Microsoft Windows Server 2012 R2 cloud benefits using IBM XIV Storage System Gen3 Real-time Compression

Eric B. Johnson

IBM Systems and Technology Group ISV Enablement
January 2016

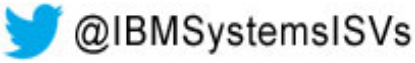

Table of contents

Abstract

With the introduction of 11.6.x code, IBM XIV Storage System Gen3 now provides Real-time Compression to help Microsoft Windows Server 2012 R2 cloud environments greatly reduce physical storage requirements and substantially lower business critical file share and application service costs while maintaining enterprise-class performance, data availability and efficiency. To demonstrate, IBM Real-time Compression implementation guidelines, examples, and XIV model 214 and 314 performance comparisons are revealed throughout this paper to help cloud administrators optimize their IBM XIV data compression experience.

Target audience

This white paper is intended for mid-to-large size businesses with Microsoft® cloud administrators, solutions architects, system administrators and storage engineers responsible for Microsoft Windows® Server 2012 R2 operation and management of Hyper-V solutions using or evaluating IBM® XIV® Storage System Gen3. Current virtualization environments looking to control data proliferation, by minimizing storage consumption without negatively impacting workloads, should benefit most from the latest IBM XIV Real-time Compression™ features.

Advanced comprehension of Microsoft Windows Server, Microsoft Failover Clustering, Hyper-V technology, Microsoft Exchange Server, Microsoft SQL Server, and XIV Storage System Gen3 administration is beneficial. However, technical reviews and supplemental references are provided throughout the paper.

Introduction

Now more than ever as companies consume increasing amounts of data, information technology (IT) departments scramble to meet storage capacity demands. The IBM XIV Storage System Gen3 helps organizations address such burgeoning requirements by adding new Real-time Compression capabilities that aid in curtailing excessive storage sprawl and promoting lower data center energy costs. This optional licensed yet fully integrated IBM XIV functionality provides an attractive option in comparison to other storage compression solutions including freestanding appliances for existing and new IBM customers.

What makes IBM XIV Real-time Compression so appealing is its ability to provide real-time primary or active data compression for a full spectrum of file types ranging from virtualization images to application databases while maintaining high performance for physical and virtual machines (VM). This allows IBM XIV to deliver superior Real-time Compression advantages versus competitive storage compression technologies that only affect secondary or backup data that is ordinarily post-processed during non-business hours. As such, Real-time Compression is just the latest addition in a long list of features that makes IBM XIV a smart choice for businesses demanding the most out of their data center investments.

IBM XIV Storage System Gen3

IBM XIV® is a high-end, grid-scale storage system that is extremely easy to use and acclaimed for its tune-free high performance, and excellent data economics including inline, field-proven IBM Real-time Compression. Built with IBM Spectrum Accelerate™, IBM XIV offers unified management and operational agility across your physical and virtual data centers.

In fact, IBM XIV is ideal for private, public, or hybrid clouds, offering high service levels for dynamic workloads, easy hyper-scaling for various environments including multi-tenant, and flexible consumption models. XIV also continues to advance its robust cloud automation and orchestration with OpenStack, RESTful API, VMware, and the latest SMI-S enhancements for Microsoft. Furthermore, it offers security and data protection through advanced mirroring, hot encryption, self-healing, and XIV storage abstraction through VMware vSphere vStorage APIs for Storage Awareness (VASA) 2.0. So, your cloud data integrity and availability is safe and assured. Also, industry benchmarks underscore exemplary XIV performance and cost benefits.

The following list highlights some of the IBM XIV Storage System Gen3 cloud-specific features and benefits:

- Full support of Microsoft technologies including cloud-friendly integrated SMI-S for System Center VMM storage automation, GeoClustering, Volume Shadow Copy Services (VSS), and Multipath I/O (MPIO)
- Disaster recovery using iSCSI or Fibre Channel (FC) synchronous and asynchronous data replication that includes up to three-way mirroring
- Real-time Compression option for capacity savings up to five times greater usable capacity of one petabyte or more per frame while maintaining high performance
- Solid-state drive (SSD) flash cache option for breakthrough performance levels
- Capacity on demand configurations
- Simple storage administration including hyper-scale management up to 144 frames
- IBM Hyper-Scale Mobility for online volume movement between XIV systems, non-disruptive to applications, for vast operational agility
- Ultra-intuitive reports on capacity growth, usage, and trends empower administrators to understand and act upon storage needs quickly
- FC and iSCSI host connectivity
- Self-encrypting hard drives

For more IBM XIV Storage System Gen3 details, refer to the following website:
ibm.com/systems/storage/disk/xiv

Introducing IBM XIV Model 314

IBM XIV Storage System Gen3 Model 314 with software version 11.6.1 builds upon the strong XIV Storage System foundation and provides greatly enhanced functional capabilities. With double the RAM and CPU resources, IBM XIV Storage System Gen3 Model 314 delivers improved IOPS per compressed capacity and user-configurable soft capacity up to 2 PB (versus up to three times the hard capacity) without performance degradation. Furthermore, each module now contains a dedicated CPU and 48 GB RAM that are used for Real-time Compression with a 2:1 compression ratio guarantee, which is now included in the base XIV license and enabled by default. At double the compression ratio, the effective compressed read/write cache capacity in a 15-module XIV system is now 1440 GB, delivering a total effective RAM of 2160 GB. The SSD read-cache capacity per system is also doubled to 24 TB.

