



IBM Adds Real Time Compression for XIV – Improves Capacity Efficiency, Retains Simplicity

Analyst: David Reine

Management Summary

Did you ever try to pack too much into a single suitcase so you wouldn't have to pay to check a second bag on an airplane? Clearly, your primary motivation was economic. Who likes paying more when you can pay less (or nothing)? The problem with this strategy is that there are airline limits to the allowed size and weight of a packed suitcase. It is hard to pack 75 pounds (34 kilograms) of stuff into a suitcase and still be able to lift it, which is why in the U.S. there now is a limit of 50 pounds. The same is true for volume; there are airline restrictions here as well. Then, of course, you only can shove so much "stuff" into a given amount of space. Of course, it just would be simpler if the airline allowed two free bags per person, as you wouldn't have to try to optimize to fit into a single suitcase, but life often isn't simple.

The issues of economics and simplicity often are entangled. This is true in the data center, where budgets often are constrained in many ways (by line item limits for equipment, software licensing, staff, energy, etc.). You want to do more for less cost, and making things simpler often is part of the solution, because visible complexity drives up costs. However, **the ultimate challenge for the data center likely is trying to put more stuff into less space.** The amount of data stored is growing rapidly and it seems destined to take up even more space, as the volume of stored data is growing faster than disks are getting fatter. In addition, more data is becoming mission-critical and thus being used more often (most of this new demand is due to analytics), circumventing the desire to move more data to less speedy and less costly storage. As a result, data center managers seem stuck with needing more containers to hold their growing mission-critical data which, of course, has significant impact on the IT budget, floor space, and energy consumption.

Back to the suitcase analogy. What if you could stuff 75 pounds into a suitcase sized to hold 50 or 25 pounds and still have it get to where it needs to go in the requisite time and without paying for extra baggage? That would be great, especially if you didn't have to jump through too many hoops to make this happen and it would even be better if you could use one of your existing suitcases. This is akin to trying to squeeze 75 petabytes of data into 50 or 25 petabytes of physical disk capacity. **If doing this compression is simple enough and without too much added cost (for the magic that makes this possible), deciding to do it would be a no-brainer,** whether you are a business line manager trying to meet business needs through analytics or a storage manager trying to do more with less budget.

That is what IBM has done with its high-performance, scale-out XIV storage product line. It accomplished this feat by adding real-time compression (RtC) to XIV. If you already have XIV, this really is a no-brainer, because evaluating, testing and using RtC is straightforward and simple, which should make the decision making a lot easier. Read on to learn more about the magic of RtC on XIV.

IN THIS ISSUE

➤ Dealing with Data Growth in the Enterprise.....	2
➤ Doing What Seemingly is Impossible ...	2
➤ IBM XIV Storage System Gen3	3
➤ How RtC Works.....	4
➤ Conclusion.....	5

Dealing with Data Growth in the Enterprise

As we are all well aware, today's enterprise data center has to store and deliver more information than ever before, and do some-to-much of it faster than before, as well. This demand has been created by applications with a need for near-immediate access to data to match the need for faster decision making. Today's enterprises are looking for answers to a question right now (while the need is greatest), not tomorrow or even in 20 minutes. At the same time, data center managers also need to identify and solve future problems before they become critical. The most glaring data center problem is being able to afford to keep, manage and deliver all of that new data that is being created each day. See Exhibit 1 (to the right) for a list of specific storage challenges.

Doing What Seemingly is Impossible

These specific challenges can be restated into three more generalized challenges:

1. **Doing more of what is hard to do or seemingly next-to-impossible** (i.e., storing and retrieving mission-critical data quickly).
2. **Doing more for less.**
3. **Doing both, without making the storage management challenges any more time consuming** (by keeping things simple).

Whether in a line of business organization or in the data center, you want more data sitting on highly responsive storage and you want to pay less for it per terabyte than you do today, all without any new operational burdens or penalties. Yes, this is like "having your cake and eating it, too!"

So how can you stuff more data into the same space and do so without sacrificing performance being delivered today and while delivering the same kind of performance for the data being "squeezed in"? That is the big question!

The answer is *intelligent compression*. Compression is not new. Many of us have been zipping or otherwise compressing files for decades. Traditional compression tends to have several problems in meeting the related requirements.

- **When you or an application seeks to compress data, you need to read something in (the data being compressed), compress it, and then write out the result.** Doing this takes time and, as a result, almost always takes

Exhibit 1 — Data Center Storage Challenges

- **Increase performance** to enable faster analytics and shorten the time for reaching valuable insight.
- **Reduce complexity** to lower operational expenses with easy to use management tools and, thus, enable the data center staff to dedicate more time to focus on valuable strategic priorities.
- **Deliver high reliability and business continuity** to satisfy the service level agreements (SLAs) to which the data center already has committed. This includes the need to improve network utilization for remote mirroring and data replication.
- **Reduce physical storage space** to help free up valuable data center floor space.
- **Lower the total cost of ownership (TCO) for storing your growing petabytes of data. Given that data growth is driven by business needs, this means lowering the TCO per petabyte stored.**

Source: The Clipper Group

place after-the-fact, that is, after the data is stored. In other words, it is not happening dynamically, as the data is first being stored. An approach that takes more time is not an option. (*Who has time to spare these days?*) It also means that you need space for both the compressed and uncompressed versions, at least temporarily, which defeats somewhat the goal of saving space.

