, Brocade and EMC have a strategic business relationship spanning many years. As a result, they have acquired unmatched levels of expertise and experience in storage networking, uniquely qualifying them to be innovators in SAN solutions.

EMC has a strong depth and breadth of storage application expertise that they are now looking to deploy certain functionality within the fabric utilizing the Connectrix AP-7420B Intelligent Switching platform.

Intelligence in the Fabric

Although SANs have excelled at boosting productivity and return on investment (ROI), many organizations have yet to realize the full potential of their SAN environments. For example, although organizations can consolidate disparate systems and storage through SAN connectivity, higher storage infrastructure management functions such as host-based volume management or data replication are often fragmented across multiple models of servers, operating systems, and storage systems. This fragmentation can lead to an overly complex, difficult-to-manage SAN environment, with the potential of arduous support issues as a result of deploying a heterogeneous environment with products from multiple vendors. Limitations presented as a result of this type of infrastructure can often result in restrictive procurement and deployment options or a complicated, disparate set of tools for managing host and array storage functions.

Intelligent storage networks represent an evolutionary step in networked storage architectures. Storage functions such network-based volume management, non-disruptive data mobility, and some replication technologies will be located on intelligent switches to deliver highly scalable networked storage services that are host- and array-independent. The storage network in effect becomes a single virtual storage pool delivering many of the benefits of a single monolithic storage device, such as centralized uniform management behaviors, providing organizations continuous economic and storage management benefits and flexibility as their enterprises grow.

Which storage services will leverage network intelligence?

Network-Based Volume Management: Volume manager software presents a flexible view of data independent of its underlying physical storage. By separating the physical devices and abstracting them to logical resources, network-based volume management software provides the following major benefits:

- Scalability and flexibility—Volumes can be presented with a 1:1 physical to logical ratio, span multiple disks per LUN, and be expanded or dynamically changed online as required
- Enhanced performance—Supports striping which enhances performance as a result of parallelizing reads and writes across a number of devices
- Improved availability—Supports dynamic reconfiguration while applications remain online

Virtualization at the network layer with intelligent fabrics will add a new layer of abstraction that will span all storage subsystems in the data center. Network-based volume management offers the ability to extend the services currently residing on several disparate host systems and expand it across all operating systems and storage throughout the entire storage network.

Invista in conjunction with the Connectrix AP-7420B can aggregate capacity from many different types of storage arrays. Storage administrators can choose which volumes from a given array they wish to virtualize then dedicate that capacity to a class of storage pools where individual volumes then get created and presented to hosts as subsets of the virtualized capacity from these pools. Volume size and structure are completely independent of structural restrictions imposed by the arrays.

Once created, hosts are presented virtual volumes (or virtual LUNs) from the intelligent switch in the network. Invista and the Connectrix AP-7420B provide a solution architected on the SCSI-3 standard which allows these virtual volumes




Figure 1 - Virtual Addressing

to appear to a host just like a standard SCSI-3 volume from any array. The core function of network-based volume management is the abstraction and mapping of the virtual LUN to the corresponding physical storage without host or network performance impact.

The EMC Invista and Connectrix AP-7420B solution provides network-based volume management becoming centralized in the infrastructure with no need to deal with individual host-based tools or to load host-based software in some cases. These functions will be available to all qualified connected hosts and work on storage from all qualified connected arrays.

Data Movement: One of the biggest IT challenges today is the ability to ensure continuous uptime to the business. This problem is getting worse, not better, in today's 24x7 environments where more mission-critical applications store data on the SAN. Consider the mechanics involved in integrating a new storage array into the existing realtime infrastructure. Storage administrators need to consult and coordinate with several IT layers to schedule appropriate planned downtime windows that successfully allow the business to meet its availability requirements while moving data across the storage infrastructure.

Once the foundation of a network-based virtual volume with a virtual address is in place, additional high-value services can be layered on top to take advantage of the flexibility that is offered. One such

service is dynamic volume migration. Dynamic volume migration allows storage administrators to move volumes from any location in the SAN to another while the host and application remain online.

To move a network-based virtual volume, Invista and the Connectrix AP-7420B together simply perform a redirection of I/O from one physical location to another while protecting the source throughout the process until the operation is complete. Despite the fact that the I/O is being redirected to a new location by Invista, the logical network-based WWN address of the virtual volume presented to the host never changes. This capability allows the data migration process to be transparent and non-disruptive to the host. Additionally, since this can be accomplished leveraging the capabilities of the virtualization solution, no host cycles are required, freeing servers to be dedicated to their application-centric functions.