IBM XIV Storage System Model 314 is available with the following:

- 9 - 15 module configurations per system

- 2 x 6-core CPUs per module (versus 1 x 6-core CPU per module in Model 214)
- 96 GB RAM per module
- 4 TB or 6 TB drives with 800 GB SSD cache per module
- Reduced minimum compressible volume size from 103 GB to 51 GB
- Support for IBM Spectrum Accelerate software licenses with XIV Storage System Gen3 Model 314
- Worldwide response support for IBM XIV Storage System Model 314, 2-hour response time, 24 hours a day, seven days a week, for severity 2 calls

For more IBM XIV Storage System Gen3 Model 314 details, refer to the following website: ibm.com/systems/storage/disk/xiv/specifications.html

System x servers

Also essential to virtual environments and offering numerous storage connectivity options while capable of driving a full variety of workloads, System x servers effectively deliver the Hyper-V cloud performance and reliability required for Microsoft's popular business critical applications. A complete range of powerful, entry-level to large-scale, uncomplicated systems is available to accommodate any budget or data center server needs. For the compression testing purposes of this paper, System x3650 M4 hosts were connected to IBM XIV Storage System Gen3 using a meshed Brocade 2498 FC storage area network (SAN).

The System x3650 M4 server is packaged in a highly dense 2U form factor to deliver the necessary performance and reliability provided by the following notable features:

- Intel Xeon E5-2600 v2 series processors
- Up to 768 GB of memory
- Quad-port integrated GbE with optional embedded dual-port 10 GbE network interfaces
- Integrated management module (IMM) that provides system administrators remote server control for common tasks that typically require physical presence

For complete details about System x servers, refer to the following website: **ibm.com**/systems/x

Microsoft Hyper-V overview

Just as IBM storage solutions, enhanced with the latest compression technologies, allow customers to do more with less, Microsoft Hyper-V virtualization also helps customers maximize System x server and other resource use. Included in Windows Server, Microsoft Hyper-V virtualization helps reduce costs by allowing a greater number of application workloads to be hosted on fewer physical servers. Consequently, Hyper-V improves IT efficiency and flexibility with rapid virtualized application deployment and maintenance for a much smaller data center footprint. Moreover, Hyper-V off-premises and on-premises cloud solutions are greatly augmented and easily managed by the robust and tightly integrated suite of Microsoft System Center offerings.

Likewise, IBM XIV and Microsoft cloud solutions seamlessly couple together to offer powerful end-to-end virtualization benefits. A partial list of Hyper-V storage-specific benefits and features are now available with the latest Microsoft Windows Server 2012 R2 and IBM XIV 11.6.x releases:

- Hosts and VMs can use thick and thin provisioned SAN volumes
- Space reclamation is available using SCSI UNMAP
- Hosts and guests can use multipath I/O for storage connectivity
- Guest virtual Fibre Channel technology using host 8Gb FC
- 1 Gb or 10 Gb iSCSI storage connectivity
- Fibre Channel over Ethernet (FCoE) storage connectivity
- VM virtual hard disk (VHD) and enhanced VHDX file formats
- Shared VHDX files for Microsoft guest clusters
- Online VHDX resizing
- Storage Quality of Service (QoS)
- Pass-through disks
- IBM XIV online volume migration using Hyper-Scale Mobility
- Cluster Shared Volume (CSV) cache
- Resource metering
- Host and guest VSS snapshots
- Storage automation support for Microsoft System Center Virtual Machine Manager 2012 including the ability to use IBM XIV pools that enable compression by default

For further information about Microsoft Hyper-V, visit the following website:

http://technet.microsoft.com/en-us/library/hh831531.aspx

For further information about Hyper-V scalability in Windows Server 2012, visit the following website:

http://technet.microsoft.com/en-us/library/jj680093.aspx

Hyper-V with IBM XIV prerequisites

Before deploying Microsoft Hyper-V solutions with IBM XIV Storage System Gen3 with Real-time Compression enabled, there are numerous recommendations. First and foremost, all system hardware should be updated to the latest supported firmware and software or drivers. Additionally, the Hyper-V hosts should run Windows update to ensure the latest Hyper-V, Microsoft Failover Cluster, Microsoft application and other critical operating system (OS) hotfixes, including security, are up-to-date. The latest Windows Server 2012 R2 Hyper-V hotfixes are located at the following website:

http://social.technet.microsoft.com/wiki/contents/articles/20885.hyper-v-update-list-for-windows-server-2012-r2.aspx

Recommended hotfixes and updates for Windows Server 2012 R2 failover clusters are also located at the following website:

http://support.microsoft.com/kb/2920151

When configuring Hyper-V hosts, Microsoft recommends limiting the roles and features to just the Hyper-V role, Failover Cluster and Multi-path I/O features. Thus, Hyper-V hosts must be dedicated to running only the necessary core stand-alone or clustered hypervisor components to help reduce non-essential resource usage, minimize solution complexity, and ease potential system-based troubleshooting. Instead, consider installing extraneous roles, features, and necessary business applications on the VMs.

Furthermore, when configuring VMs, the Hyper-V host system drive must not be used for the default virtual hard disk location. This helps to avoid host system disk latency and free space depletion issues. Preferably, select a non-system IBM XIV or other volume path that is more appropriate to individual stand-alone or clustered Hyper-V configurations.

Also to avoid similar performance impacts, Microsoft recommends only running antivirus software in the guest OS rather than in the Hyper-V host OS. If this is impractical due to company security policies, ensure that the Hyper-V host VM file directories are excluded from the antivirus scans to help prevent performance degradation and stability issues. For further Hyper-V antivirus exclusion details, see the following website:

http://social.technet.microsoft.com/wiki/contents/articles/2179.hyper-v-anti-virus-exclusions-for-hyper-v-hosts.aspx

Even though this is an impartial list of Hyper-V guidelines and individual environments might require further considerations, the aforementioned suggestions can help to alleviate many of the common virtualization pitfalls. For a complete list of Hyper-V prerequisites and detailed best practices, see the Microsoft TechNet website for countless excellent resources such as the Hyper-V pre-deployment and configuration guide at:

http://blogs.technet.com/b/askpfeplat/archive/2013/03/10/windows-server-2012-hyper-v-best-practices-in-easy-checklist-form.aspx

