

# IBM XIV with Real-time Compression feature in VMware vSphere environment

*A technical report*

*Mandar J. Vaidya*  
*IBM Systems ISV Enablement*

*January 2016*

 @IBMSystemsISVs



## Table of contents

<b>Abstract</b> .....	<b>3</b>
Scope .....	3
<b>Introduction</b> .....	<b>3</b>
<b>Prerequisites</b> .....	<b>4</b>
<b>Introduction to IBM XIV Storage System</b> .....	<b>4</b>
<b>Introducing IBM XIV Model 314</b> .....	<b>4</b>
<b>Introduction to VMware vSphere</b> .....	<b>5</b>
<b>IBM Real-time Compression</b> .....	<b>6</b>
IBM Real-time Compression with XIV: Technology overview.....	6
Architectural overview.....	7
Planning for Real-time Compression in XIV system.....	8
Estimating compression savings in XIV.....	9
<b>Implementing real-time compressed volumes on IBM XIV Storage System</b> .....	<b>10</b>
Enabling Real-time Compression feature in XIV .....	10
Creating compressed storage pool.....	11
Creating a compressed volume .....	12
Converting an existing uncompressed volume to a compressed volume .....	15
Configuring XIV mirroring with compressed volume.....	16
<b>Benefits of XIV with real-time compressed volumes in a VMware virtualized environment</b> .....	<b>17</b>
VMware vSphere 6.0 virtual machines on a compressed XIV volume .....	17
Performance comparison between compressed volumes on XIV Model 314 and non-compressed volumes on XIV Model 214.....	19
Application score comparison between XIV 314 and 214 models .....	20
Infrastructure score comparison between XIV 314 and 214 models .....	21
VMmark total score comparison between XIV 314 and 214 models .....	21
<b>Summary</b> .....	<b>23</b>
<b>Acknowledgements</b> .....	<b>23</b>
<b>Appendix A: Test environment</b> .....	<b>24</b>
<b>Appendix B: Resources</b> .....	<b>25</b>
<b>About the author</b> .....	<b>25</b>
<b>Trademarks and special notices</b> .....	<b>26</b>

## Abstract

---

*This technical paper demonstrates how the IBM XIV Storage System, with the IBM Real-time Compression feature enabled, can reduce the storage capacity needed for the primary data and provides enterprise-grade performance, without any relative notion. Real-time Compression can be transparently used with VMware vSphere to effectively decrease storage consumption of VMware virtual machines and the application data they contain.*

*This paper introduces the integration of Real-time Compression technology with XIV grid architecture, outlines the compression benefits in a virtualized environment, and describes the deployment strategy of Real-time Compression feature in IBM XIV with VMware vSphere.*

*The target audience for this paper is system and storage administrators looking to implement the XIV with Real-time Compression feature in VMware vSphere environments*

## Scope

This technical report:

- Discusses the IBM® Real-time Compression™ feature with IBM® XIV® Storage System and its configuration in VMware vSphere environments.
- Discusses the benefits of the Real-time Compression feature with IBM XIV Storage System.

This technical report does not:

- Discuss the configuration of the IBM XIV Storage System and VMware vSphere.

## Introduction

---

Industries continue to demand effective utilization of IT infrastructure including servers and storage without affecting the performance of the application. Server virtualization provides many benefits of using server resources effectively and this leads to its explosive adoption in today's data center. Many customers are taking benefits of server virtualization to effectively use server hardware resources, reducing management tasks.

IBM XIV Storage System, with its next-generation grid architecture, automated load balancing, and exceptional ease of management provides best-in-class enterprise storage for virtual servers. IBM XIV provides end-to-end support for VMware solutions, including vSphere and vCenter and provides optimal storage performance and resource use.

Many organizations are facing enormous challenges in addressing primary storage growth where primary storage is growing by 30% to 60% and many consumers prefer data primarily on disks instead of offloading it to tapes for faster access\*. IBM XIV Storage System introduces Real-time Compression feature in 11.6.0 release. IBM Real-time Compression uses IBM Random Access Compression Engine (RACE) technology which is seamlessly integrated with its storage system software stack to compress data before writing it to disk resulting in up to 80% storage capacity savings. IBM XIV Storage System with Real-time Compression feature increases the usable capacity per rack typically to one petabyte or more reducing the overall cost per capacity without additional hardware.

---

\* Source - IT Market Clock for Storage - Gartner report



The solution presented in this paper builds upon server virtualization from VMware and storage virtualization provided by XIV Storage System. Real-time compressed volumes are implemented for VMware vSphere virtual machine data, including operating system (boot disk) and application data.

## Prerequisites

---

This paper provides deployment and management guidelines for customers who are planning or have already decided to implement VMware vSphere on the IBM XIV system using Real-time Compression.

### Technology skills prerequisites

- VMware vCenter Server management and administrator
- VMware ESXi installation
- VMware Virtual Machine File System (VMFS)
- Guest operating system installations in a VMware environment
- Storage subsystems and terminology

## Introduction to IBM XIV Storage System

---

IBM XIV is a next-generation high-end open disk system. An XIV system has a grid architecture designed for enterprise reliability, performance, functionality, and scalability at a lower overall cost of ownership, while eliminating the complexity and providing ease of management.

The XIV system is ideal for cloud and analytics, offering high service levels for dynamic workloads, easy hyper-scaling including in multi-tenant environments, flexible consumption models, and robust integration with cloud orchestration platforms through OpenStack, RESTful application programming interface (API), VMware, and other means. The modular nature of XIV enables transparent and scalable storage providing on-demand capacity growth with little or no storage administration and without the need of manual data migration or performance tuning.

The IBM XIV system offers the following notable benefits:

- Grid architecture and data distribution approach provides tuning-free high performance.
- Exceptional availability using three-way mirror capability, self-healing with the capability to rebuild in less than an hour and five-nine design.
- Ease of management and provisioning, using simplified user interface.
- Elastic and robust multi-tenancy storage system designed for the needs of cloud environment.
- Environmental friendly storage system with higher effective capacity per rack resulting into lower cost per terabyte.

For additional information about the IBM XIV Storage System, refer to the following URL:  
[ibm.com/systems/storage/disk/xiv/index.html](http://ibm.com/systems/storage/disk/xiv/index.html)

## Introducing IBM XIV Model 314

---

In addition to all the benefits of XIV Storage System, IBM has announced XIV Model 314 which increases the high utilization and power efficiencies of XIV 4 TB and 6 TB disk drives, delivering outstanding data economics to your high-end storage. With double the RAM and processor resources, IBM XIV Storage



System Gen3 Model 314 delivers improved input/output operations per second (IOPS) per compressed capacity and effective capacity up to 2 petabytes without performance degradation.