There are a number of very valuable real-world uses for this dynamic volume migration capability:

- Lease roll-overs and technology refreshes
- Volume moves or changes in response to changing performance or cost requirements in support of ILM strategies (application reprioritization)

Lease Roll-Overs/Technology Refreshes

One significant challenge organizations have today is getting new arrays deployed in their environment. Today's migration solutions can be highly complex, manual, and disruptive, making the process extremely labor-intensive and inefficient. Dynamic volume migration capabilities as part of a network-based virtualization system can significantly improve the process.

An organization can bring a new array into the environment, virtualize enough capacity to hold the source data, and then set a mobility job in motion to copy the data from the source to the new array. All of the data copying occurs in




Figure 2 - Dynamic Volume Migration

the background without host impact. Once the copy is complete and consistent, the cutover point is chosen for the I/O to be redirected to the new array.

Once the data has been successfully moved to the target array, the old array can be removed from the environment. All of this happens while the applications accessing the virtual array and volume remain online. Furthermore, there is no performance penalty on the hosts as there would be with a host-based volume manager.

Changing Performance or Cost Requirements

Now let's look at how using Invista and the Connectrix AP-7420B can facilitate performance moves. Picture a configuration with medium and high-performance storage pools. A storage administrator observes that a volume in the medium pool is not meeting a specified SLA. Therefore, to remain in compliance with the service level guaranteed to the application owner, a dynamic volume mobility job is set in motion to migrate this volume into the high-performance pool.

Just as in the lease-rollover example, a copy is performed in the fabric with no host impact and no application interruption. Once the copy is complete and consistent, I/O is redirected to the new volume in the high-performance pool.

In the future, policy-based automation capabilities will transparently and automatically optimize the performance of the storage infrastructure in accordance with business requirements. This seamless mobility capability based on changing performance requirements will be key in supporting advanced information lifecycle management strategies.

Heterogeneous Replication: Replicating copies of data across the storage infrastructure is accomplished today via array replication technologies. Some vendors support only homogeneous replication which can provide challenges and be disruptive. However, the capability of replicating "production" data across heterogeneous storage environments can now be best accomplished at the network layer with Invista and the Connectrix AP-7420B.

Intelligent storage networks are ideally positioned to address the requirement for replication services that span multiple storage systems for every day storage tasks like backup and data recovery scenarios. An example is local copy-based replication between arrays of different types in order to reduce costs at a disaster recovery site by using less expensive storage where the data is there only for emergencies and not primary operations.

What are the benefits of placing this intelligence in the network?

Especially for large enterprises, fabric applications on intelligent switches provide a wide range of benefits, including:

- **Increased management scale and efficiency:** Fabric applications are deployed on a few scalable, intelligent fabric switches that can span large numbers and mixes of SAN-connected servers and heterogeneous storage systems. As a result, organizations can consolidate the number of "management images and footprints" for their storage functions, which will enable a small number of IT personnel to efficiently manage the growing number of devices and ever-increasing storage capacity.
- **New levels of management functionality and flexibility:** Although the standalone functionality is often similar to that provided by current storage and server systems, fabric applications deliver new levels of management flexibility and efficiency. For instance, organizations can consistently deploy storage functions independent of and across mixed storage and server systems. Point-in-time (PIT) copy volumes can reside in different storage cabinets from the source copy, and data volumes can be migrated effortlessly between dissimilar storage systems. Higher flexibility and functional scalability are key enablers for implementing storage as a "utility service." As a result, organizations can leverage a new class of data management application that is practical only through a fabric-based platform.
- **New capabilities for data management:** The fabric application architecture significantly increases data mobility across multiple classes of storage, and enables fundamentally advantageous methods for data

management.

- **Simplified management:** Centralizing storage management functions in the fabric can simplify administration by reducing the number of management touch points. As a result, organizations can operate more efficiently and increase the level of their productivity. Capital and operational cost reduction through asset optimization and deployment flexibility: The unification of storage functions across mixed storage environments optimizes the use of storage capacity and performance against varying application requirements increasing storage procurement flexibility.

Locating the essential functions of storage resource management, device presentation, and configuration at the storage network level provides for unprecedented management flexibility and resource utilization.

Brocade and EMC are building new applications, introducing intelligence and automation into the SAN fabric to provide new ways to simplify data management and reduce overall storage costs. As Brocade and EMC develop new options for increasing SAN fabric intelligence, they are helping SAN architects by providing them with practical examples of how they can begin planning today to implement new data movement applications without disruption to their existing SAN environments while ultimately reducing overall storage costs.

