May 2008

Optimizing Storage with EMC Invista

The new generation of storage virtualization – EMC Invista V2.1

Thomas Trainer
Managing Partner
Evaluator Group, Inc.

7720 East Belleview Ave.● Suite 210 ● Greenwood Village CO 80111 ● (303) 221-7867 ● FAX (303) 221-1615 ● www.evaluatorgroup.com

Executive Summary

EMC Invista V2.1, the new generation of storage virtualization, is proving to be a technologically sound offering for business production environments. It is delivering high value functions, as well as preserving and extending storage investments, both within IT enterprises and for corporations as a whole.

New developments in architecture, availability, scalability, upgradeability, and data protection capabilities enable users to deploy a more flexible, network-based storage virtualization infrastructure.

Storage environments challenged by age-old dilemmas…and new hurdles

Clients report that data migrations continue to cost them significant time and money. Valuable human resources are being applied to what should be an automated process. This is delaying the implementation of new technologies and not meeting the need for speed in today's data center.

More and more, IT organizations are held to service-level agreements (SLAs) with their clients, be they internal departments or external customers. These SLAs are, in part, supported by the rapidly maturing concept of tiered storage. However, just simply implementing tiered storage does not ensure that the information will be moved to the right tier at the right time. This requires automation and ease of information movement between tiers.

Invista V2.1 in production

Purdue University, a statewide university system in Indiana, USA, is truly a case study for information technology usage from many angles, such as medical research and treatments, information technology research, and technology management of a multifaceted business environment.

While Purdue uses many server and storage technology products within its various research departments, they have standardized on Invista for their application production environments. This includes human resources; finance (receivables, payables, cash flow, balance sheet, and income and expenses); email; staff development (tenure tracking, pay scales, career development) and student enrollment and registration (which are handled by Purdue as a subscription service with associated accounting).

The university's servers and applications run 24x7x365. Though humans may not be onsite for some of the lights-out periods, storage managers rely on 99.999% uptime for their servers and storage environments to support their critical applications.

Purdue's storage growth, while episodic like many users, typically grows an average of 30% per annum. EMC's storage footprint at Purdue has grown from 9 terabytes (TBs) in 2002 to over 225 TBs at the end of 2007.

Purdue deployed EMC Invista (after evaluating other technology) to complement their VMware virtual server environment. Their choice to implement Invista was clear, for the following reasons: (1) substantial time would be saved on migrations and storage management, (2) data could be moved non-disruptively across heterogeneous storage, without impacting users or applications, and (3) storage utilization would be greatly improved.

For this review, measurements of implementation success included: pre-installation planning, implementation execution, ongoing management, scaling, training availability and overall satisfaction. Purdue rated EMC at 100 out of 100 in all six areas.

Purdue commented that "EMC's planning was clear and the implementation was done quickly, without error. Our experience with Invista has proven its advantages and supported our decision to implement it". Additionally, Purdue said ". . .there was no impact to the business production environment to implement Invista".

EMC Invista V2.1…next-generation technology

Over the past few years, the industry has watched Invista evolve from an early-version implementation to the new generation of SAN virtualization. Invista V2.1 provides new capabilities, while continuing to build on its unique split-path architecture. Invista now provides customers with greater redundancy, more robust high availability, and increased preservation of storage investments via enhanced protection capabilities (more on this later). More scalable performance and non-disruptive upgrades to all future platforms are also key components of this product.

Understanding Invista's architecture

Split-path, simply put, means there are separate paths for the I/O being written from host to storage (data path) and the control and configuration commands for the intelligent switches (control path). In Invista's split-path architecture, the I/O data moves seamlessly from host to storage through the fabric - just like SANs already move data today. Leveraging its redundant intelligent switches, Invista moves the I/O data at wire speed for superior performance, scalability and high availability. In order to fully grasp its unique architecture though, we need to also understand Invista's "stateless" approach.

Stateless means there is no additional caching stage for I/O writes between servers and storage – data continues to move directly and uninterrupted through the SAN fabric. This is the opposite of in-band, stateful solutions, like IBM's SAN Volume Controller (SVC) and Hitachi Data Systems' Universal Storage Platform V (USP V / VM). In these in-band caching solutions, I/O data moving through the SAN must first be cached in an appliance, stored briefly, and then forwarded on to the physical storage. This creates additional latency and introduces the potential for more problems.

