Significantly Improving Storage Efficiency — IBM Delivers Real-time Compression on Unified Storage

Analyst: David Reine

Management Summary

When you go into the supermarket do you ever wonder how many different products are on the shelves? Do you wonder how they determine what goes on each shelf? All of the shelves are shared among many products. *What gets placed on the bottom shelf? What goes on the top?* One week your favorite item might be at eye level, and the next time there may be a sale item in that same spot. *How do they determine the value (such as return per foot of frontage) of each shelf?* The supermarket often positions high-margin (impulse buying) items close to related items; for example, premium buns, wraps and condiments at the deli counter, trying to get the greatest return from that positioning. The overall goal is to get the greatest return from all of the store's shelf space both premium and hard-to-reach, so this is an ever-changing proposition, with grocery stockers making regular adjustments to maximize the latest strategy and promotions. Maximizing sale of more profitable items is an important measure of *effectiveness*.

However, maximizing the use of shelf space is another important goal, and is a measure of *efficiency*. Any empty space is a lost opportunity. Thus, stockers move the cans and boxes around to keep the shelves as full as possible. Packaging with a smaller footprint increases the density, so that is one way of getting more product into the same space. Additionally, some food companies offer freeze-dried items or condensed items, such as instant coffee, soup-in-a-cup, condensed milk, concentrated juice, and powdered drinks. All you have to do is add water to the contents to reconstitute it. This also improves shelf efficiency plus there is less weight and volume to transport. As a good rule of thumb – the better the space efficiency, the better is the opportunity to maximize the return on the shelf space. In order to make a profit, the supermarket has to balance the goals of effectiveness and efficiency. Sometimes, the efficiency and the effectiveness are at odds on a cost per unit basis. However, if you can improve the efficiency without affecting the effectiveness, then it is a winning proposition.

Trying to fit more "product" into a given amount of space becomes even more significant as we transition from our personal grocery shopping to our professional IT lives, especially as we try to figure out what to do with all of the information that needs to be stored. When it comes to storing data, we find the same concepts of effectiveness and efficiency. Often, these are at odds with each other. Usually, when you make something *more effective* (think "higher qualities of service"), you do so by spending more per unit of storage. Conversely, when you make storage *more efficient* (costs less per unit of storage), you often do so by lessening the quality of service in one or more ways. **What we all seek is that elusive deal where we don't affect the quality of service, yet are able to store a lot more (without spending a lot more).**

This is exactly what IBM has done with its latest improvement to its *Storwize V7000 Unified*

mid-range storage system. By adding a new feature called *Real-time Compression*, you can store more (probably, a lot more) without having to buy any more hardware and without any impact on storage operations. The cost of this is a modest software licensing charge. To learn more about Real-time Compression, please read on.

The Quest for Storage Efficiency

The long-term storage of data deals with information that needs to be preserved and protected as a result of application requirements or enterprise, industry, or even governmental regulations and standards. The value of data to the business, however, may change over time, along with the urgency to access it, usually in terms of "how fast" or "how often".

- *How can we store it economically, yet maintain rapid availability when necessary?*
- *How can we reduce the amount of physical storage required to preserve and protect a given amount of data?*
- *How can we deploy storage more effectively in terms of different tiers of storage, yet also deploy it more efficiently in order to reduce the total cost of ownership (TCO)?*

First, we can separate this pile of bits and bytes into two basic categories: data that is "likely needed" and "less likely needed", loosely corresponding to *active data* and *archived data*. Active data may be needed at any moment and any significant delay (think "a second or two, at the maximum") will slow down the pace of business. While archived data may be needed at any moment, the expectations for instantaneous delivery are few, because it is realized that it is just one of thousands or millions (or more) of data that has been archived, usually in a manner that places cost minimization ahead of near-instantaneous retrieval speed.

Of course, this isn't a binary world most of the time, as there may be many levels of quality of service (platinum, gold, silver, etc.) and cost (from very high to very low). Optimally placing data at the best storage tier (where quality of service matches cost of delivery) is the name of this game. Whether you do this through automation or manual administration, let's assume that it gets done and done well.

While IT management may give you a good pat on the back, more than likely it will be followed by that stinger of a question "What have

you done for me lately?" This, of course is driven by the rapid growth of data outstripping the near stagnant storage budget. Wouldn't you like to be able to pull something out of your professional bag of tricks, besides just saying that you are waiting for the next refresh of technology to lower your average costs? Well, now you can. Exactly how you will do this depends on what storage now is sitting on your floor.