The Platform: Intelligent Switch

Shifting from one technology paradigm to the next, introducing rapid change in the storage infrastructure, the presentation of new intelligent storage networks should be additive to the process. Existing storage environments should be able to upgrade existing components relatively seamlessly and gradually to take advantage of the numerous benefits that intelligent storage networks provide. Existing fabrics can add intelligent switches without “rip-and-replace.”

EMC has chosen to leverage the Connectrix AP-7420B Intelligent Switching platform for deployment of network-based storage virtualization. Intelligent switches have additional processing power that enables them to run storage applications. They do this by using specialized port-level processors that can perform the basic operations of virtualization (I/O redirection) at wire-speed. The intelligent switch is thus a high-performance platform that can be easily incorporated into the SAN infrastructure that exists today. EMC is working with Brocade to deploy Invista on an intelligent switch platform.

Brocade XPath Technology

Brocade XPath Technology is a comprehensive set of functions and subsystems for deploying a new class of fabric applications. The EMC Connectrix AP-7420B leverages XPath Technology to provide significant advantages over alternative approaches for the deployment of fabric applications.

A More Effective Model for Fabric Applications

By hosting certain storage application functions directly within a SAN fabric, XPath Technology enables an application instance to span multiple SAN-connected host and storage systems. This consolidated deployment model reduces management costs while extending application functionality and flexibility. In contrast, existing approaches for deploying application functions within SANs have caused difficult technical tradeoffs and cost-of-ownership issues resulting in their limited success.

For example, in-band appliances using standard computing platforms do not scale effectively because they require a general-purpose server to process every storage data stream “in-band.” Common scaling limits include PCI I/O buses limited to a single 2 Gb/s data stream and contention for centralized processor and memory systems that are inefficient for data movement and transport operations.

Likewise, out-of-band appliances distribute basic storage virtualization functions to agent software on custom host bus adapters (HBAs) or host operating system drivers in order to avoid a single data path bottleneck. However, high-value functions such as storage volume sharing and data replication and migration must be performed on an off-host appliance platform with similar limitations as in-band appliances. Moreover, the installation and maintenance of custom drivers or HBAs on every host introduces a new layer of host management and performance impact. In contrast, XPath Technology addresses these architectural limitations while remaining compatible with the software applications developed for these varied approaches.

XPath Technology makes this type of environment practical to deploy by providing an off-host, off-storage application platform that achieves the following design goals:

- Network-based virtualization allowing performance transparency through line-rate data processing with the throughput and latency required to avoid incremental performance or administrative impact to existing hosts and storage systems
- Linear scalability for simple and complex processing of storage I/O operations
- Compact and cost-effective deployment footprints
- Data movement scalability to fulfill the vision of site-wide, always-on data replication
- Transport-neutrality across Fibre Channel, IP, and other protocols to make transport an administrative choice rather than an architectural decision
- Investment protection by interoperating with existing equipment while leveraging new and existing storage software tools

- Greater data reliability over appliances that cache data. By processing data at wire-speed through the switch, all writes are safely stored on attached storage resources before any writes are acknowledged to the host. Appliances that cache data run the risk of data corruption with the loss of the appliance or array performing virtualization.

XPath Design Goals

The target environment for XPath Technology is a next-generation SAN with the following characteristics:

- Network storage virtualization to increase management and resource efficiency for larger numbers of servers and storage systems.
- Dramatically increased, site-wide data replication and movement across a hierarchy of storage systems that enable significant improvements in data protection and disaster recovery.

