

## EMC Invista: The Easy to Use Storage Manager

Invista delivers centrally managed LUN Virtualization, Data Mobility, and Copy Services running within a SAN. Invista LUN virtualization services are similar to those provided by host-based volume management solutions. Volume management is the most desired feature users most want to see in their SANs. Early Invista adopters found that volume management reduced complexity and cost compared to volume management software running on each server.

Invista Data Mobility services eliminate the pain and cost of data migrations. Invista-powered data mobility services eliminate the application downtime associated with traditional data migration methods. Applications continue to run while Invista moves data between arrays in the storage network. This works extremely well as customers add or subtract arrays in the network.

Invista copy services can make working clones of existing storage volumes. Clones of production storage volumes can be used for application testing and data mining. A database clone can be used by an application development team to test enhancements. Similarly, a copy of a month-end database can be used for search intensive tasks like recognizing buying habits and price sensitivity.

Invista uses a technology approach known as Split Path Architecture (SPA). The Invista SPA uses intelligence found in intelligent storage switches as a data mover. The Invista software, which runs on clustered Invista platforms, directs traffic and handles exceptions.

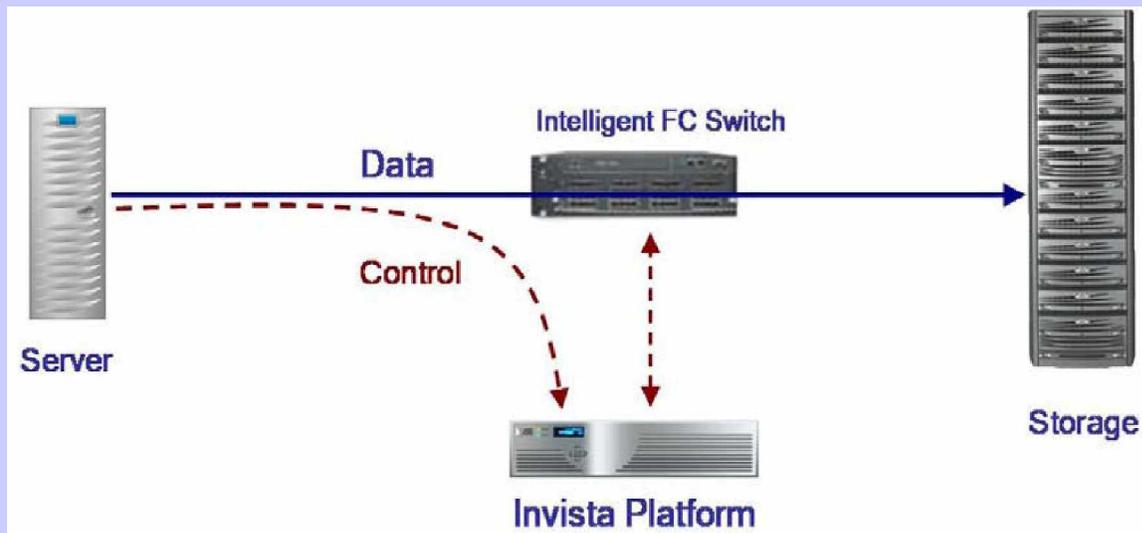


Figure One: The Invista Split Path Architecture

## ESG Lab Validation Report

Invista software running on an Invista hardware platform communicates with the switch using an application programming interface provided by the switch vendor. After Invista software has downloaded an I/O mapping to an intelligent switch, the switch, using its high degree of port level processing power and bandwidth takes over.