Introduction to IBM Storwize V7000 Unified

If you already have an *IBM Storwize V7000 Unified* storage system, you are way ahead in your quest to make something seemingly magical happen. But, let's assume that you haven't done that yet, and explain what Storwize V7000 has been delivering since it was announced in 2010. If you are familiar, just skip to the section "New to Storwize V7000 Unified" on the next page.

Unlike some vendors, IBM has a multiplicity of storage solutions for the data center. This includes the *DS3000*[1] family for the entry-level data center, up to the IBM *System Storage DS8000*[2], *XIV Storage System*[3], and the *Scale-Out Network Attached Storage* (*SONAS*)[4] for the largest enterprises. Most importantly for the mid-sized business or segregated outpost of a larger one, IBM has a variety of mid-range storage solutions, in the $25K to $250K range, including the *Storwize V7000 and Storwize V7000 storage systems*. IBM's customers do have many options and this allows them to choose one or more solutions that best fits *their* storage requirements rather than trying to force fit a single, "universal" solution into their data center, unless that is what is desired.

In an era of rampant data growth (with more to store), there are increasing requirements for

---

[1] See **The Clipper Group Navigator** entitled *Delivering Enterprise Features for the Mid-Range – IBM Introduces DS3500 with TPC MRE*, dated June 7, 2010, at http://www.clipper.com/research/TCG2010027.pdf.

[2] See **The Clipper Group Navigator** entitled *Maximizing the Business Value of Information and Lowering Energy Consumption with IBM's DS8700*, dated April 13, 2010, at http://www.clipper.com/research/TCG2010018.pdf.

[3] See **The Clipper Group Navigator** entitled *XIV Gen 3 – IBM Lowers TCO and Raises Performance and Functionality*, dated August 5, 2011, and available at http://www.clipper.com/research/TCG2011028.pdf.

[4] See **The Clipper Group Navigator** entitled *IBM SONAS – A New Kind of Infrastructure to Meet Large-Scale Storage Challenges*, dated March 15, 2010, at http://www.clipper.com/research/TCG2010010.pdf.

scalability at each tier. Simply put, an array selected must be able to scale to handle growing business requirements. However, the ability to scale is not the only requirement. Today, more so than ever before, enterprises of all sizes are looking for the storage management functionality once only found in the most expensive, high-end systems.

In an attempt to meet both the scalability and functional requirements, IBM Storwize V7000 was introduced in October 2010 as a multi-tiered, scalable mid-range storage solution with enterprise-class capabilities and outstanding performance.[5] It had outstanding power, flexibility, ease-of-use and functionality, in many ways similar to that of IBM's DS8000 enterprise-level systems. It had the following features and capabilities.

- **SSD devices** with very high IOPS – for database applications requiring high performance;
- **High-performance SAS** drives – to satisfy mission-critical Tier-1 application requirements;
- **High-capacity nearline SAS** drives –to fulfill the rapidly growing needs of Tier-2 business-critical applications, such as email and other web-facing requirements;
- **Easy-to-use graphical user interface** for storage administrators;
- **Easy Tier** – for automatic and dynamic migration between tiers; and
- **Virtualization** – to improve storage utilization rates, in much the same way that data centers have consolidated and virtualized their server infrastructure to improve the TCO of the IT environment.

Six months later, IBM added two-way clustering, 10Gbps iSCSI, additional drives, and improved *VMware* management. In 4Q11, IBM added unified block and file storage, greater flexibility for remote mirroring, enhanced *Tivoli* support, and even more new drive options.

It can have a range of storage capabilities (i.e., multiple tiers to satisfy multiple needs like active and archived data, to name just two), uses automation to move data between the tiers, is easy to administer, handles both file- and block-oriented access (that's what makes it "unified"), can virtualize legacy arrays behind it, and is

---

[5] See **The Clipper Group Navigator** entitled *IBM Brings Enterprise Functionality to Mid-Range Storage* dated October 7, 2010, and available at http://www.clipper.com/research/TCG2010047.pdf.

Exhibit 1 — Storwize V7000 and Storwize V7000 Unified Management Components and Capabilities

- Virtualized unified storage capable of enjoining external arrays of many types
- IBM *FlashCopy Manager*
- *Tivoli Storage Productivity Center*
- *Tivoli Storage Manager*
- *Tivoli Storage Manager FastBack*
- IBM *Systems Director*
- IBM *Active Cloud Engine*
- Block and file replication
- Asynchronous remote block mirroring
- Thin provisioning
- Storage performance monitoring
- Capacity planning and forecasting
- Real-time Compression [new]

Source: IBM

---

priced competitively for the mid-range. While that sounds like everything that you might want (and the list of capabilities is much longer, see Exhibit 1, above), remember that your boss is going to ask, sooner or later, for another round of improvements.