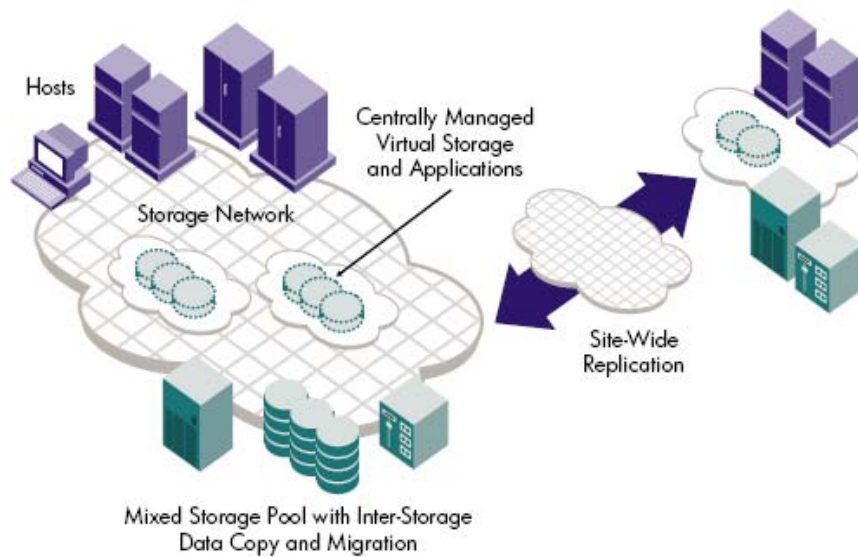


Figure 3 - Key characteristics of an environment that can benefit from XPath technology

XPath Technology Components

XPath Technology includes four major components that jointly support the performance, scalability, and software deployment goals described above:

- **XPath Partitioned Processing:** Utilizes distributed control and data path processors to scale storage network software
- **XPath Storage Processors:** Provide line-speed processing of storage data by using a rich set of storage-optimized hardware acceleration engines
- **XPath Multi-Protocol Fabric:** Provides a low-latency, protocol-neutral interconnect that integrally links all components with any-to-any non-blocking throughput
- **XPath OS and APIs:** Unify a distributed platform and provide tools and services for third-party storage applications to exploit XPath Technology

XPath Partitioned Processing

To effectively scale storage applications, XPath Technology combines the key techniques of functional partitioning and distributed processing across multiple levels of processing elements. Storage transport and application functions are deployed at the optimal level or levels based on the required performance and functions.

For instance, storage transport layer functions are implemented primarily in custom hardware; storage application functions are implemented in a mix of custom hardware, realtime firmware, and general-purpose software; and storage network management functions are implemented in software running on a general-purpose control processor.

Split Control and Data Path Processing

Split control and data path processing exploits the general nature of fabric applications to greatly increase their scalability and performance. This split processing isolates the most frequent and performance-sensitive functions, and physically distributes them to a set of replicated, hardware-assisted data path processors leaving more complex configuration coordination functions to a smaller number of centralized control processors.

The performance and scalability of fabric applications are determined by a few performance-sensitive data path operations such as I/O reads and writes that account for the vast majority of traffic in the storage network. In addition, the relatively simple logic of these operations can be reduced to functional primitives common to many types of storage applications. XPath Technology defines data path processors (also referred to as “XPath Storage Processors”) that are specifically optimized for these performance-sensitive functions.

The remainder of network traffic consists of control path operations for configuration, administration, and exception operations. While accounting for the majority of the functions and complexity of storage applications, these operations are relatively infrequent and can be handled by a smaller number of general-purpose control path processors.

XPath Storage Processors “split” storage network data streams into control and data path operations in real time by directly processing the majority of data path operations within the processor. Data path operations that cannot be successfully handled as well as any explicit “control” operations are forwarded to a control processor for completion.

Forwarding boundary case operations to control path processors enables XPath Storage Processors to be optimized for a low-overhead, hardware-assisted processing model for data path operations. Conversely, isolating the operations of control path processors avoids consuming a centralized CPU and memory system with bandwidth- and copy-intensive workloads, as is the case in appliance- and host-based processing models.

Distributed Data Path Processing

Another scaling mechanism of XPath Technology, distributed data path processing exploits the effectiveness of parallel processing against the largely independent data streams within a storage network. Distributed processing works in conjunction with control/data path separation and the high-scale XPath Multi-Protocol Fabric interconnect to provide a system with much greater application performance and scalability. Because the XPath processing units are performance-optimized for storage data traffic, they can be more efficiently aggregated into a single control and management image that delivers millions of virtual I/Os¹ and multiple gigabytes of concurrent throughput in a compact footprint.

XPath Storage Processors

XPath Storage Processors are customized storage network processing units dedicated to performance-sensitive storage application data flow. These processors implement large parts of the optimized XPath Technology by integrating the following features in every unit:

- Multi-protocol, multi-Gb/s storage network I/O ports
- In-line RISC processors with local memory and frame buffers
- A rich set of storage operations engines and hardware assists

¹ Virtual I/Os are defined as one complete round trip of storage control and data traffic required to pass a 512-byte I/O read request through a network processing element to a physical target.

The Solution: Invista

EMC's network-based virtualization product is called Invista. We will now turn our discussion to how it delivers this functionality. Key to its ability to deliver is the product design and architecture.