- **When you or an application seek to use the data that you have compressed, you almost always have to first decompress it and then store the expanded result before you can use it.** This also takes more time, which you don't have. As noted above, it also means you need space for both the compressed and uncompressed versions, at least temporarily.
- **When you do compression, you want it invisible.** You don't want to have to change your processes or programs so that they become "compression aware" and, more importantly, you don't want the users of the data

aware of the extra steps (by their response times increasing). Unfortunately, this often is the case with many compression solutions.

- **If you do compression, you want your compression solution to be universally applicable** (rather than having many solutions for many applications or data types). While there are some applications that can operate on compressed data, this tends to be true only for specific types of data that the application controls.
- **Before you do compression, you want to know whether it is going to be worthwhile.** With many compression solutions, you won't find out until after the targeted data has been compressed. A crystal ball would be nice; a guaranty of what is to come would be even sweeter.
- **When you decide to compress, you want the process of doing it to be simple – as simple as flipping a light switch.** With traditional compression solutions, this often is not the case. You don't want to have to hire another specialist or train someone to become one.
- **When you turn on compression, you don't want anyone to notice the compression that is going on in the background, either at initial ingestion of existing data or as new data is added or data is formed.** With traditional compression solutions, this usually is not the case.
- **When you turn on compression, you don't want your application servers burdened with new compression and decompression chores.** With many traditional compression solutions, your well-tuned and balanced application servers already have been deployed to make something happen. If you have to get more servers and/or spend resources to keep them in balance (because of the added burden of doing compression and decompression), you are reducing the budgetary savings.

What you need is a high-performing storage solution that offers invisible (both unaware users and no slow down), in-line (operating in

real-time) compression that is simple to administer and delivers significant economic savings.

Where are you going to get that? If you already have *IBM XIV Storage Systems*, then you can achieve all these goals while avoiding all of the headaches with a flip of a switch. If XIV is new to you, then it's time that you became aware of it and consider the possibilities.

IBM XIV Storage System Gen3

The *IBM XIV Storage System*, introduced several years ago, has a parallel, scale-out architecture that delivers consistently high performance with automated, balanced data placement across all key resources. You just keep adding modules to scale capacity and performance linearly.

XIV was designed for simplicity and reliability, providing data protection and availability along with ease-of-use, to lower the TCO of the IT infrastructure. With five 9s availability, and a full 6TB rebuild in typically less than an hour, XIV provides resiliency and, with ENERGY STAR certification, is one of very few enterprise storage systems endorsed as such for energy efficiency.

The minimum single XIV frame contains six modules with 72 disks, with each of the modules housing a dozen drives in 1TB, 2TB, 3TB, 4TB, or 6TB configurations. (All of the drives need to be the same.)¹ This minimum XIV frame provides the data center with a net (redundant) capacity of up to 169TB. A single XIV frame can grow to 15 modules with 180 disk drives. Each module contains an Intel CPU and up to 48GB of memory per module. Using *IBM Hyper-Scale* technologies, it is possible to manage up to 69.8 PB in an XIV Storage System, deployed across 144 frames with 15 modules in each, and managed through a single pane.

Clipper last looked at Generation 3 of the XIV technology in August 2011.² At that time, IBM announced a 20 times increase in XIV efficiency and performance with the introduction of *InfiniBand* and the latest Intel microprocessor technology (*Westmere*). Since then, over the past four years, IBM has devoted significant resources

¹ The number of drives directly determines the performance of the cluster. If you want more performance at the expense of more capacity, you should use smaller-capacity drives. If performance is somewhat less critical, then you should use larger drives.

² See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2011028.pdf) entitled *XIV Gen3 – IBM Lowers TCO and Raises Performance and Functionality*, dated August 5, 2011, and available at <http://www.clipper.com/research/TCG2011028.pdf>.

Exhibit 2 — XIV's Embedded Compression Prediction Tool

The screenshot shows a user interface for the XIV's Embedded Compression Prediction Tool. At the top, it displays 'P_Thin_1' and 'XIV 6011956a' with a progress bar at 47%. Below this, it indicates 'no-domain' and '516 GB Hard, 0 GB Saved (0%), 110 GB Potential saving (45%)'. The main part of the screenshot is a table with five rows, each representing a volume (V_T1_001 to V_T1_005). Each row shows the volume name, its current size, the size after compression, the potential saving percentage, and the final status (e.g., '54% Potential saving' or 'Uncompressed').

Volume	Current Size	Compressed Size	Potential Saving	Status
V_T1_001	103	86	54% Potential saving	54% Potential saving
V_T1_002	103	51	61% Potential saving	61% Potential saving
V_T1_003	120	51	63% Potential saving	63% Potential saving
V_T1_004	86	51	Uncompressed	Uncompressed
V_T1_005	120	0	Uncompressed	Uncompressed

Source: IBM

to continually improve XIV's architecture, most notably with SSD caching (to accelerate performance), Hyper-Scale mobility (to efficiently replicate or relocate data across the modules and frames), encryption, 3-site mirroring, and 6TB drives.