New to Storwize V7000 Unified

IBM has continued to improve the functionality offered in Storwize V7000, which is part of IBM's *Smarter Storage Strategy*. And now, **IBM has announced *Real-time Compression* for Storwize V7000 Unified. This is the surprise that you need to pull out of your bag of tricks.** There are other additions concurrent with compression. These include four-way clustering for even more scalability, support for FCoE, VMware support for file access, and even more software interoperability. These features are added to other innovations from the full spectrum of IBM storage products, including an innovative GUI from XIV, *Easy Tier* from the DS8000, and the *Active Cloud Engine* from SONAS. These all represent capabilities that distinguish Storwize V7000 Unified from its competitors. **While all of these are important, what you really want (right now) is the special effect of Real-time Compression.**

Real-time Compression

Now, with Real-time Compression, the data center can store more data without having to buy any more storage, simply by obtaining a license to turn on RTC, an expenditure with a quick

Exhibit 2 — Components and Capabilities

• DB2 and Oracle databases		• Up to 80%
• Virtual Servers (VMware)	• Linux virtual OSes	• Up to 70%
	• Windows virtual OSes	• Up to 50%
• Office	• 2003	• Up to 60%
	• 2007 or later	• Up to 20%
• CAD/CAM		• Up to 70%

Source: IBM

ROI. There is no additional hardware to buy and nothing to change; simply a license to procure.[6] If you already have Storwize V7000 or Storwize V7000 Unified, simply obtain a license and improve your physical storage capacity. If you have a block-based array other than Storwize V7000, consider putting it behind Storwize V7000 or the IBM SAN Volume Controller (which also can do Real-time Compression), enabling the IT staff to achieve compression on the non-Storwize-V7000 array and getting the advantages of common management and virtualization, plus Real-time Compression. Truly, this a triple benefit.

Storwize V7000 Unified takes advantage of IBM innovations introduced on its larger platforms, to enable the data center to store even more active and less active data, more efficiently, in the same unified space. **Real-time Compression works behind the scenes to compress and decompress your data, dynamically, in real-time, making compression, literally, a "no-brainer". No compromises need to be made between "hot", active data, and less active data.**

How Real-time Compression Works

Real-time Compression uses a time-tested and patented *Random Access Compression Engine* (*RACE*), used previously in IBM Real-time Compression Appliances. This technology, according to IBM, allows you to store up to five times more data for DB2 and Oracle databases, with significant savings for other kinds of data, too (as shown in Exhibit 2, above.) **By using storage more efficiently, the data center can slow the increase in capacity and reduce acquisition costs, rack space, energy consumption, and software costs for additional functions.** IBM has even made available a *Compresstimator* tool to evaluate the expected benefits in your environment. Your mileage will vary, depending on your mix and volumes of data, as you would expect.

Real-time Compression is an innovative, simplified compression that is fully integrated into Storwize V7000 and Storwize V7000 Unified to further improve their storage efficiency with a delivery compression of 50% or better. **It is a high-performance implementation that has been designed to support active, primary workloads**, in support of SSDs and high-performance SAS devices, **and also for less active data**, using high-capacity nearline SAS drives. **In other words, it won't slow down what you now are doing; it will just do it with less physical storage space.** (This sounds like magic but is no illusion!)

The inclusion of this dynamic compression capability reduces the acquisition budget, saves rack space, reduces energy costs, and minimizes licensing costs for software with charges based upon total physical storage capacity. In short, compression will help the data center freeze storage growth, or at the least, enable the enterprise to slow down additional storage purchases. If additional acquisitions are still required, a new Storwize V7000 Unified system will provide the data center with more new capacity at a lower

---

[6] Of course, if you don't already have Storwize V7000, you'll have to get one, but with the space-saving benefits of compression, it shouldn't be hard to convince your boss.

Exhibit 3 — Traditional Compression

- Traditional compression engines take a fixed input and produce a variable output depending on compressibility
- Some or all data written is stored uncompressed and only compressed later
  – Reduces efficiency because additional storage required
- Compression ratios depend on chunk size used
  – Big chunks can have poor I/O performance
  – Small chunks offer lower compression ratios
  – Compression typically degenerates over time creating fragmentation requiring more garbage collection
  – Poorer I/O performance