IBM XIV Storage System Gen3 Model 314 delivers:

- Two 6-core processor per module (versus one 6-core processor per module in Model 214)
  - One 6-core processor is dedicated to Real-time-Compression
- 96 GB RAM per module (versus 48 GB RAM per module in Model 214) - 48 GB of RAM is dedicated to Real-time-Compression
- 1 to 2 PB of effective capacity without performance degradation
- Improved IOPS per compressed capacity
- User-configurable soft capacity up to 2 PB
- Reduced minimum compressible volume size from 103 GB to 51 GB

The additional processor and 48 GB RAM are dedicated to Real-time Compression. At double the compression ratio, the effective compressed read/write cache capacity in a 15-module XIV system is doubled to 1440 GB, delivering a total effective RAM of 2160 GB. The solid-state drive (SSD) read-cache capacity per system is also doubled to 24 TB. For more details, refer XIV Model 314 specifications at: [ibm.com/systems/storage/disk/xiv/specifications.html](http://ibm.com/systems/storage/disk/xiv/specifications.html)

## Introduction to VMware vSphere

---

VMware vSphere is a virtualization platform capable of transforming a traditional data center, consisting of industry-standard hardware, into a shared mainframe-like environment. Through the use of virtualization technology, hardware resources are abstracted into pools which might be allocated to a variety of workloads.

VMware vSphere empowers users to virtualize any application with confidence, redefines availability and simplifies the virtual data center. It provides highly available, resilient, on-demand infrastructure that is the ideal foundation of any cloud environment. This can drive down data center cost, increase system and application uptime, and drastically simplify the way IT runs the data center. VMware vSphere is purpose-built for the next generation of applications and serves as the core foundational building block for the software-defined data center.

VMware vSphere provides following benefits:

- Powerful server virtualization - It virtualizes the server resources and aggregate them into logical pools for allocation of multiple workloads. It improves hardware utilization to as much as 80 percent or more without sacrificing performance.
- Efficient storage integration - It enables external storage arrays to operate in a more VM-centric manner by increasing efficiency and performance of virtual machine operations.
- Maximum uptime and high availability – It enables to maximize uptime across your cloud infrastructure, reducing unplanned downtime and eliminating planned downtime for server and storage maintenance.
- Simplified management – It provides a simplified and powerful management interface to manage the creation, sharing, deployment, and migration of virtual machines.

- Lower IT costs - It reduces capital expenditures and operational expenditures by providing powerful virtualization platform.
- Network services – It provides optimized network services along with simplified administration and management for virtual environment.
- Automation - Provides automation framework to streamline routine tasks.
- Intelligent operations management - It adapts to your specific environment, giving you better insights in time to take proactive action.
- Enhanced security - It provides secure server virtualization platform to protect data and applications.

For additional information and details about VMware vSphere, refer to the following URL:  
<http://www.vmware.com/products/vsphere>

## IBM Real-time Compression

---

IBM Real-time Compression technology takes storage efficiency to a new level using state-of-art technology, which is designed to improve efficiency by compressing data up to 80%. It enables users to store data up to five times as much data in the same physical disks with minimal impact on performance. IBM Real-time Compression is designed to be used with an active primary data such as production databases and email applications.

IBM Real-time Compression uses the industry-standard Lempel-Ziv (LZ) compression algorithms. It uses a technique called *random access* to read or write specific blocks of compressed data. This ensures that only blocks that require read or modification are accessed, and updates to data can be completed with no disruptions to other blocks in a compressed volume.

### IBM Real-time Compression with XIV: Technology overview

IBM Real-time Compression implementation in XIV Storage System is based on RACE technology. It compresses the data that is written into the storage dynamically. The compression occurs transparently to the hosts and applications and each write is compressed as it passes through Storage System to the disks. Implementation of IBM Real-time Compression provides the following benefits:

- It provides compression for the primary active data and supports inline compression of data.
- Implementation supports online conversion of uncompressed volumes to compressed volumes without the need of cleaning or archiving of existing data which helps administrators to reclaim the space while system is online.
- No change to the existing environment is required. Compression is in-built feature of XIV system and can be implemented without changes to application, hosts, and fabric networks.
- Real-time Compression implementation in XIV supports mirroring of compressed volumes. If source volume is compressed, remote copies are always compressed. It reduces bandwidth utilization by sending less (compressed) data across the network.

- Compression can result in reduction in cost per capacity as it provides increased usable capacity per rack. More data can be stored in less rack space reducing power and cooling requirements and thus resulting in the reduction of overall operational expenses.

### Architectural overview

Real-time Compression implementation in XIV Storage System uses *above cache* architecture where data is compressed or decompressed between the I/O interface and the cache. The compression runs on every module of XIV taking advantage of parallel architecture of XIV. It compresses the portion of volume which only belongs to the module and thus distributing compression workload across all the modules of XIV. Hence, Real-time Compression implementation in XIV has minimized impact on the performance delivered by XIV.

Figure 1 shows a pictorial representation of IBM Real-time Compression architectural overview.