IBM specific solution support should be reviewed at the IBM System Storage Interoperation Center (SSIC). The SSIC allows customers to validate interoperability for the most popular IBM multi-vendor hardware and software combinations. The SSIC support matrices are located at the following website:

ibm.com/systems/support/storage/ssic/interoperability.wss

IBM XIV compression-specific requirements

IBM XIV Real-time Compression has the following requirements:

- IBM XIV 11.6.x or later code
- Thin provisioned pools with compression enabled for the XIV volumes
  Note: Sufficient pool free space is required for volume compression and decompression including snapshots.
- XIV volumes to be compressed must be a minimum of 103 GB (now 51 GB with 11.6.1.x code) and up to a maximum of 10 TB in size
  Note: To compress existing smaller-than-51 GB volumes, customers can resize the volume to 51 GB before the conversion. To accommodate the larger volume, the related thin pool soft capacity may need to be enlarged too. Keep in mind, a host rescan may also be required. Otherwise, there is no significant operational impact. The actual hard capacity needed for this volume will be reduced by compression as expected.
- XIV supports a maximum of 1024 compressed volumes (including snapshots)
- Synchronous, asynchronous, and three-way mirroring is supported for compressed volumes if all XIV systems have 11.6.x or later code

Hyper-V volume configurations

Before examining compression concepts closer, typical Microsoft Hyper-V volume types should be considered. For application virtual machines, there are two types that are distinguished by their volume presentations and definitions. IBM XIV volumes can be mapped directly to a Hyper-V host or to a guest, and quite often, both types are used for a single VM. Hyper-V common volume configurations include:

- Host volume VHD or VHDX files that can be used for the guest OS system files or dedicated application data files
- Guest or VM direct access storage using host pass-through disks or virtual Fibre Channel for VM SAN volume access that is equivalent to physical machine SAN volume access

Note: Both volume types are popular in small-to-large sized Microsoft cloud production environments and are simply a matter of preference.

While it is common to use VHDX files for the guest OS system files, it is also common to place multiple VMs and their corresponding VHDX files on a single XIV volume presented to a Hyper-V host or failover cluster. In general, IBM recommends using 2 TB or larger XIV volumes with no more than 20 VMs each. This helps to mitigate risk in the event of catastrophic file system, software, or hardware corruption or failure that impacts a single volume. In other words, do not place all of your eggs in one basket. Likewise, administrators may want to create separate application data and log VHDX files to potentially expedite disaster recovery. Using separate application data and log VHDX files is not an XIV performance measure but more of an administrative granular organization or precautionary measure.

If maximum performance is a top priority, as with some business critical database applications, it is best to use direct access storage whether it is virtual FC, iSCSI using 10 GbE, or pass-through disks. In this case, VM direct access to the XIV Storage System occurs through the underlying physical host components including FC host bus adapters (HBA), network adapters, and so on by presenting extra hardware-based I/O paths to the Windows software virtual hard disk stack. So, it is basically the performance equivalent of physical machines accessing SAN storage. As such, it is the physical server, hardware bandwidth and throughput limitations that determine how many VMs may run on any single host and what type of performance each VM can expect. Thus when it comes to virtual application data and log volumes, they should also be sized according to traditional physical server application methods. However, it is advisable to test VM application data and log VHDX files in comparison to VM direct access storage before production deployments because performance is just one of the determining factors. Ultimately, XIV compression does not dictate the performance behavior as much as the overall business critical application and Hyper-V VM configuration variables.

For more information about the best practices for using Microsoft Hyper-V with IBM XIV Storage System Gen3, visit the following website: **ibm.com**/support/techdocs/atsmastr.nsf/WebIndex/WP102456

Compression concepts

Many people are familiar with general compression concepts through experience with popular file compression software. Applications such as *WinZip* or *gzip* provide similar space-saving results (albeit slow software-based results for inactive files) when compared to IBM XIV Real-time Compression. Just like popular software file compression, not all file types make good candidates for IBM XIV Real-time

Compression. Files that are already highly compressed do not benefit from IBM XIV Real-time Compression. Many multimedia and software encrypted files fall into this latter category including already highly compressed video, audio, and image files. Although it is worth mentioning that self-encrypting drives and other back-end encryption methods are supported with IBM XIV Real-time Compression. Regardless, virtual or physical data centers that contain good compression file candidates including uncompressed file share data and many Microsoft Exchange or SQL Server databases can realize compression savings commonly in the 40-80% range that can result in up to five times as much effective data storage.

This is because many virtual and physical file shares or databases, whether Microsoft Exchange or SQL Server, consist of file data that contains large amounts of repetitive alphanumeric characters that are ideal for IBM XIV Real-time Compression. Moreover, IBM XIV Real-time Compression can readily compress Microsoft Exchange and SQL Server database white space and free space. To better appreciate this concept, a basic review of Exchange database white space and SQL Server database free space is useful.

Exchange white space and SQL free space is rather straightforward, once the database activity is examined. Clearly as data is entered and accumulates, the physical size of a database increases. When data is removed such as when deleting data, migrating data (for example user mailboxes) to other databases and so on, the amount of database data often decreases yet the database file sizes may not decrease. As a result, databases contain the equivalent of empty space called *white space* in Microsoft Exchange and *free space* in Microsoft SQL Server.

Most of the time, database white space or free space is inconsequential considering Exchange or SQL Server maintenance routines, in newer releases, can be used to reduce the database size and increase disk capacity. However, in the event, white or free space becomes problematic or reclaiming disk space is suddenly required (particularly for older Microsoft Exchange or SQL versions), then IBM XIV Real-time Compression provides a simple, alternative solution for databases that are migrated to or reside on compressed volumes.

Furthermore, IBM XIV Real-time Compression can also help improve data protection speed and lower capacity requirements. IBM XIV Real-time Compression can be used to reduce backup and restore capacity requirements because many file shares and databases can be significantly compressed, including white or free space. Therefore, the destination backup storage capacity requirements can be reduced. Accordingly, this same logic applies to data replication involving mirrored volumes between multiple IBM XIV systems running 11.6.x code or greater. As there is less data to replicate due to compression, the performance gains can also be realized by freeing up inter-storage array network bandwidth, thus enhancing disaster recovery point objectives (RPO).