Now, IBM has announced the availability of *Real-Time Compression (RtC)*³ for XIV to simplify and accelerate the storage of more data in less space than ever before, delivering up to five times more capacity, usually without affecting throughput performance.

Along with the new RtC functionality, IBM also has added support for *Microsoft Azure Site Recovery* for simplified on-premise disaster recovery coordination. Azure Site Recovery simplifies the orchestration/management of replication, improving data protection.

How RtC Works

RtC uses IBM's patented *Random Access Compression Engine (RACE)*, the same technology found in IBM's *Storwize* and *SVC* platforms, to compress active primary data. More than 70 patents enable RACE technology to redefine how to make industry standard LZ compression operate in real-time and allow random access without affecting application performance. RtC is available on XIV Gen3 modules; no additional hardware is required.

Unlike other compression techniques, XIV RtC allows storage administrators to activate compression on a per volume basis, compressing data only when it makes sense. It can expand a 485TB rack of usable storage so that it can hold up to 2.4PB of real data, with the amount depending upon the data content. XIV compresses up to

80%, enabling the storing of up to five times as much data – that is, 400% more – in the same space.

XIV provides this level of grid-optimized, enterprise-grade compression and corresponding performance with a guarantee. Yes, IBM guarantees the compression and performance levels you will receive, based on results from using IBM's *Comprestimator Utility* and *Disk Magic* on your actual data. (See what this looks like in Exhibit 2 at the top of this page.) In fact, RtC on XIV is a no-risk opportunity, as IBM provides a 45-day try-and-buy trial program with no license required. After the trial period, a license is available on any newly acquired or existing XIV Gen3 platform.

As noted, RtC compresses data, including *DB2* and *Oracle* databases, and all other kinds of compressible data, inline, with no background processing required. See Exhibit 3, on the next page, for estimated savings by type of data stored. (Be careful, as your actual "mileage" may vary, but that is the beauty of the *Comprestimator Utility*, as you get to see what you are going to save before you compress it.)

RtC enables storage administrators to manage compression with no additional hardware and, as mentioned before, to scale capacity simply by adding modules. Deploying RtC enables the data center to reclaim much-needed capacity now, while compressing existing application data non-disruptively. In addition, RtC frees up bandwidth as you can replicate compressed data using less bandwidth. With a user-friendly GUI that contains an embedded compression prediction tool, the IT staff can make sound decisions, based upon an ongoing display of predicted/actual capacity savings.

Using this inline compression capability reduces the acquisition budget, saves rack space, reduces energy costs, and minimizes licensing

³ See [The Clipper Group Navigator](http://www.clipper.com/research/TCG2012013.pdf) entitled *Significantly Improving Storage Efficiency – IBM Delivers Real-time Compression on Unified Storage*, dated October 2, 2012, and available at <http://www.clipper.com/research/TCG2012013.pdf>.

**Exhibit 3 —
Examples of How Much More Data You Will Typically Get in the Same Space**

DB2 and Oracle databases		5x
Virtual Servers (VMware)	Linux virtual OSES	3.3x
	Windows virtual OSES	2x
Office	2003	2.5x
	2007 or later	1.3x
CAD/CAM		3.3x

Source: IBM

costs for software with charges based upon total physical storage capacity. In addition, IBM expects RtC on XIV to deliver the same five 9s of availability that is delivered on XIV without RtC. In short, compression will help the data center freeze storage growth (for a while, due to all of the freed-up space that can be reclaimed and repurposed) or, at the least, enable the enterprise to slow down additional storage purchases.

What is the difference between IBM RtC and other compression technologies? Well, for one thing, IBM RtC is applied to active primary data, thus enabling more data collections to be compressed than can be done with most other compression technologies. This significantly expands the range of data eligible for compression, providing increased benefits as a result of the increased savings from reclaimed storage at enterprise-level performance. Once switched on, RtC is applied immediately and it is easy to manage (i.e., requires little oversight). There is no requirement to schedule post-process compression, thus eliminating the requirement to reserve additional space for the uncompressed copy of the data. In addition, the data center can accurately predict these savings through IBM's Comprestimator – and the forecast compression results are guaranteed!

With RtC, the storage administrator can extend XIV system soft capacity beyond hard capacity to levels such as 140%, and with great confidence, since XIV compression ratios are predictable and stable.

Conclusion

If you already have deployed IBM's enterprise-class XIV Storage System, the decision to activate real-time compression should be a no-brainer. Turn on the trial and see what you might be able to save. If you have not deployed XIV – maybe now you have a good reason to take a closer look.



About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-third year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- **The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.**

About the Author

David Reine is a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, he joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- **Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial “123” when you hear the automated attendant.)**

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, Clipper Notes, The Clipper Group Calculator, and “*clipper.com*” are trademarks of The Clipper Group, Inc., and the clipper ship drawings, “*Navigating Information Technology Horizons*”, and “*teraproductivity*” are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosures

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.