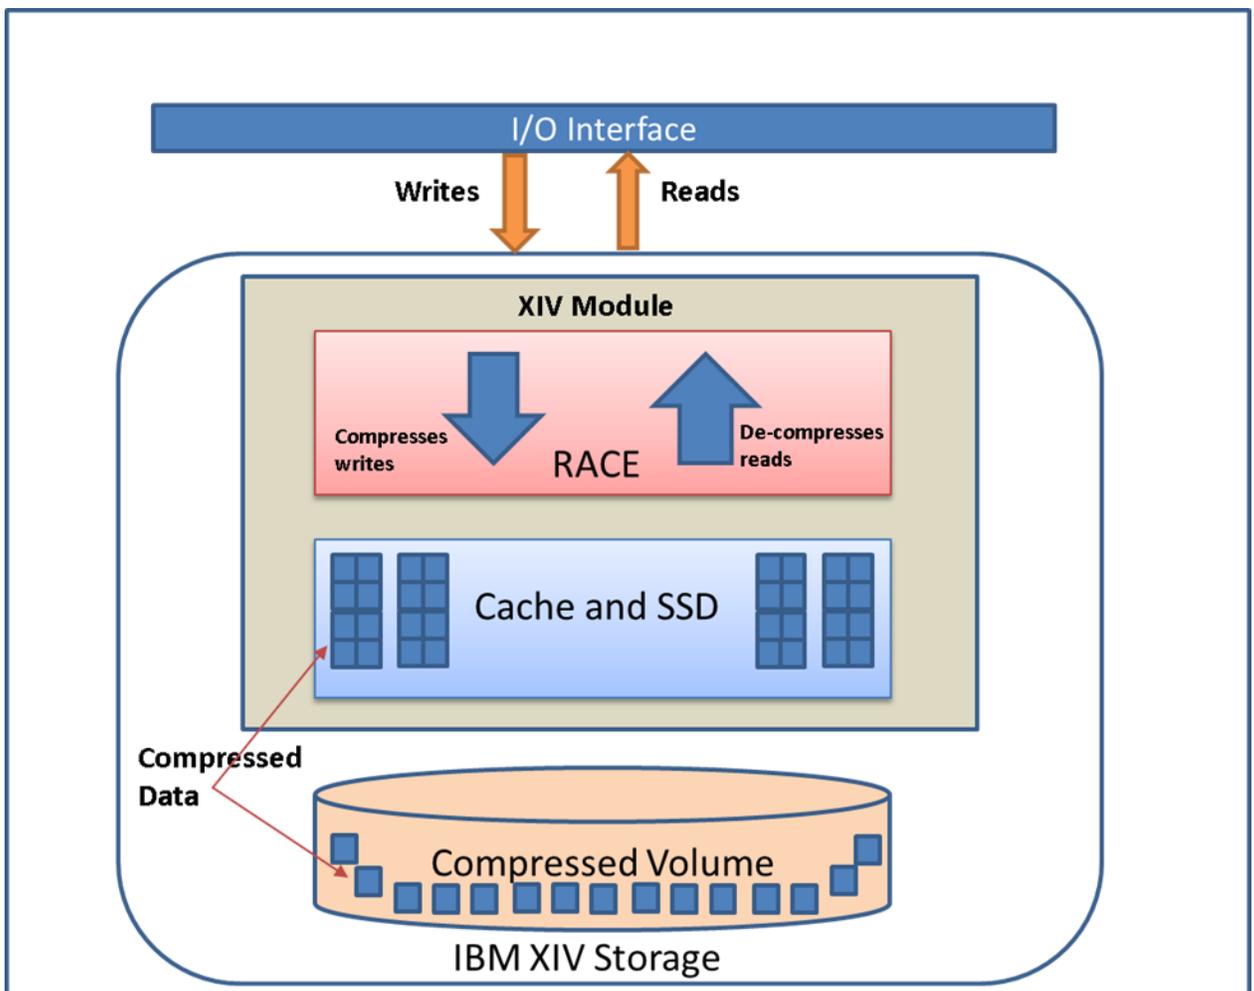
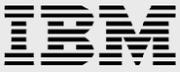


Figure 1: Overview of IBM Real-time Compression with XIV

Whenever write operations happen, data is compressed before it enters the cache. During the read operations, reads are stored compressed in cache and data is decompressed when it is read from cache using RACE, before passing it to the host.



During XIV mirroring operation on a compressed volume, data is compressed only at source and compressed data is sent across the network reducing network bandwidth.

## Planning for Real-time Compression in XIV system

Some applications workloads or data types can benefit more from Real-time Compression than others. Before implementing compression, it is important to analyze the type of data and volumes used on the XIV system. Real-time Compression is designed to work best with random access data such as:

- Database environments – IBM DB2®, Oracle, MS-SQL, and so on
- Database applications – SAP, Oracle applications, and so on
- Server/Desktop virtualization – KVM, VMware, Hyper-V, and so on
- Other compressible workloads – seismic, engineering, and so on
- Email – Microsoft® Exchange, and so on

In general, consider the following guidelines before implementing compression:

- Type of data where compression savings are expected more than 25%.
- Volumes which do not have already compressed data (such as compressed files, videos, and so on).
- Data where application-based encryption is not used.

Also, consider the following recommendations for implementing compression:

- For smaller number of compressed volumes, configure them in one I/O group.
- For larger number of compressed volumes with more than one I/O group, distribute compressed volumes across I/O groups.
- Pools should have extra 5% capacity to hold the additional metadata for the compressed volumes.

Compression requires dedicated hardware resources within the nodes which are assigned when compression is enabled. Compression is enabled when the first volume in an I/O group is created and is disabled when last compressed volume is removed from the I/O group.

Compression feature in XIV system is designed to work on thin volumes that are part of thin provisioned pools. Table 1 lists minimum and maximum volume sizes which can be configured in an XIV system.

	Compressed	Uncompressed
<b>Minimum volume size</b>	103 GB	17 GB
<b>Maximum volume size</b>	10 TB	161 TB

Table 1: Minimum and maximum volume sizes

To compress existing volumes smaller than 103 GB volumes, customers can resize the volume (to 103 GB or larger) using XIV graphical user interface (GUI) or command-line interface (CLI), before the conversion. To accommodate the larger volume, the related thin pool soft capacity may need to be enlarged too. Host

rescan may be required. Otherwise, there is no significant operational impact. The actual hard capacity needed for this volume will be reduced by compression, as expected.

## Estimating compression savings in XIV

On an XIV system supporting compression, the compression ratio for all uncompressed volumes in the system is continuously estimated, even before enabling compression. The decision to use compression can be based on the expected storage savings of the compressed data and the expected effect on performance (throughput and latency) of the compression processing overhead. Expected storage savings can vary from 5% higher or lower than the estimated compression ratio.

XIV offers CLI to estimate compression on an existing volume. By using this command-line utility, customers can determine whether the data can be a good candidate for Real-time Compression. In order to view estimated compression on an uncompressed volume, login to XIV XCLI, and run the command-line utility.

```
ISV XIV 3>>
ISV XIV 3>>
ISV XIV 3>>vol_comprestimate vol=MV_VMW_Thin_UnC_vol01
Command executed successfully.
ISV XIV 3>>
```

Figure 2: Command line utility to estimate compression

After the command is run, potential compression savings can be viewed from volume view.

SQL2012SP1VM01_data_vol_01	2,231	09	2,231 GB	4.3%	21 GB
SQL2012SP1VM01_Log_Vol_01	103	0	103 GB	56%	0 GB
MV_VMW_Compressed_vol03	258	4	258 GB	42%	3 GB
MV_VMW_Thin_UnC_vol01	154	7	154 GB	48% Potential saving	0 GB

Figure 3: Potential compression savings display from Volume view

Potential compression savings can also be viewed in detail by right-clicking on the volume and selecting **Properties** of a volume.



Figure 4: Volume properties to display estimated compression savings

## Implementing real-time compressed volumes on IBM XIV Storage System

After Real-time Compression feature is enabled, it can be used to create compressed volumes. Enabling compression on XIV does not affect non-compressed volumes. Use the following guidelines before working with compressed volumes:

- *Real capacity* is the extent space that is allocated from the storage pool. The real capacity is also set when the volume is created and, similar to thin-provisioned volumes, can be expanded or shrunk down to the used capacity.
- *Virtual capacity* is available to hosts. The virtual capacity is set when the volume is created and can be expanded or shrunk afterward.
- *Used capacity* is the amount of real capacity used to store customer data and metadata after compression.
  - **Compression savings** is the amount of capacity (in percentage) that is saved by using compression.
  - **Real** is the amount of space from the storage pool allocated to allow the volume data to be stored.
  - **Total** is the capacity of the volume that is available to hosts.