EMC Invista is based on innovative new technologies to help customers achieve non-disruptive enterprise operations, eliminate planned downtime, and to centralize and streamline storage management. The integrated hardware and software solution delivers these benefits through the simple and non-disruptive dynamic volume migration, local copy and migration capabilities of data across multiple tiers of heterogeneous storage arrays.

EMC Invista provides for the creation of network-based "virtual volumes" within the intelligent storage network, forming a dynamic environment in which physical storage resources can be moved and changed rapidly and non-disruptively. Invista groups distributed physical storage devices into a common pool of storage resources. From that pool, customers can easily provision and manage their disparate information resources.

Built with a highly scalable "Split-Path" architecture, Invista delivers the full value of network storage virtualization with the levels of performance, reliability, and integrity required for deployment in enterprise data centers. Invista provides the foundation for customers to take advantage of network storage virtualization today, while serving as a platform to deploy new network storage virtualization capabilities easily and cost effectively.

Invista is the only network-based virtualization solution that enables enterprise customers to preserve all of the native performance of their current storage infrastructure without limiting the environment to a single common denominator like appliance and array-based solution. Businesses can opt to virtualize their heterogeneous storage environment while continuing to use supported array-based replication or other storage software functionality currently running on their storage systems. This provides investment protection across the infrastructure that no other solution can.

Provides Non-disruptive Migrations, Eliminates Repetitive Volume-Management Tasks

Currently storage administrators can spend 20-30 percent of their time on individual operating-system-specific host-based volume-management tasks. Adding Invista and presenting a network-based virtualized volume allows a single tool to be used across multiple heterogeneous operating systems and storage devices. EMC Invista enables IT organizations to significantly reduce the amount of time spent on these manual tasks. With centralized network-based volume management in place, Invista customers can reduce the repetitive volume management activities that must occur individually in the server environment today.

Invista's volume mobility enables storage administrators to dynamically and without disruption move storage volumes from any storage device to another—without application disruption—for normally disruptive processes such as lease rollovers, technology refreshes, data movement across multi-tiered heterogeneous environments or to respond to rapidly changing performance needs.

Invista's network-based local replication provides additional flexibility and choice. For example, businesses can create additional data copies non-disruptively for backup, data warehousing, or other secondary uses. Because the data is replicated via the storage network and the intelligent switches, it can be copied to and from heterogeneous and tiered storage arrays.

Innovative Architecture, Open and Integrated

EMC believes that virtualization will exist throughout the IT infrastructure and offers the industries most comprehensive solutions through host (VMware®), NAS (Rainfinity®) and SAN (Invista). Invista's Split-Path architecture is also critical for delivering value to customers. Unlike in-band (appliance or array-based) storage virtualization architectures that introduce significant bottlenecks by performing all of their processing within the data path, Invista is based on an Split-Path approach that places the virtualization intelligence in the storage network to eliminate impact on server or application performance.

Through the use of open APIs (application programming interfaces), Invista prepares enterprises for the emergence of the Fabric Application Interface Standard (FAIS) API for the rapid deployment of network-based applications. Invista also is integrated with EMC ControlCenter®, EMC's industry-leading family of storage resource and device management software.

Deployment configurations

Deploying Invista involves hardware and software. The Invista software runs on a highly-available system called the Control Path Cluster (CPC). The CPC configures and interacts with intelligent switches. The intelligent switches may be new switches or customers may be able to upgrade their existing switches to make them intelligent. Figure 4 presents a simplified look at the components of Invista and how they could be deployed in a SAN.




Figure 4 Deploying the EMC Connectrix AP-7420B with EMC Invista

Hosts can connect directly to the intelligent switches or connect through a standard Layer 2 (non-intelligent) front-end fabric for fan-in (as depicted). Arrays can be connected to the intelligent switches directly (as depicted) or via a back-end fabric.

Conclusion

Fabric-hosted, network storage virtualization solutions are the next natural phase in the evolution of SANs which have already contributed significantly to the consolidation benefits in data centers. By leveraging the intelligence within the SAN fabric, EMC Invista on the Connectrix AP-7420B Intelligent Switch will provide a variety of new capabilities. Enterprises that deploy and implement EMC Invista on the Connectrix AP-7420B have the opportunity to gain an unprecedented level of control over their data center storage infrastructure: simplifying management of complex infrastructures, providing a non-disruptive operations capability, and facilitating critical elements of a proactive ILM strategy. This powerful combination will help organizations build highly available, scalable networked application architectures with unified management.