To shed more light on the topic of IBM XIV Real-time Compression data protection, whether additional physical capacity is required for backups depends on the media. For practical purposes, three types of backup media can be cited– tape, IBM XIV compressed volumes, and IBM XIV or other storage uncompressed volumes. For backup tape capacity, this does not change when using IBM XIV Real-time Compression. Assuming a Microsoft VSS backup environment, IBM XIV compressed source data is decompressed before backing up the data to tape. Thus the backup and restore process is just like other familiar conventional data protection processes. In other words with tape backups, it is business as usual. However, if one performs IBM XIV compressed volume-to-volume backups or copies, all Real-time

Compression benefits are realized with the capacity savings and the like. Finally, backup data is decompressed if it is placed on any uncompressed volume so it is essentially the same as backing up to tape.

IBM XIV Real-time Compression overview

For existing and new customers considering IBM XIV Real-time compression for their organizations, a 45-day free trial is available with model 214. This allows administrators to evaluate and test their Microsoft cloud applications and individual environments. Just enable compression for the IBM XIV system and run the evaluation guide. After the trial period, customers can elect to purchase the required licenses or disable the feature. For compression assistance during the evaluation period, contact your local Business Partner, IBM representative or use the email address for the compression team at XIV_RTC@il.ibm.com. For highly compressible data volumes, organizations can expect to see the following immediate and long-term IBM XIV Real-time Compression green IT solution benefits:

- Helps curb storage sprawl and promote smaller data center footprints
- Lowers energy consumption by reducing electrical and cooling costs thus minimizing environmental impact
- Helps avoid unnecessary hardware purchases and compliance with IT budget constraints
- Decreases and streamlines software costs for data compression

To determine the best Microsoft cloud application or other data volumes that contain ideal compression file candidates, customers should use the Comprestimator utility which is included in 11.6.x or newer code combined with the latest IBM XIV management tools download at IBM Fix Central. For regular and thin pools with uncompressed volumes, the Comprestimator utility can be accessed using the IBM XIV Storage System command-line interface (XCLI) to perform the compression savings calculations. However, for customers using other storage arrays, a downloadable version of the Comprestimator utility is also available to estimate compression savings for data that may be migrated to IBM XIV Storage System. The utility is accurate within 5 percent based on IBM patented Random Access Compression Engine (RACE) formulas. For more information, see the following website:

www-304.ibm.com/webapp/set2/sas/f/comprestimator/home.html

After using Comprestimator on volumes to determine if they are good candidates, it is simply a matter of taking advantage of the new IBM XIV Real-time Compression feature and immediately realizing the compression benefits. After the feature is enabled, IBM XIV Real-time Compression commences at the volume level. Once an IBM XIV compressed volume is created, it uses the same proven RACE code as stand-alone IBM Real-time Compression appliances, IBM Storwize® family, and IBM FlashSystem™ products. Embedded in new software of each IBM XIV module, RACE drives host transparent compression of active primary data in real time with negligible performance degradation. Just like thin provisioning, Real-time Compression uses on demand block allocation of data based on host writes versus allocating all of the blocks during the initial volume creation. However, volumes have supplemental Real-time Compression benefits during host I/O operations. So there are multiple sides to Real-time Compression.

After thin provision-like capacity savings are factored in, a closer examination of Real-time Compression reveals the additional compression advantages. At a high level using temporal, lossless data compression

algorithms, the RACE literally shrinks the primary active host data as it is written in real time. It is important to note that the RACE manages only the parts of the compressed volume that belong to its corresponding interface or data module and profits from the highly parallel IBM XIV design. Thus, the IBM Real-time Compression processor and memory requirements are balanced across up to 15 modules resulting in negligible performance impact. Likewise and unique to the IBM XIV, Real-time Compression takes place before reaching cache and after the compressed data is written into cache, the write is acknowledged back to the host. The resultant minimal latency is proportionate to a maximum block size of 64 KB from the I/O interface. Furthermore, the interface splits the I/O if it is larger than 64 KB.

When it comes to host read operations, RACE behaves differently. Reads are decompressed in-flight when they are read from cache before they are transferred back to the interface. To expedite this process, IBM Real-time Compression uses the XIV prefetch mechanism for compressed application data. Moreover, regarding host random read requests, they also become storage random reads because RACE needs to decompress each read request. Because the cache is also compressed, the effective disk cache is much larger. As a result, the cache hit ratio and random read performance potentially improves. Also worth mentioning, less data being read or written to disk means less mechanical movement of disk actuator arms during host I/Os, resulting in increased storage efficiency. Not to mention, when there is more data in cache and IBM XIV's global read cache extension using SSDs, compressed volumes can take advantage of improved cache hit ratios combined with less latency.

Finally as expected, uncompressed volumes simply bypass the RACE software nodes and follow the standard IBM XIV data paths. Figure 1 illustrates the data route comparison between IBM XIV conventional and compressed volumes.