### Enabling Real-time Compression feature in XIV

Before using Real-time Compression feature for the first time in XIV, it needs to be enabled. Perform the following steps to enable Real-time Compression feature.

1. Log in to XIV management GUI and click Settings.

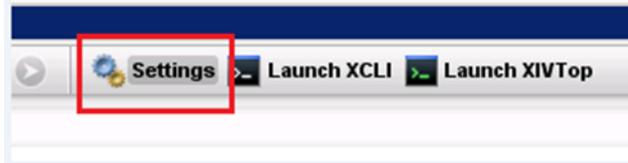


Figure 5: Enabling Real-time Compression

2. In the **Settings** window, click **Parameters** and then select **Enabled** from the **Compression Capabilities** drop-down list. Settings wizard will request licensing confirmation from the user. Click **Update** to enable the compression feature.

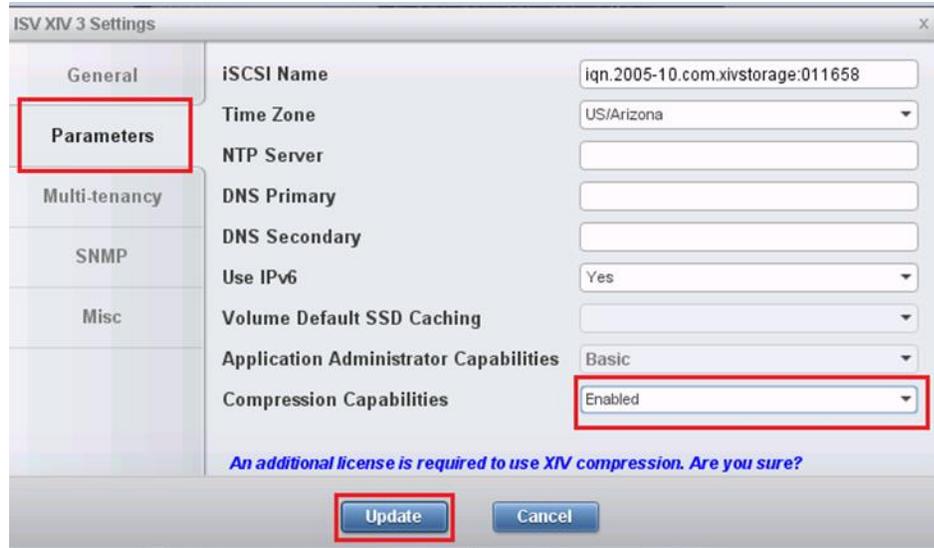


Figure 6: Enabling Real-time Compression inside settings

## Creating compressed storage pool

The compression feature in the XIV system is designed to work on thin volumes that are part of a thin provisioned pool. Perform the following steps to create compressed storage pools on XIV.

1. Log in to XIV management GUI and make sure compression is enabled on an XIV system. Click **Pools** → **Storage Pools** and then click **Create Pool**.



Figure 7: Selecting the options to create a pool

- In **Create Pool** wizard, select the pool type as **Thin Pool**, specify the pool soft and hard size, snapshots size and pool name. Select the **Compression Default** check box and click **Create** to create a compressed storage pool.

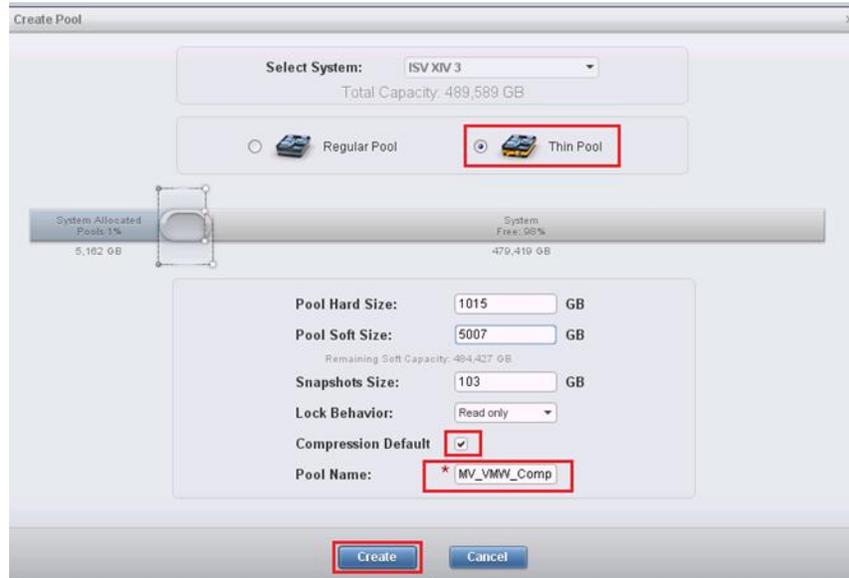


Figure 8: Create pool wizard

## Creating a compressed volume

After the compressed storage pool has been created, users can create a compressed volume using the compressed storage pools.

- To create a compressed volume on XIV, right-click the compressed storage pool and click **Create Volumes**.

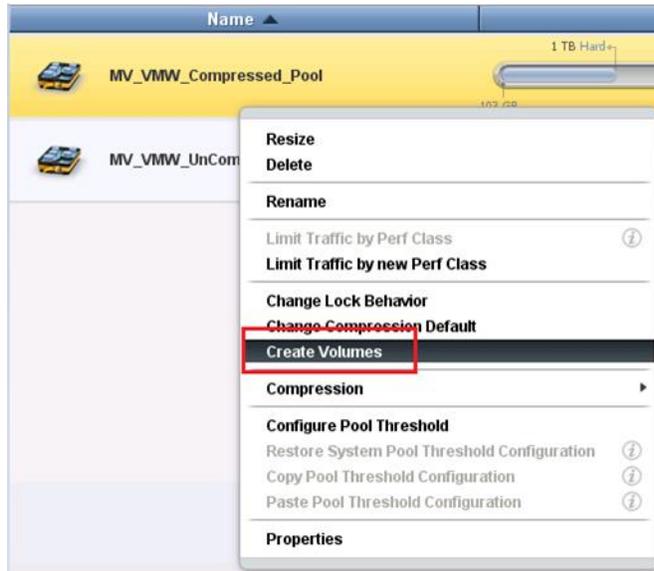


Figure 9: Invoking the Create Volumes wizard

- In **Create Volumes** wizard, select storage pool, specify the volume size and the volume name and select the volume type as **Compressed** by selecting the check box and then click **Create** to create a new compressed volume.