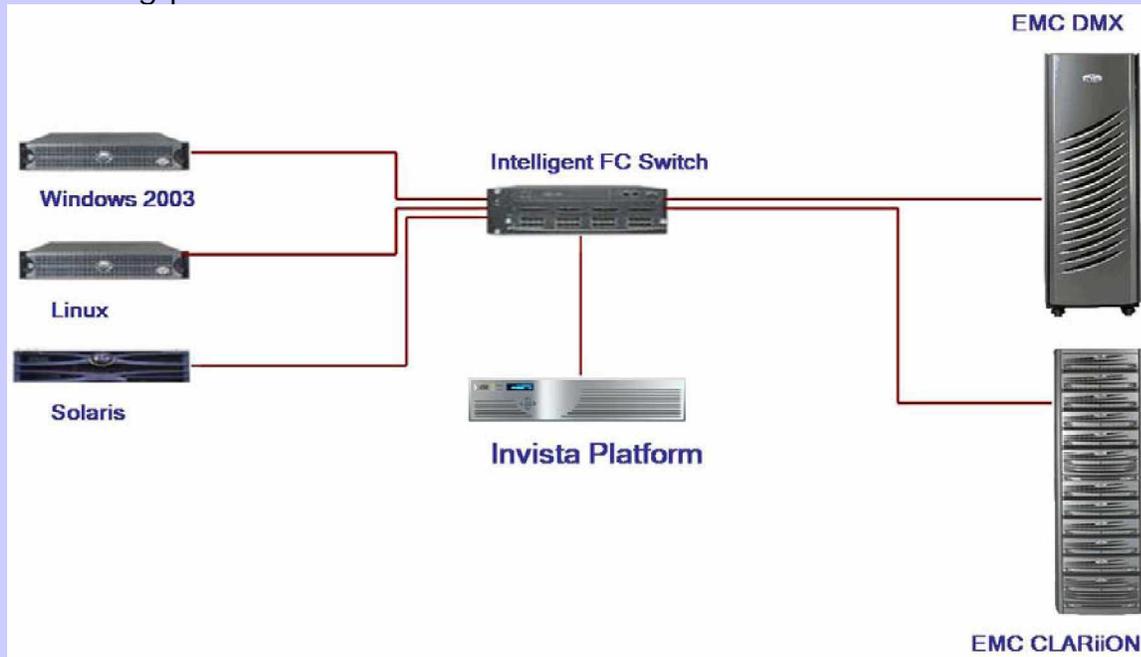


Figure 2: Invista Test Bed

EMC's Enterprise Research Group (ESG) has been tracking the evolution of intelligent FC switches. ESG is impressed with the high throughput and low latency that they can provide, and has concluded that intelligent switches introduce negligible latency. Thus they can be to drive data movement at extremely high throughput rates with little or no impact on servers, applications, and users. ESG calls this negligible latency "processing I/O requests at wire speed."

Application specific integrated circuits or "ASICs" embedded in an intelligent switch perform I/O redirection at extremely high speeds. Virtualization solutions employing a split-path architecture do not cache I/O requests and as a result, do not incur performance bottlenecks or data integrity risks.

Split-path architectures introduce negligible latency when performing I/O redirection. Because I/O redirection occurs at wire speed, split-path architectures acknowledge writes only after data has been written to an

attached storage system. This eliminates any risk of data loss associated with losing a virtualization system with cached data.

It is important that an intelligent storage network solution work with existing infrastructure within the data center and doesn't lock customers out of choices in the future.

Invista interoperates with a wide variety of storage systems, switches, and operating systems. Supported storage systems include systems from EMC (Symmetrix, CLARiiON), Hitachi (99XX), and HP (EVA). EMC has announced plans for the support of more storage systems. Invista is currently qualified to interoperate with intelligent switches from Brocade (AP7420) and Cisco (MDS 9000 series with Storage Service Module [SSM]). McData has also announced that Invista-qualified switches are scheduled to ship later this year. Additionally, Invista has been qualified to work with servers running HP-UX, Linux, Solaris, and Windows. AIX and VMware ESX are scheduled for release in 2006. As shown in the following diagram, Invista software running within the storage network creates a single level of management where virtual volumes, data migrations, and copy services can be managed regardless of the number of servers, operating systems, or storage systems in the data center.

## ESG Lab Validation Analysis

ESG Lab did hands-on analysis of the Invista virtualization architecture and features in order to validate its value proposition to customers. This ESG Lab report focused on the following features and functions of the EMC Invista:

- Ease of configuration and management
- Heterogeneous LUN virtualization
- Non-disruptive data migration
- Full volume on-line copy services

The goal of ESG Lab testing of the EMC Invista virtualization platform was to focus on fundamentals, including ease of use and non-disruptive virtualization, data mobility, and copy services. ESG Lab found that the Invista platform has an intuitive user interface that facilitates fast and non-disruptive virtualization, migration, and copy services. The testing was done at EMC headquarters in Hopkinton, MA.

Much of the testing was performed as the Windows server ran an assortment of applications, including Windows Media player video streams, file system level copies, and an Oracle application. Testing was performed on a pair of intelligent Cisco MDS 9216a Fibre Channel switches attached to

Symmetrix and CLARiiON storage systems. Invista was observed operating with storage arrays from Hitachi Data Systems and IBM and switches from Brocade and McData, as well.

Because EMC professional services performs initial installation and configuration services for Invista customers, ESG believes that a well planned Invista installation can be completed in a half-day or less. Planning may take longer than the actual installation.

ESG Lab testing began with a quick tour of the test bed shown in Figure Three, using the Invista Element Manager graphical user interface. An intuitive navigation tree which displays physical and logical resources was navigated using a Web browser.