Data corruption can and has occurred with in-band storage virtualization solutions if the appliance fails before all of its cached data has been de-staged and sent to the physical array. See, once the appliance receives the I/O into its cache, it sends a "received" acknowledgment back to the server. But, in fact, the I/O has not yet made it to storage disk; the I/O must still be routed and sent to the appropriate backend physical array. To avoid these "man in the middle" drawbacks, Invista uses its unique architecture.

Don't get me wrong though, caching data is a good thing – but the backend storage arrays already have high performance cache. So when you implement storage virtualization, why would you use an in-band appliance to cache the data *before* the I/O gets to the better-performing array cache? The clear alternative to this is Invista, which can route the I/O from host to storage without caching, as depicted in Figure 1 below.

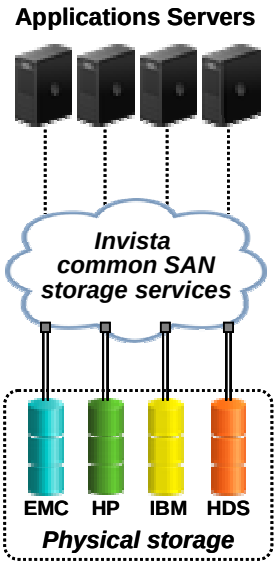

Figure 1: Invista Transfers I/O at Wire Speed

Scaling with Invista…more virtual volumes, more mobility sessions, and back-end load balancing

With a maximum usable capacity of 15.6 petabytes (PBs) and support for up to 8,000 virtual volumes, along with virtual volumes scalable from 100 megabytes (MBs) to 2 TBs, Invista V2.1 provides

even greater scalability than other storage virtualization alternatives.

Concurrent data mobility sessions now include up to 40 simultaneous sessions with a combined throughput of 2,800 megabits per second (MB/s). Clearly this kind of throughput enables users to move information from one tier of storage to another and even from one storage platform to another, at exceptionally high transmission rates. This scaling in capability remains true to EMC's technological background of making information available at the right place at the right time. Figure 2 (below) illustrates a simplified view of data mobility sessions possible within the Invista instance.

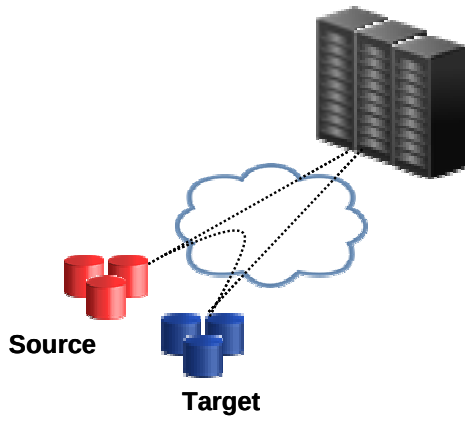

Figure 2: Data Mobility with Invista

Once the virtual LUNs are created in Invista, they are presented to the hosts from the storage area network (SAN). Simply stated, the hosts see Invista as storage. Invista then becomes the storage manager within the SAN, optimizing storage management. This provides a common management platform and greatly simplifies storage management. With Invista, the storage is optimized so managers can spend less time managing the intricacies of heterogeneous storage and reduce the costs associated with storage management software purchases from multiple vendors…it's simply no longer needed in the Invista environment.

Figure 3 (below) illustrates virtual LUNs that are being presented to hosts by the Invista application. LUN management has been improved and no longer resides at the host or storage subsystem level.