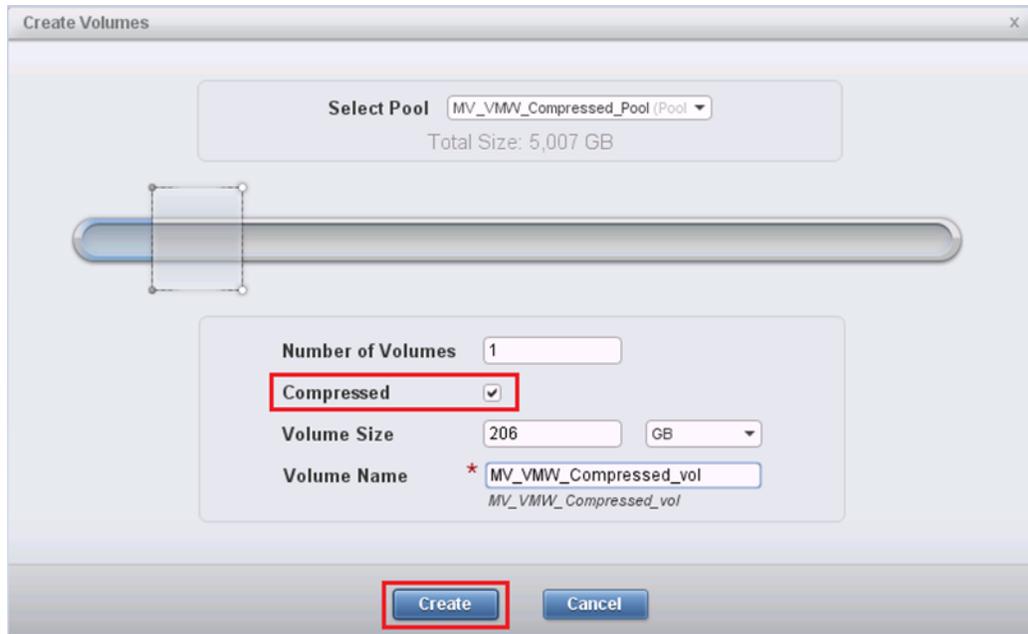


Figure 10: Creating volume wizard

- Newly created volume is displayed under Volumes and Snapshots listing. User can customize volume listing by selecting the table columns appropriately to display compression savings, used space, and compression savings (in percentage).

Name	Size (GB)	Used (GB)	Size (Disk)	Compression Saving (%)	Compression Saving
MV_VMW_Compressed_vol	206	0	206 GB	0%	0 GB

Figure 11: Compressed volume display

- XIV system will start inline compression of data as it is written to the volume. Compression savings will be immediately reflected on the XIV system.

systems (2) > ISV XIV 3 > Volumes and Snapshots > Volume (1), Snapshot (0)

Name	Size (GB)	Used (GB)	Size (Disk)	Compression Saving (%)	Compression Saving
MV_VMW_Compressed_vol	206	4	206 GB	42%	3 GB

The Used space of the compressed volume is 4 GB. You saved 3 GB.

Figure 12: Compression savings display in volume listing

- Detailed volume properties can be seen by right-clicking the volume and clicking **Properties**. Details about compression saving and compression ratio are displayed in the Volume Properties dialog box.

Volume Properties

System: ISV XIV 3

Name: MV\_VMW\_Compressed\_vol

Size: 206 GB (403,316,736 Blocks)

Used Capacity: 4 GB (4,043 MiB)

Size on Disk: 206 GB

Compression Status: **Compressed**

Compression Saving: 3 GB (2,977 MiB)

Compression Ratio: 42 %

Snapshots Used Capacity: 0 GB (0 MiB)

Serial Number: 156

Consistency Group: None

Pool: MV\_VMW\_Compressed\_Pool

Locked Status: No

WWN: 001738002D8A009C

ID: 89e14d0009c

OK

Figure 13: Displaying detailed volume properties

## Converting an existing uncompressed volume to a compressed volume

XIV supports converting of an existing uncompressed volume to a compressed volume without the need of archiving the data and while data is online. This process is completely transparent to the VMware environment and can be performed while ESX host is accessing the volume. It helps storage administrator to reclaim the space while data is online.

Only thin uncompressed volumes can be converted to compressed volumes. Before converting the volume, understand the data stored in the volume, its compressibility and estimate compression savings as mentioned in the *Estimating compression savings in XIV* section.

Perform the following steps to convert existing uncompressed volume to compressed volumes:

- Right-click the uncompressed volume which needs to be converted and click **Compression** and then **Compress**.

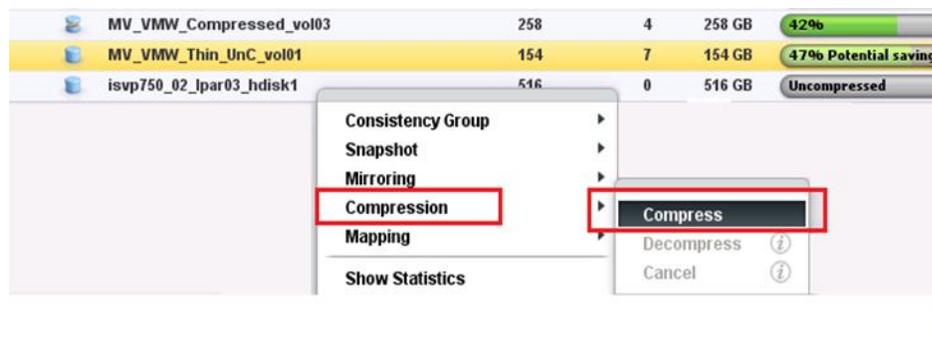


Figure 14: Converting an uncompressed volume to a compressed volume

- Compress wizard will ask for user confirmation and select whether you would like to keep source volume and click **OK** to convert.



Figure 15: Compress volume wizard confirmation

- After the volume is converted to the compressed volume, compressed volume along with the original volume (if chosen to be kept) is listed under volume display.

Volume Name	Size (GB)	Compression	Potential Saving	Original Volume
MV_VMW_Compressed_vol03	258	4	258 GB 42%	3 GB
MV_VMW_Thin_UnC_vol01.original	154	7	154 GB 47% Potential saving	0 GB
isvp750_02_lpar03_hdisk1	516	0	516 GB Uncompressed	0 GB
MV_VMW_Thin_UnC_vol01	154	4	154 GB 44%	3 GB

Figure 16: Compressed volume display along with the original volume

## Configuring XIV mirroring with compressed volume

XIV Storage System supports synchronous and asynchronous remote mirroring for disaster recovery. Remote mirroring can be used to replicate the data between geographically remote sites. Real-time Compression in XIV supports mirroring for compressed volumes. While configuring remote mirroring, if source volume is compressed then remote copies are always compressed. It helps in transmitting less (compressed) data across the network and reduces bandwidth utilization. Mirroring of compressed volume is supported on XIV version 11.6.0 onwards.