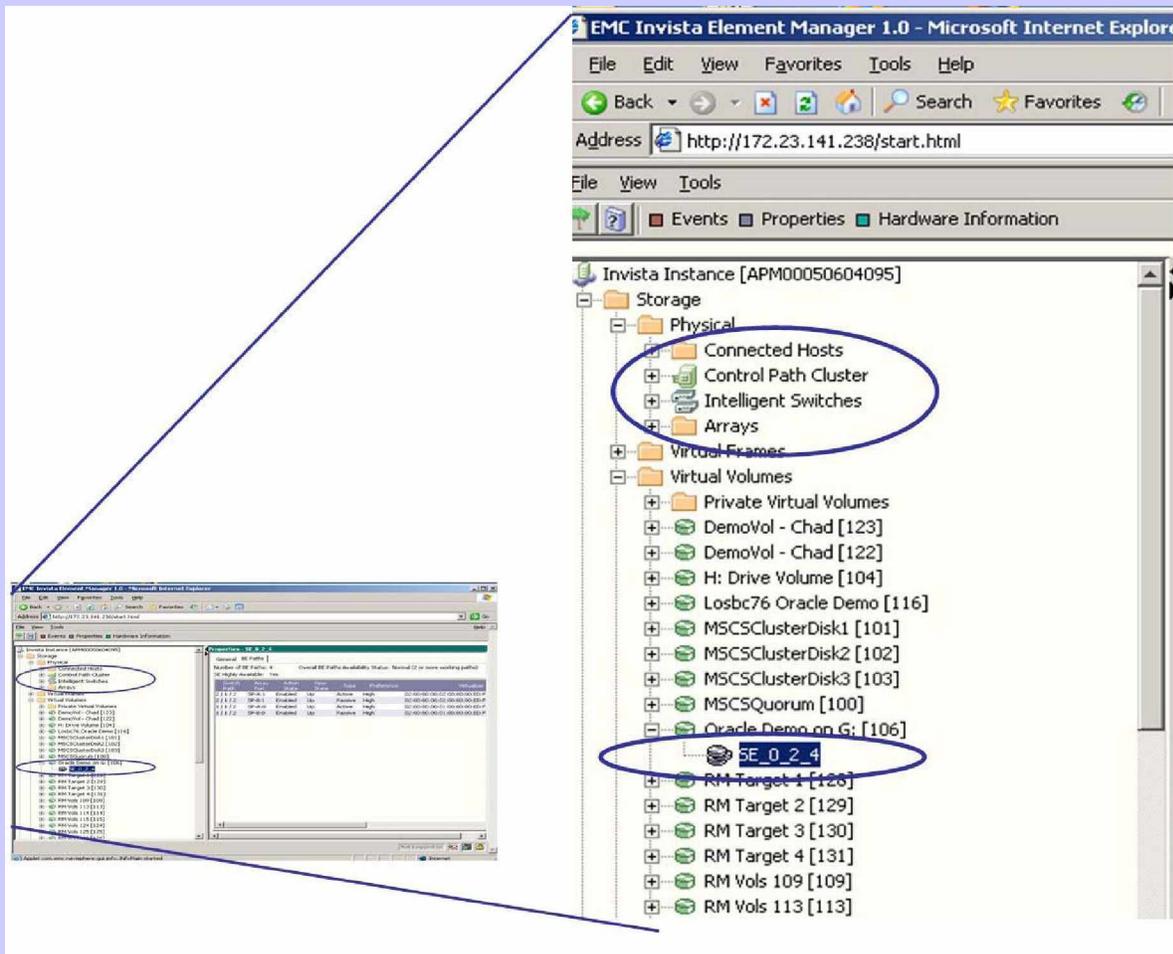


Figure 3: Invista GUI from a Browser Interface

On the left of Figure 3, note the physically Connected Hosts, Control Path Cluster, Intelligent Switches and Arrays. Physically connected resources which were already configured and virtualized by Invista were explored

using the Virtual Volumes view. As shown here, the physical properties of each virtual volume (in this case, the Oracle demo volume G: volume) can be viewed on the right. The Oracle volume is actually accessed through four independent and fault tolerant Fibre Channel interfaces.

ESG Lab tests found the Invista GUI to be intuitive and easy to navigate. It is no small feat that Invista provides a single interface for the presentation of complex Fibre Channel concepts, including multi-path drivers, host bus adapters, switch zones, access control, and storage devices. EMC claims to have re-invented the product development cycle with this solution while investing in a long and controlled Beta program to refine usability based on extensive customer feedback. ESG believes that the fruits of that labor are obvious and that users familiar with Fibre Channel concepts will find the system easy to use.

## ESG Lab Validation Report: LUN Virtualization

Virtualization is not a new concept; we have all used Virtual memory and telephone virtual circuits. Some have used Virtual networks and Virtual Machines in their work.

Like virtual memory and virtual machines, virtualization of physical storage provides an abstraction layer between physical and logical resources. Virtualization creates a single layer of management that makes it possible to add, change, and move physical resources while minimizing the impact on applications. To applications and users, a virtual storage device looks and feels just like a physical device.