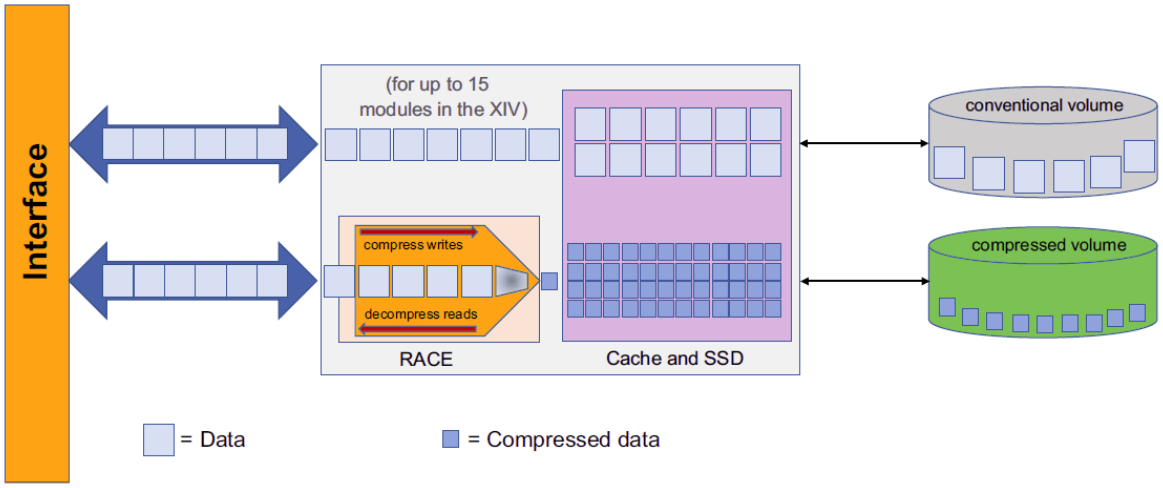

Figure 1: Data route comparison between IBM XIV regular and compressed volumes.

Testing

Two separate real-time compression test phases were run on two IBM XIV Storage Systems with 11.6.1.x code, including models 214 and 314. The first test phase focused on compression savings and two different VM file types were examined - virtual Microsoft SQL Server 2012 SP1 online transaction processing (OLTP) data files and VM Windows Server 2012 R2 system files. Both VHDX and virtual FC

volume types were tested for compression savings. Basic compression testing for Microsoft SQL Server 2012 SP1 OLTP data files included separate VHDX files and guest virtual Fibre Channel XIV volumes. Essentially, the identical SQL data files were copied to both VHDX and virtual FC volume types for comparison. Both VHDX and VM direct-accessed volumes achieved 73% compression savings as illustrated in Figure 2.

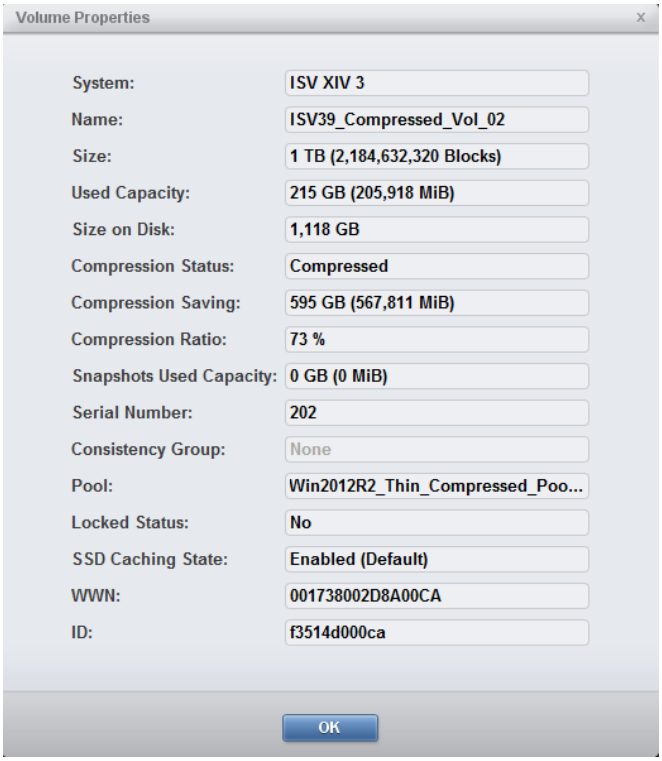

Figure 2: Microsoft SQL Server 2012 OLTP data volume properties with resulting compression savings

Windows Server 2012 R2 VM system files were placed on default VHDX files. Keep in mind when it comes to Hyper-V VM system files, the compression ratio calculations vary based on the types of Microsoft Windows Server 2012 roles and features [for example, Internet Information Server (IIS) files and website content, Active Directory databases and so on], as well as application or other files present on the root drive. Virtualization image file compression also depends on how much highly compressed data is contained in the image files. Figure 3 depicts the IBM XIV graphical user interface (GUI) compression properties for a Microsoft CSV with multiple types of Hyper-V VMs. For Microsoft Windows OS VM VHDX files with a variety of roles and features, an average compression ratio of 40% is common.

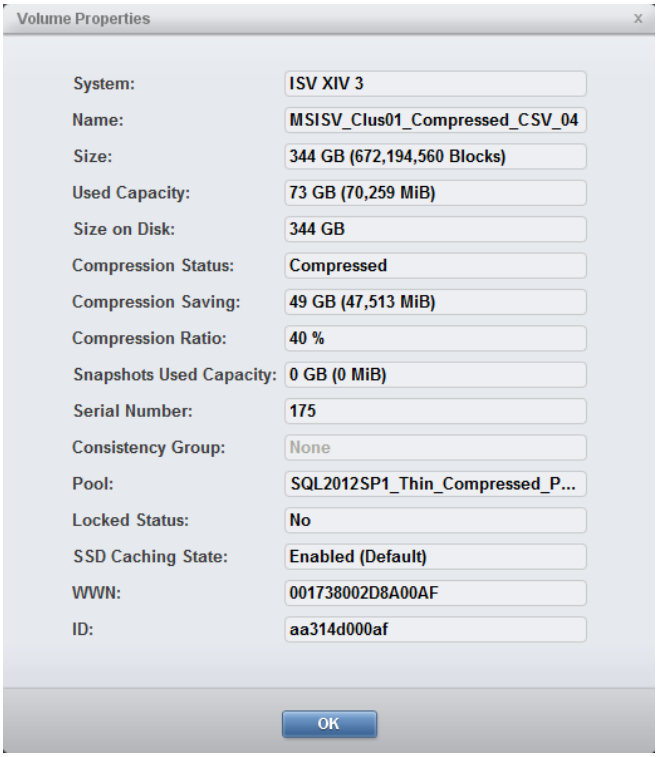

Figure 3: IBM XIV GUI compression properties of Microsoft CSV with Hyper-V VMs

For combined performance and compression testing, the Benchmark Factory utility is used to create Microsoft SQL OLTP databases (as close to real-world as possible) to simulate production workloads that could capture the IBM XIV real-time compression behavior. Specifically, the tool is used to generate Microsoft SQL TPC-C workloads that were run multiple times to confirm and achieve consistent test results to aid customer deployment planning.