- Before configuring compressed mirror, establish a mirrored relationship between the XIV systems.



Figure 17: Establishing mirror relationship between XIV systems

- After XIV mirrored configuration is established, configure mirrored volumes across the XIV systems using the compressed pools.

Figure 18: Configuring mirrored volume

- After the mirrored volume is configured, compressed data is replicated to the target XIV system. Volume compression information is available on the source volume.

Name	Size (GB)	Used (GB)	Size (Disk)	Compression Saving (%)	Compression S...
HAFS_SCVMM_Lib_Vol_01	223	0	223 GB	Uncompressed	0 GB
LU1-3ce38621-xiv_clone_test_vm	120	10	120 GB	37% Potential saving	0 GB
MVVVol	2,013	0	2,013 GB	Uncompressed	0 GB
MV_VMW_Compressed_vol02	103	3	103 GB	Compressed	
Win2012R2_HA_T01_Vol_01	120	8	120 GB		
Win2012_AD_VM_Vol_01	120	9	120 GB		
Win2012_HA_VM_Vol_02	120	8	120 GB		

The compression information is not available on the destination volume in a mirrored/mobility relation. Check the source volume for the actual compression information.

Figure 19: Displaying mirrored compressed volume

## Benefits of XIV with real-time compressed volumes in a VMware virtualized environment

Implementation of VMware virtualized environment on XIV with Real-time Compression can take advantage of compression benefits while creating VMFS datastore on compressed XIV volume, converting uncompressed XIV volumes to the compressed volumes and also by replicating compressed VMware data across multiple XIV storages using XIV mirroring functionality.

For this technical paper, two workloads and their corresponding data were used to measure the effective compression in XIV with Real-time Compression feature enabled.

The workloads that were used for this lab validation are outlined in the following section.

### VMware vSphere 6.0 virtual machines on a compressed XIV volume

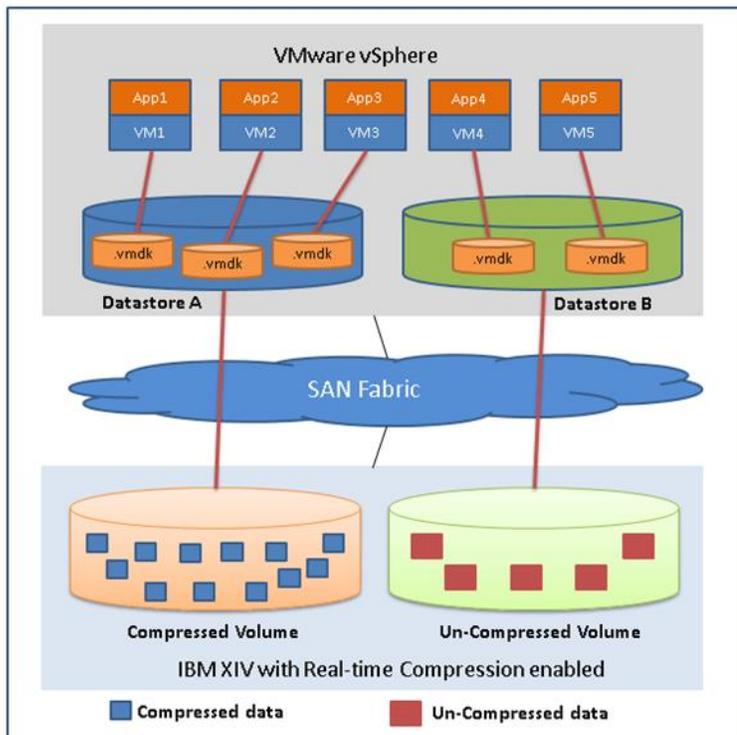


Figure 20: Virtual machines on a compressed XIV volume

IBM XIV with Real-time Compression feature in VMware vSphere environment  
© Copyright IBM Corporation, 2016



VMware vSphere 6.0 supports creating virtual machines using traditional VMFS datastore. When virtual machine is created using VMFS datastore in vSphere 6.0 environment, it is stored in a VMware VMDK file. VMDK files can be stored on a Network File System (NFS) mount point or within a VMFS datastore residing on a block storage volume such as those provided by XIV Storage System. Along with the operating system VMDK file, the minimum files for a virtual machine include a configuration (.vmx) file and various log files.

Real-time Compression is enabled at the XIV volume level. So, in a vSphere 6.0 environment, all virtual machines and data stored within a VMFS datastore that resides on a XIV volume is compressed. This includes operating system files, installed applications, and any other data.

Lab validation testing was performed to analyze the compressibility of virtual machines running two popular server operating systems. The details of the virtual machines are as follows:

- Microsoft Windows® 2008 R2 Enterprise Edition
  - Thin provisioned 50 GB VMDK
  - Base operating system files only
- Red Hat Enterprise Linux® 6.5 (64 bit)
  - Thin provisioned 16 GB VMDK
  - Base operating system files only

The test was performed by creating virtual machine using VMFS datastore created on a compressed XIV volume.

Operating system	Capacity used (in GB)	Compression savings (in GB)	Compression savings (in percentage)
Microsoft Windows 2008 R2	4	3	42 %
Red Hat Enterprise Linux 6.5 (64bit)	3.3	3.6	52 %

Table 2: Observed compression savings with VMware virtual machines

It is expected that many virtual machines would be deployed on a VMFS datastore. Second lab test was performed to measure the impact of deploying multiple virtual machines on a single compressed VMFS datastore. The same base virtual machine images were used from the previous test, and each virtual machine was cloned four times to create a total of 10 virtual machines on a single VMFS datastore.

Table 3 illustrates the measured compression savings by deploying the multiple virtual machines on a compressed VMFS datastore.

Virtual machines	Capacity used (in GB)	Compression savings (in GB)	Compression savings (in percentage)
Windows 2008 R2 and Red Hat Linux 6.5	26	26	50 %

Table 3: Observed compression savings by deploying multiple VMware virtual machines on a compressed datastore

This result observed during lab testing demonstrates that virtual machines running Windows and Red Hat Enterprise Linux are highly compressible, with an average compressibility of 50 %.

### Performance comparison between compressed volumes on XIV Model 314 and non-compressed volumes on XIV Model 214

VMware VMmark is a free tool from VMware to measure the performance and scalability of applications running in a virtualized environment. VMmark generates a realistic measure of platform performance by incorporating a variety of platform-level workloads in addition to traditional application-level workloads. It is designed to benchmark the performance of the virtualization software and hardware. VMmark calls a unit of work for a benchmark as a *tile*. A tile is a group of eight virtual machines concurrently running a collection of diverse workloads. Each of these workloads represents a common application workload found in today's data centers. Included in each tile are a mail server, a Web 2.0 database and web system, an e-commerce back-end database and front-end web layer, and an idle system. The total number of VMmark tiles that can be accommodated in a multi-host platform gives a coarse-grain measure of that platform's consolidation capacity. VMmark includes the following workloads.