An Invista virtualized volume can be made up of one or more physical volumes using storage from one or more systems and vendors. Capacity can be added to an existing Invista volume if more space is needed for an application. An Invista volume can be moved or copied while applications are running.

**A LUN is a physical "chunk" of storage within a Fibre Channel storage array.** A LUN is often presented to applications and users as a logical drive (for example, an F: drive on a Windows server). Invista implements LUN virtualization within the Storage Area Network (SAN) between servers and storage.

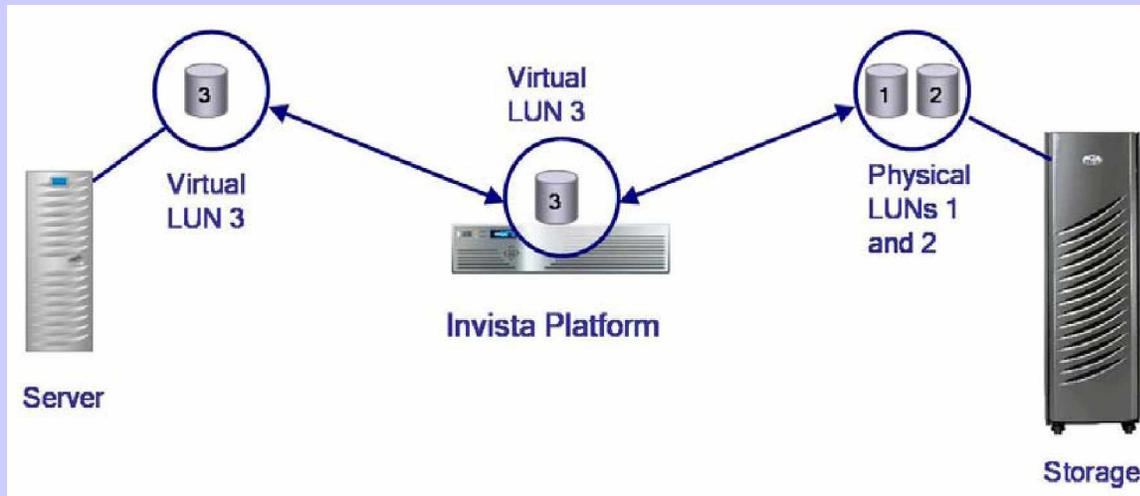


Figure Four: Invista LUN Virtualization

Invista software running on an Invista hardware platform in the storage network works in concert with an intelligent switch to provide virtualization services.

As shown in Figure 4, applications see LUNs that are virtualized by Invista using one or more physical LUNs presented by a storage array. In this case, an application server sees an Invista virtualized LUN 3 which is physically located on LUNs 1 and 2 within a storage controller. The server would see the virtualized drive as a normal part of the file system; e.g., a Microsoft server would see this Virtual LUN 3 as a mapped disk drive.

## ESG Lab Tested Invista Virtual Volumes

ESG Lab created an Invista virtual volume and then used a Windows server to store and access a movie located on that virtual volume. Three configuration operations were performed using the GUI to create and provide access to the virtual volume:

1. Import a Storage Element
2. Create a Virtual Volume
3. Provision the Virtual Volume to a Virtual Frame

Configuration operations were initiated using an intuitive right click on elements in the navigation menu on the left. For example, the following screen shot was taken as a 2 GB virtual volume was being created after performing a right click on a storage element to be virtualized.