The test environment includes the following key VM and benchmark variables. The Microsoft SQL Server VM is configured with 6 virtual processors and 16 GB of RAM. For the Benchmark Factory utility, the Microsoft SQL Server TPC-C scale was set to a 32,768 maximum user limit for a 5 TB database and resulted in a data compression ratio of 71% that aligns with the previous compression-only results. As a reminder, for TPC-C, performance is measured by response times to preset transactions and number of transactions per minute achieved, listed as a tpmC value.

The charts in **Error! Reference source not found.**, **Error! Reference source not found.**, Figure 6, and Figure 7 reveal the performance comparison between IBM XIV model 214 non-compressed and 314 compressed Microsoft SQL VHDX and vFCA data volumes as the user workload ramps up to the maximum user limit. Due to lab resource and benchmark maximum user limitations, the performance charts do not reveal that the IBM XIV model 314 can easily handle workloads of 150,000 IOPS and greater. However, the lab tests were sufficient to allow administrators to extrapolate higher XIV model 314 performance with compression enabled and provide a good comparison to model 214.

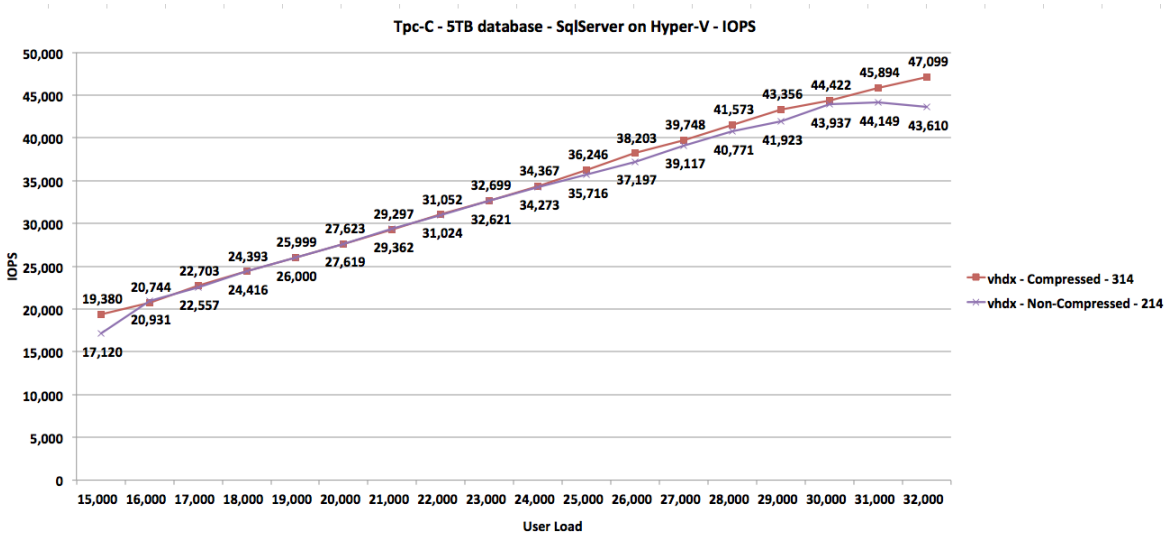

Figure 4: Microsoft SQL Server 2012 SP1 Hyper-V VM VHDX IOPS during TPC-C benchmark testing

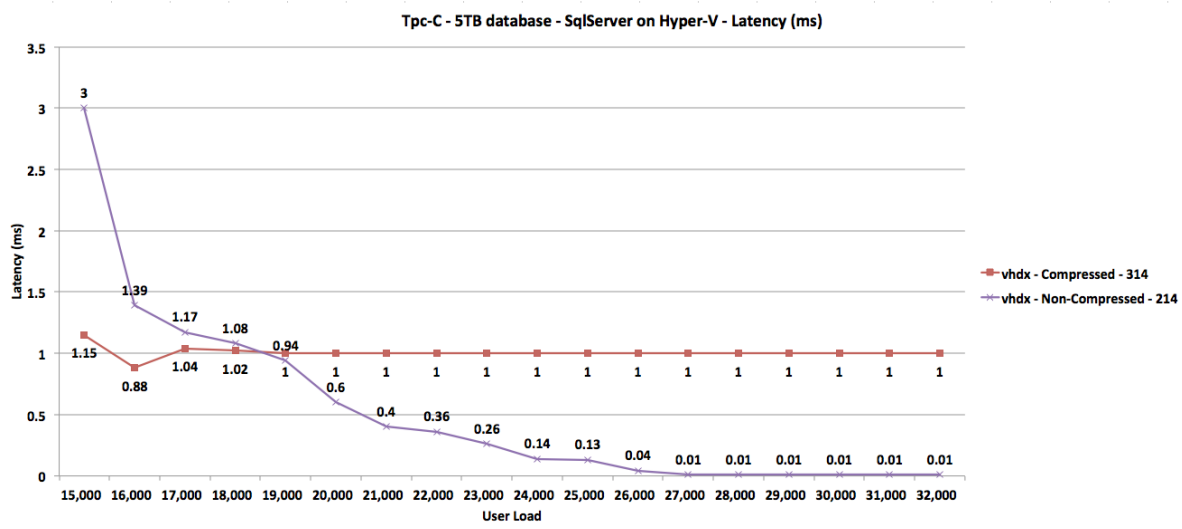

Figure 5: Microsoft SQL Server 2012 SP1 VHDX application response time during TPC-C benchmark testing