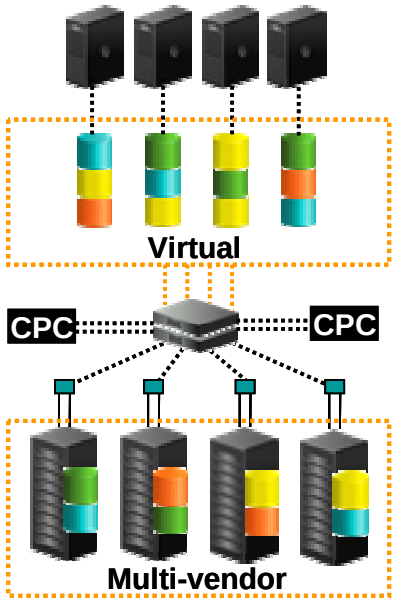

Figure 3: Invista Balances I/O Loads

Non-disruptive upgrades within a network virtualization solution

Many host, appliance, and array-based storage virtualization solutions require some amount of downtime to upgrade their hardware or software. In contrast, Invista V2.1 provides users with the capability, and flexibility, to upgrade any of its hardware or software components non-disruptively, without any impact to the application or user. By leveraging its intelligent switches, upgrades can be done without downtime or stoppage in I/O writes, much like how you would upgrade dual redundant fabrics in a SAN today.

High availability reaches new levels with Invista V2.1

The distributed Control Path Clusters (CPCs) in Invista V2.1 offer users the ability to separate the dual redundant fabric and control components by

up to 500 meters, in order to protect and provide continuous operations in the event of a local site failure.

Additionally, Invista supports RAID1 across third-party storage arrays, enabling users to construct highly available storage configurations for their most critical applications.

Figure 4 (above right) illustrates the distributed CPC node implementation. In earlier versions of Invista, the CPCs had to be located in the same physical location rack (depicted by the dotted line box in Figure 4). Now, each CPC node can be located in a separate location (up to 500 meters apart at this time).

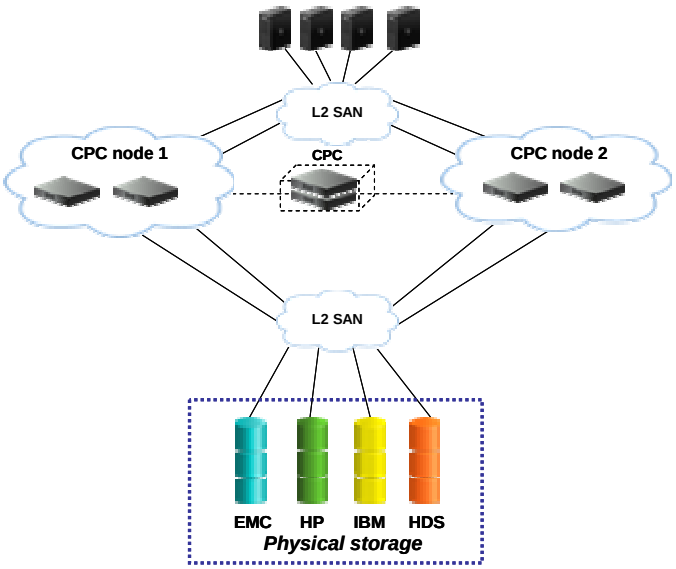

Figure 4: CPC Nodes Can Now be Separated by up to 500 Meters for High Availability

EMC RecoverPoint V3.0…Disaster Recovery and Business Continuity for the virtualized Invista environment

EMC RecoverPoint V3.0 provides replication and protection for the Invista virtual environment. EMC RecoverPoint is an enterprise-class, disaster recovery solution that provides continuous data protection (CDP) and continuous remote replication (CRR), as well concurrent local and remote (CLR) data protection. This provides a single, common solution for both protection and replication of the same data across heterogeneous storage and server environments. RecoverPoint is the only replication solution that leverages fabric-based storage services via EMC Connectrix (the same intelligent switches from Brocade and Cisco that Invista uses) for high performance, high availability and scalability. Its SAN integration makes it easy for the user to deploy, and it protects both physical and virtual source volumes, with no impact to hosts or applications.

RecoverPoint also gives users any-to-any functionality, allowing them to replicate between any combination of virtual or physical sources and targets.

Figure 5 (below) illustrates RecoverPoint replicating a virtual volume from one site to another while running Invista.