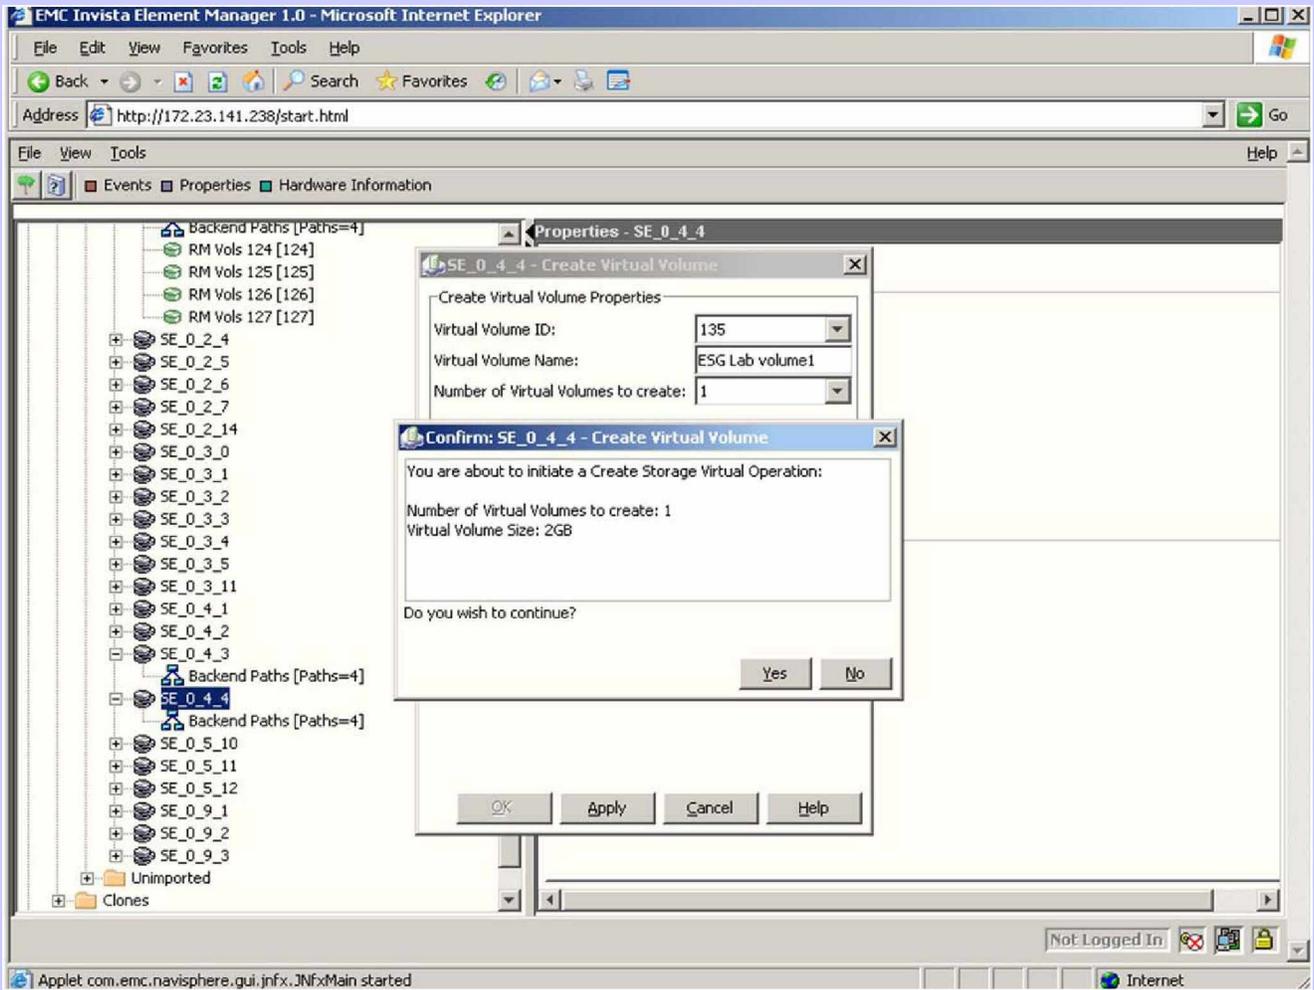


Figure Five: Creating a Virtual LUN

Microsoft server administrators often use the "Disk Administrator" to manage and change drive parameters. A Windows disk administrator rescan option was used to find the new device and format it as an F: drive. Discovering and managing the device using disk administrator was a transparent process and felt exactly the same as if Invista were not there. A right click on the F: drive in disk administrator and the hardware properties tab revealed that the device was being presented to Windows as an "EMC Invista SCSI Device" instead of an array specific device name (e.g. "EMC Symmetrix").

In other words, once a virtual device has been configured using Invista, administration of the storage on the server was exactly the same as if Invista were not present.

## How the ESG Tested for Data Mobility

Data mobility is a service provided by the Invista platform that is often referred to in the industry as data migration. A data migration is the movement of storage currently in use by an application to a new set of storage devices, often on a different storage array. Data migrations occur when organizations add arrays to the storage network to satisfy new application demands, because of lease expirations, and for a variety of other reasons.