- Mail server load simulation using Microsoft Exchange 2007 application
- Web 2.0 load simulation using Olio application
- E-commerce load simulation using the DVD Store 2 application

The workloads chosen for use in VMmark are a representation of the popular applications commonly run by VMware customers.

For more details on VMmark benchmarks, refer: <https://www.vmware.com/products/vmmark/>

Figure 21 shows the lab topology used for running VMmark benchmarks against XIV Models 214 and 314 to find out performance penalty, if any, by configuring compressed volumes on model 314. With XIV Model 214, non-compressed volumes were used. Compressed volumes were configured on XIV Model 314. Both models of XIV consist of 15 modules with each module containing twelve 4 TB disks and one 800 GB SSD.

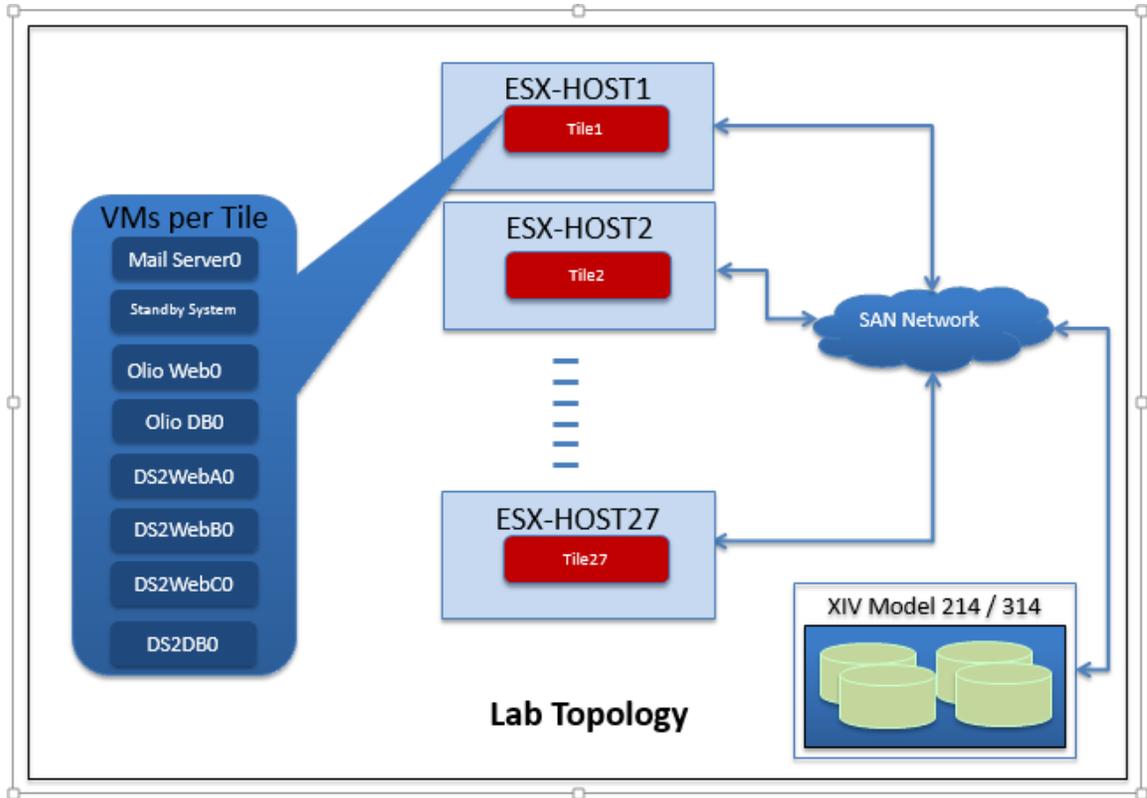


Figure 21: VMmark benchmark lab topology

For VMmark benchmarking workloads, each ESX server was configured to run one tile. Each tile had eight virtual machines. For each benchmarking run, 27 such tiles were used and benchmarking results were gathered from VMmark output.

The benchmark helps to measure the virtualization overheads of the individual workloads as well as the scalability of the entire system. Therefore, results for multi-tile runs are reported as the aggregate score for all tiles, the individual scores for each of the tiles, and the scores for the workloads within the tiles as well as the individual scores for each infrastructure workload.

The following sections summarize the results observed.

### Application score comparison between XIV 314 and 214 models

VMmark runs workloads for applications that are typically used in a virtualized environment. They are run in tiles, with each tile consisting of an exchange application, an Olio application, and DVD Store 2 application workloads. After running VMmark benchmarks, the metrics of each tile is computed and aggregated. Then, a geometric mean is computed as the final score for the tile. With all of the per-tile score added, to create the application workload portion of the final metric to provide application score for a tile.

Figure 22 shows the application score comparison between a compressed volume on IBM XIV Model 314 and an uncompressed volume on XIV Model 214. As per the results, the VMmark application score was almost the same across both the XIV models.

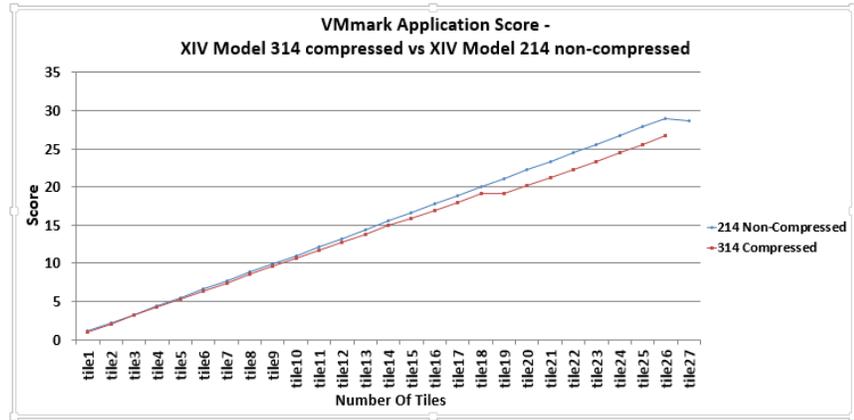


Figure 22: VMmark application score comparison

### Infrastructure score comparison between XIV 314 and 214 models

While running application benchmarks, VMmark also runs infrastructure workloads such as cloning and deploying of VMs, automatic VM load balancing across a datacenter, VM live migration (VMware vMotion) and dynamic datastore relocation (storage vMotion) to determine infrastructure score. These infrastructure operations complement conventional application workloads.

Figure 23 shows the infrastructure score comparison between a compressed volume on IBM XIV Model 314 and an uncompressed volume on XIV Model 214. Results show that the infrastructure score observed on XIV Model 314 was relatively better compared to the score on XIV Model 214.