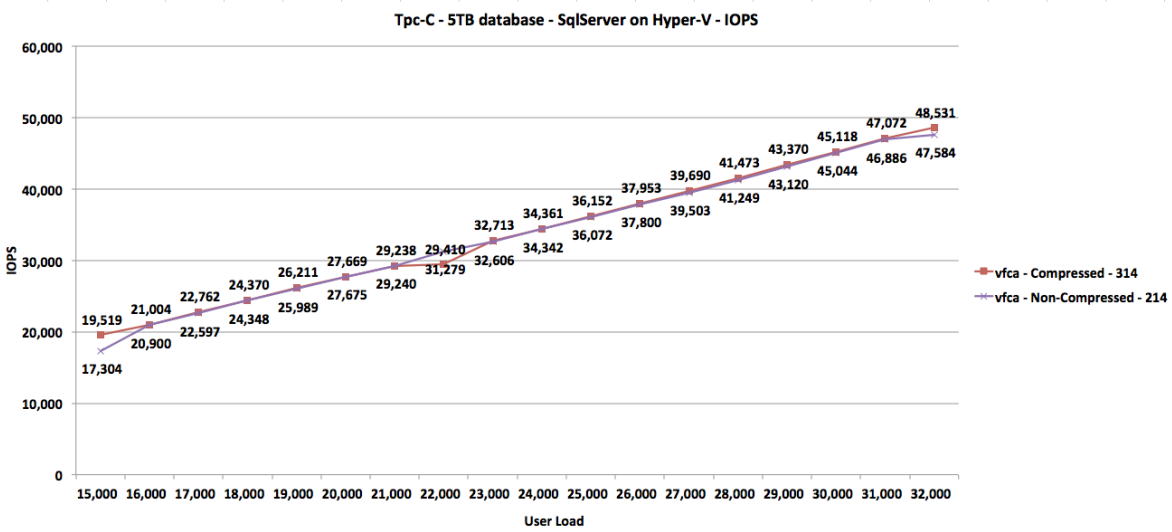

Figure 6: Microsoft SQL Server 2012 SP1 Hyper-V VM vFCA IOPS during TPC-C benchmark testing

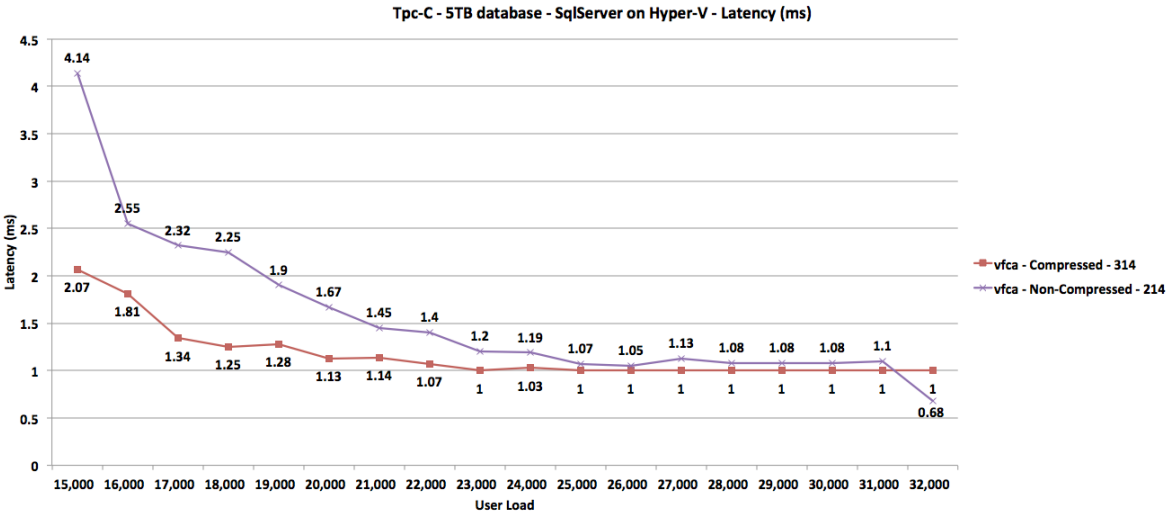

Figure 7: Microsoft SQL Server 2012 SP1 vFCA application response time during TPC-C benchmark testing

As the charts indicate, the IBMX XIV model 314 compressed performance compared to model 214 non-compressed clearly reveals the enhanced hardware upgrades of the newer model. With a single CPU and 48 GB RAM dedicated to real-time compression for each XIV module, the model 314 compressed benchmarks outperform model 214 IOPS. Additionally, application latency for model 314 is 2-3 times better for the lower user loads and never exceeds 1 ms for the higher user loads which is fantastic when considered for real world database production performance. Furthermore, with the addition of 4 and 6 TB hard drives with up to 970 TB of effective capacity in a single 15 module frame with a 2:1 compression guarantee, you can see how attractive the compression performance and large capacity savings becomes. So fundamentally as demonstrated, the Benchmark Factory test results reveal IBM XIV model 314 real-time compression provides substantial capacity savings while maintaining high performance in comparison to its predecessor XIV during typical virtual Microsoft SQL Server OLTP workloads.

For additional IBM XIV Real-time Compression performance or other information, review the following IBM Redpaper™: www.redbooks.ibm.com/abstracts/redp5215.html?Open

IBM XIV Real-time Compression implementation

Once an IBM XIV has been updated to 11.6.x code, it can be enabled for compression by a few different methods. To enable Real-time Compression from the XIV GUI, simply use the right-click menu on the system or from the system **Parameters** tab as seen in Figure 8. If the command line interface is preferred, use the XCLI command, *system_compression_enable*.