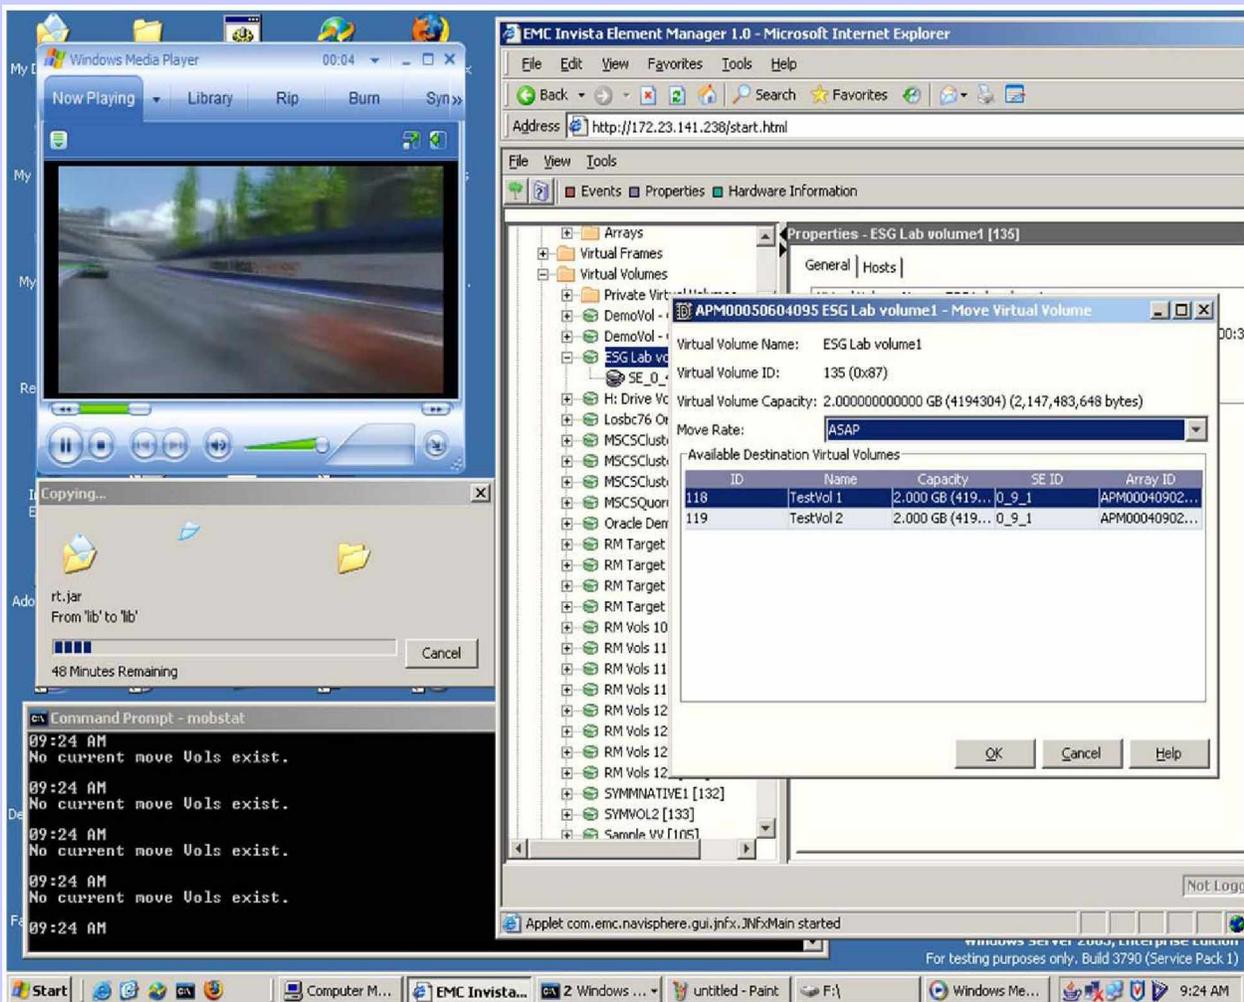


Figure Six: Invista Data Mobility

Data migrations happen frequently across organizations of all sizes and platforms. Applications must be shut down as data is moved using traditional data migration methods. Data is often moved from one system to another using backup tapes or a network. Regardless of how the data is moved, the migrations require extensive downtime, consume significant resources, and are prone to error.

The Invista platform provides data mobility services that enable users to change the physical storage behind a virtual Invista volume while applications are running. Dedicated hardware within a SAN switch is used to move data under the direction of Invista software. Invista performs data migrations with no application downtime.

ESG Lab performed a series of tests to learn how to configure Invista Data Mobility services, to verify that applications run without interruption as data is moved, and to measure the rate at which data can be moved.

The first migration was of a Windows volume running on a Symmetrix storage system to a CLARiiON storage system. As shown in the following diagram, Windows media player was used to view a movie and the contents of the Program Files directory were duplicated as the migration ran. Using the Invista GUI to configure and execute the move was easy and intuitive. Right clicking on an existing virtual volume and selecting the move option revealed the panel shown on the right in the previous screen shot. Note that Invista automatically found and listed available volumes with sufficient available capacity. A move rate of ASAP was selected to get the migration done as quickly as possible. Additional move rates of low, medium, and high, which can be used to minimize the impact on running applications, were noted. The migration completed in less than a minute as the movie and the file copy ran without error.

A second migration was performed as an Oracle 9i batch job ran in a loop while watching a movie. The Oracle batch job added, selected, and removed rows from a test database. The migration completed in just over a minute as the movie and Oracle ran without error.