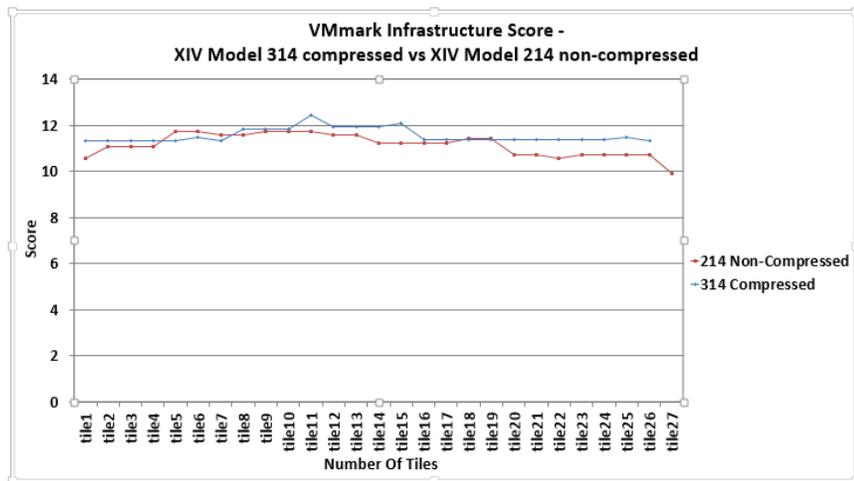


Figure 23: VMmark infrastructure score comparison

### VMmark total score comparison between XIV 314 and 214 models

The final benchmark score is computed as a weighted average of the application-workload component and the infrastructure-workload component. VMmark gives a weightage of 80% to the application-workload component and 20% to the infrastructure-workload component. These weightages are chosen to reflect the relative contribution of infrastructure and application workloads to overall resource demands.

Figure 24 shows the total score comparison between a compressed volume on XIV Model 314 and an uncompressed volume on XIV Model 214. Benchmarking score showed negligible difference between both the models of XIV.

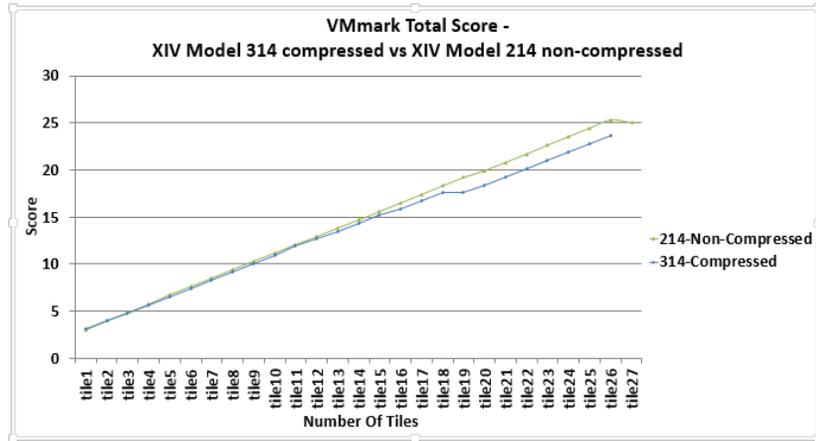


Figure 24: VMmark total score comparison

VMmark benchmarking scores observed from the three comparison charts (in Figure 22, Figure 23, and Figure 24) indicate that XIV Model 314 with compressed volumes produce similar benchmarking scores as that of XIV Model 214 with non-compressed volumes. With a single processor and 48 GB RAM dedicated to Real-time Compression for each module of XIV Model 314, there is no performance impact to the VMware workload on the XIV system. Furthermore, with the addition of 4 TB and 6 TB hard drives and increasing the effective capacity up to 2 PB in a single 15-module frame with a 2:1 compression guarantee, XIV Model 314 delivers large capacity savings without affecting the performance.

## Summary

---

The IBM Real-time Compression integrated with the IBM XIV Storage System, can be used to effectively reduce the storage capacity required for many types of data. The patented algorithms implemented in the RACE technology of the Real-time Compression feature allow data to be compressed in real time. Real-time Compression is fully transparent to hosts and applications, and compression rates are consistent over time. IBM XIV grid architecture allows compression workload to spread across multiple modules of XIV and hence Real-time Compression has minimized performance impact on the applications running on XIV Storage System. Furthermore, by dedicating processor and RAM for Real-time Compression, XIV Model 314 delivers significant compression savings without affecting performance. The XIV Model 314 system also reduces the data center footprints and power consumption as a result of higher effective capacity per XIV system.

VMware vSphere virtual machines can be seamlessly deployed on the compressed volumes, often with the compression savings of 50% to 75%, allowing customers to reduce the storage capacity required for virtualized environments.

## Acknowledgements

---

The author would like to thank Itay Raviv and Moshe Kalish for running the VMmark benchmark tests.



## Appendix A: Test environment

---

The following information provides details about the test environment used for testing the solution.

IBM XIV Storage System : 11.6.0

VMware ESX Version : 6.0

## Appendix B: Resources

---

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld®  
[ibm.com/partnerworld/systems](http://ibm.com/partnerworld/systems)
- IBM Redbooks®  
[ibm.com/redbooks](http://ibm.com/redbooks)
- IBM Publications Center  
[www.elink.ibmink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US](http://www.elink.ibmink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US)
- IBM XIV Storage Knowledge Center  
[ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv\\_kcwelcomepage.html](http://ibm.com/support/knowledgecenter/STJTAG/com.ibm.help.xivgen3.doc/xiv_kcwelcomepage.html)
- IBM XIV with Real-time Compression feature in Oracle environment  
[ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102551](http://ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102551)
- Microsoft Windows server 2012 R2 cloud benefits using IBM XIV Storage System Gen3 Real-time Compression  
[ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102553](http://ibm.com/support/techdocs/atmastr.nsf/WebIndex/WP102553)
- VMware vSphere 6.0 Documentation Center  
<https://www.vmware.com/support/pubs/vsphere-esxi-vcenter-server-6-pubs.html>
- VMware VMmark benchmark overview  
[www.vmware.com/products/vmmark](http://www.vmware.com/products/vmmark)

## About the author

---

**Mandar J. Vaidya** is a Storage Technical Enablement Specialist in IBM Systems ISV Enablement group. He has more than 15 years of experience working with various storage and systems technologies. Mandar holds a Bachelor of Engineering degree from the University of Pune, India. You can reach Mandar at [mandar.vaidya@in.ibm.com](mailto:mandar.vaidya@in.ibm.com).





## Trademarks and special notices

---

© Copyright IBM Corporation 2016.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and [ibm.com](http://ibm.com) are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.



Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.