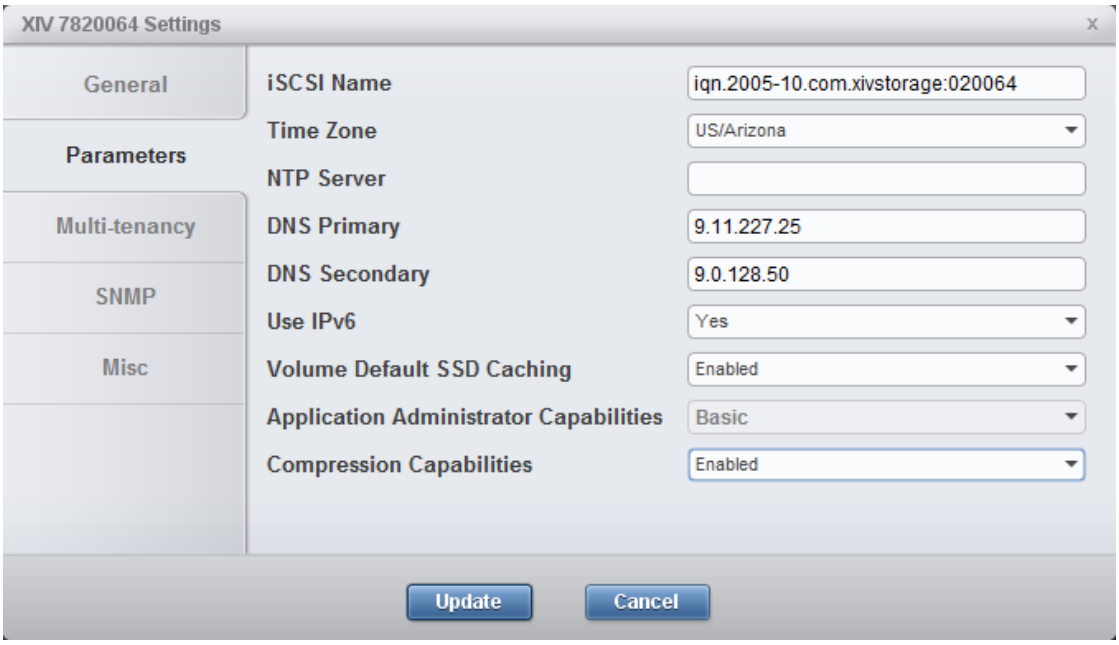

Figure 8: Enabling IBM XIV compression capabilities

Once compression is enabled for the system, there are a few considerations to keep in mind. To compress data for an existing regular or uncompressed volume, just use the XIV GUI to drag and drop the volume from a regular pool to a thin pool with compression enabled that has sufficient free space for the volume addition. In order for this to work properly, the volume must be a minimum size of 103 GB (now 51 GB with 11.6.1.x code) and up to a maximum of 10 TB which is the standard IBM XIV requirement for all compressed volumes. If an administrator attempts to compress a newly moved volume that does not satisfy the size requirements, the operation fails with a pop up error noting the compressed volume size infraction. Also note, IBM XIV supports up to a total of 1024 compressed volumes. Nevertheless, once the regular volume has been moved to a compressed thin pool, just right-click the volume and select compress. Notably, a failsafe mechanism provides administrators the option to keep a copy of the source uncompressed volume during the conversion process. After testing the integrity of the new compressed volume by using duplicate test VM applications (mirrored after production environments), decommission or delete the uncompressed source copy, if present, to free up space for new capacity requests. Once a volume completes the conversion process or a new compressed volume is created, its corresponding reads and writes can take advantage of the RACE.

Summary

The IBM XIV Storage System Gen3 simple and intuitive user manageability, flexibility, reliability and performance deliver an immediate return on investment (ROI) with the addition of Real-time Compression capacity savings in the 40-80% range or up to 5 times the effective storage. By remarkably reducing storage requirements, IT environments can use the enhanced efficiency of IBM XIV Real-time Compression for instant higher data capacity while deflating virtual and physical data center storage costs. Consequently, maintaining a smaller data center footprint helps to ensure less administrative burdensome tasks, compliance with IT budget constraints, and reduced energy consumption for a lower environmental impact.

It all starts by analyzing current cloud and physical environments for optimal compression candidates such as Microsoft database or file share volumes using the simple Comprestimator utility. Once this task is completed and compression is enabled for the IBM XIV Storage System Gen3, customers realize immediate cost and data capacity savings with negligible performance degradation. Lastly, unlike many competitors, IBM XIV performs Real-time Compression for active primary files which is preferential to compression solutions that only process secondary or backup data after it is written to disk at a later time.

For additional information about IBM XIV Real-time Compression, review the web-based pointers in the Resources section.

Resources

For further information about the various solutions products, visit the following websites:

- IBM XIV Storage System
  ibm.com/systems/storage/disk/xiv/index.html
- IBM XIV Storage System Architecture and Implementation Redbook
  www.redbooks.**ibm.com**/abstracts/sg247659.html?Open
- IBM Real-time Compression on the IBM XIV Storage System Redpaper
  www.redbooks.**ibm.com**/abstracts/redp5215.html?Open
- Systems x servers
  ibm.com/systems/x
- Compretimator Utility Version 1.5.2.1
  www-304.**ibm.com**/webapp/set2/sas/f/comprestimator/home.html
- Microsoft Exchange Server
  technet.microsoft.com/en-us/library/aa996058(v=exchg.150).aspx
- SQL Server Best Practices
  technet.microsoft.com/en-us/sqlserver/bb671430.aspx
- Hyper-V Getting Started Guide
  technet.microsoft.com/en-us/library/cc732470(WS.10).aspx
- Failover Clustering Overview
  technet.microsoft.com/en-us/library/hh831579.aspx

@IBMSystemsISVs

Trademarks and special notices

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 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.

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

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

INFINIBAND, InfiniBand Trade Association and the INFINIBAND design marks are trademarks and/or service marks of the INFINIBAND Trade Association.

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.