A third migration was performed using a traditional server-based copy method in order to compare this more traditional method with Invista. The scenario was designed to mimic the process used during a storage upgrade of an existing production Oracle application. To avoid lost database updates, the Oracle application had to be shut down before the migration began. The contents of a 2 GB Oracle volume residing on a CLARiiON volume were copied to a Symmetrix volume using a Windows copy command. The copy completed in 25 minutes. The following table summarizes the results of ESG Lab Data Mobility testing.

Method /Application	Time to Migrate	MB/sec	GB/hour	Down Time
Invista Movie Copy	40 seconds	50	180	None
Invista Movie Oracle	70 seconds	29	102	None
Traditional N/A	25 minutes	1	4	25+ Minutes

Table One: Migration Method Results

## What the Numbers Mean

- Invista migration rates of 180 and 102 GB/hour were observed.
- Actual migration rates surpassed the conservative 60 GB/hour guidance from EMC.
- Invista migrations ran almost 50 times faster than a traditional migration using a server-based copy.
- Invista Data Mobility eliminated application downtime.

## Why This Matters

Data migrations using traditional methods incur application downtime and can have a significant business impact. ESG has found that most customers can tolerate very little downtime without impacting revenue.

ESG Research indicates that the majority of respondents could not tolerate more than one to four hours of downtime before experiencing a significant impact to their businesses (23 percent).

Nearly as many said that they could not tolerate more than an hour (22 percent) and many couldn't tolerate any downtime at all (14 percent)

ESG Lab has verified that Invista performs data migrations with no application downtime.

## Copy Services and How ESG Tested Them with Invista

Invista copy services are very similar to data mobility services – both make copies of a volume while applications are running. The difference can be seen in what happens when the job is done. When the data mobility duplication process has completed, users and applications are moved to the new volume.

Copy services creates one or more extra versions of a volume which can be used by another application. Copy services are typically used to duplicate

volumes for use during application testing or data mining. Creating a second copy for application testing reduces the risk of bugs and data corruption. Creating a second copy for data mining moves resource-intensive tasks, like long-running database queries and reports, off of mission critical servers, host bus adapters, storage switches, and storage systems.

Legacy copy services methods include restore from tape and server-based copies over the network. An alternative to legacy methods is storage system-based copy services. Most primary storage systems provide internal copy-based services and some provide remote copy services that replicate data between storage systems. Operation of storage system-based copy services is vendor specific and services between storage systems from different vendors are rare. Invista provides a single level of management for copy services between storage systems from a variety of vendors.