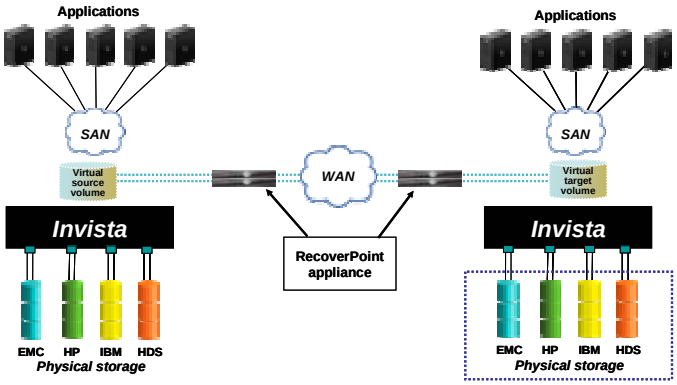

Figure 5: Invista Business Continuity

In this continuous virtualized replication environment, virtual source volumes at the primary site are replicated using RecoverPoint appliances to virtual target volumes at the remote business continuity process site.

Figure 6 (below) is another illustration of virtual volume replication to a physical target at a remote disaster recovery (DR) site.

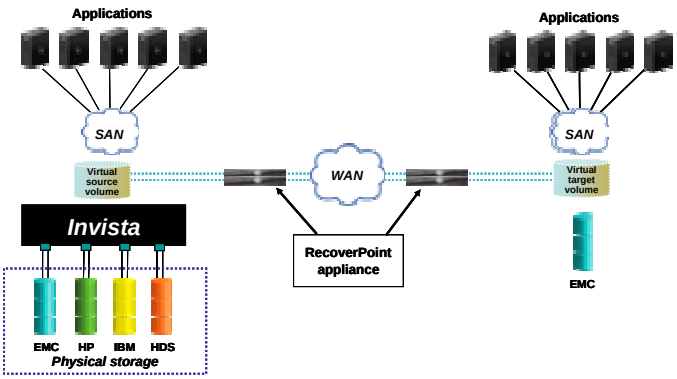

Figure 6: Invista Disaster Recovery with RecoverPoint

By using Invista with RecoverPoint, primary site virtual volumes can be replicated over distance to a remote DR site. The data at the DR site can then be used to restore the primary site, or another site of similar configuration, at the appropriate time subsequent to a business outage.

As with all high availability solutions the kind and type of solution selected is greatly dependent upon the critical nature of the business and the applications that run the business. Recovery time objective (RTO) and recovery point objectives (RPO) are critical metrics for users to consider when designing any highly available process environment.

VMware leverages advancements in Invista V2.1

Beginning with Invista V2.0, VMware now recognizes Invista within the ESX 3.x environment. Invista is VMware–certified and listed within the VMware Compatibility Guide.

Replication Manager integration is also supported within the VMware ESX 3.O environment and we believe some highly creative VMware and Invista storage virtualization solutions are bound for the market in 2008- 2009.

Figure 7 (below) illustrates Invista within a VMware environment.

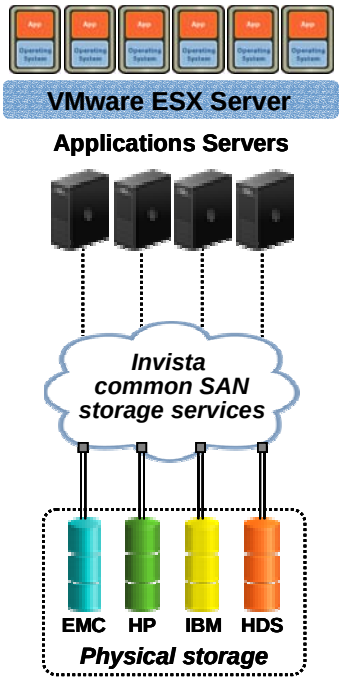

Figure 7: Invista in the VMware Environment

Summary

EMC's Invista V2.1 demonstrates that it is a proven and technologically sound offering for business production environments. It provides valuable functions and preserves storage investments.

New developments in Invista's architecture, availability, scalability, upgradeability, and protection capabilities are enabling users to have more choice in deploying flexible network-based storage virtualization infrastructures.

7720 © East Belleview Ave.● Suite 210 ● Greenwood Village CO 80111 ● (303) 221-7867 ● FAX (303) 221-1615 ● www.evaluatorgroup.com