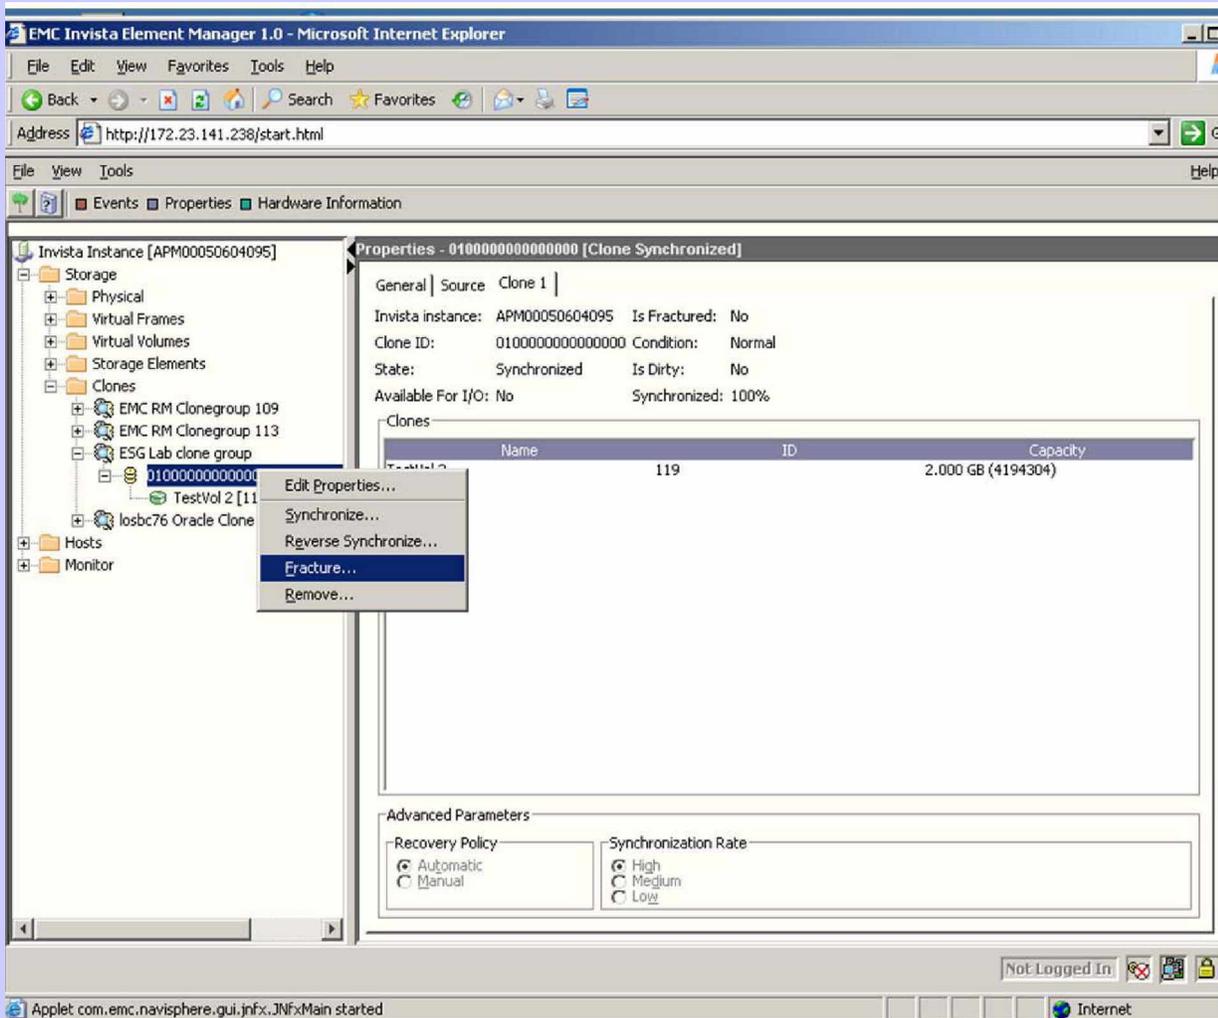


Figure Eight: Fracturing an Invista Clone

ESG Lab verified that a Windows volume cloned using Invista Copy Services could be accessed from a second Windows server. The Invista GUI was used to configure and execute the copy. A four step process was executed: a clone group was created, a clone was added using the start synch option, the clone was fractured (screen shot above), and the clone was assigned to a server.

Two minutes after the cloning process had begun, a rescan option from the Windows disk administrator utility was used on the second windows server to recognize and verify the contents of the cloned volume.

## Why This Matters

Data center managers have told ESG that copies of production storage volumes are often needed for application testing and data mining. Making copies of production volumes reduces the load on mission critical servers and storage systems so that primary applications can run at full speed.

Running in the network and connected to a heterogeneous mix of storage devices, Invista provides a single user interface for the creation of copies within and between storage systems from multiple vendors. Compared to legacy copy methods, Invista copy services reduce the impact on servers and storage systems.

## ESG Lab Validation Highlights

1. Virtualized volumes were created, moved, and cloned during ESG Lab testing.
2. Testing was performed on an intelligent Cisco MDS 9216a Fibre Channel switch attached to Symmetrix and CLARiiON storage systems. Invista was observed operating with storage arrays from Hitachi Data Systems and IBM and switches from Brocade and McData.
3. ESG Lab found the Invista GUI to be intuitive and easy to navigate for users familiar with Fibre Channel concepts.
4. Invista Data Mobility services were used to perform online migrations. Oracle, Windows Media player and Windows-based copies ran without error.
5. With applications active, ESG Lab observed data migration rates of 102 GB/hour and 180 GB/hour, which exceed EMC conservative guidance of 60 GB/hour.
6. Invista data migrations ran up to 50 times faster than a server-based copy.

7. A legacy data migration using a server-based copy caused 25 minutes of application downtime. Invista incurred zero downtime.
8. Invista copy services were used to create and access a copy of an existing volume on a Windows server from a second Windows server. The elapsed time from Copy Services setup to virtual volume access on the second Windows server was two minutes.

## ESG Lab's View

EMC opted to wait for fast and intelligent platforms from major switch vendors to mature before providing data services in the storage fabric. EMC leveraged the switch platform for a variety of reasons.

**First**, EMC believes that split path architecture with dedicated switch hardware for data movement is architecturally superior to an in-band appliance-based solution.

**Second**, a heterogeneous virtualization solution that mixes storage, an intelligent control path appliance, and switches, all from different vendors, needs to be thoroughly tested for interoperability to meet the reliability and serviceability needs of enterprise-class customers.

**And finally**, EMC re-invented the product development cycle with this solution and invested in a long and controlled Beta program to refine features, usability, and reliability, based on extensive customer feedback.

ESG research demonstrates that the benefits of an intelligent storage network solution are real, and that early adopters have achieved significant savings. ESG Lab validated that the Invista platform has an intuitive user interface and can be used to perform fast and non-disruptive virtualization, migration, and copy services.

The storage virtualization market is still emerging and with EMC now in the game, ESG expects that storage virtualization will gain even more momentum. The EMC Invista exceeded our expectations, providing an easy to use, high performance solution with a useful feature set.