Traditional Compression

Fixed input → Variable output

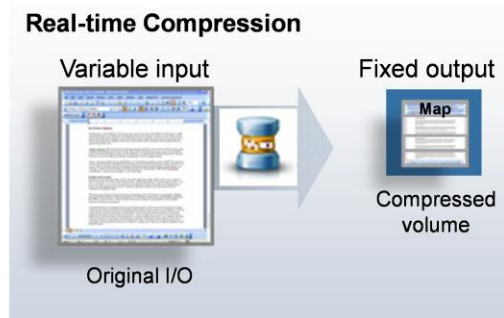Original volume → Compressed volume
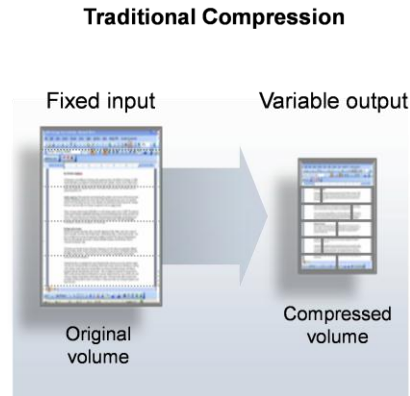
Source: IBM

cost per TB. The previous version of Storwize V7000 (without compression) already was an attractively priced mid-sized offering. Even with the added licensing list price of $9,000 per shelf of disk, Storwize V7000 Unified is now an even better bargain on a TCO/TB basis. Other vendors, such as EMC and NetApp, provide compression as well. Traditional compression engines take a fixed input and produce a variable output depending on compressibility (See Exhibit 3, above). Active data is first stored uncompressed and compressed later through a pre-scheduled post process activity (i.e., it is not compressed in real-time, as it is first being stored).

Traditional compression ratios depend upon the size of each "chunk"; while larger chunks can have a poor I/O performance and higher

Exhibit 4 — IBM's "Compression Without Compromise"
Variable Sized Compression Enhances Performance and Efficiency

- RACE flips this approach, taking a variable data stream and producing a "fixed" output
  – Compressed volume has a consistent logical layout
  – Temporal locality: data that's accessed together is compressed together
  – Variable sized input chunks get better compression
  – Less data is compressed and decompressed
  – Fewer disk I/Os
  – Better performance

1 MB "chunk" → 100 byte update

Real-time Compression

Variable input → Fixed output

Original I/O → Map → Compressed volume

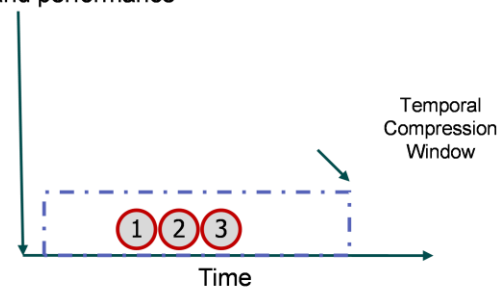|	Traditional Compression	IBM Real-time Compression
	1 MB Read	0 MB Read
	1 MB Decompress	0 MB Decompress
	100 Byte Update	0 Byte Update
	1 MB Compress	100 Byte Compress
	1 MB Write	<100 Byte Write
	Total I/O: 2 MB	**Total I/O: <100 Bytes**

Source: IBM

Exhibit 5 — IBM's "Compression Without Compromise" Explained
Designed for Real-Time Random Access to Active Data
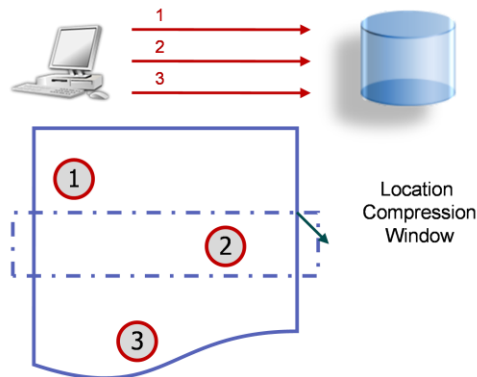
- Applications make multiple updates to data
- Traditional compression uses fixed-sized chunks and compresses each update based on its physical location on a volume
- In this example, **three** separate compression actions

- RACE compression acts on data that is written around the same time ("temporal locality") not according to physical location
- Temporal locality is more related to real system operation
  - Applications may make related updates to different parts of a volume
- RACE takes advantage of the structure of the data and its relations
- In this example, only **one** compression action
- Better compression and decompression efficiency and performance

Source: IBM

ratios; smaller chunks offer lower compression ratios but better performance. In fact, traditional compression methods typically degenerate over time, creating fragmentation requiring even more administrative clean-up. On the other hand, IBM's Real-time Compression is a dynamic procedure, operating immediately, in an unobtrusive, easy-to-manage process.

What makes IBM's implementation superior?

Using RACE, IBM provides compression for both active and inactive data, not simply compressing secondary copies of the data, as do some compression schemes. **Real-time Compression always is active, there is no need to schedule a post-process compression task, freeing up administrative personnel to be doing more productive work.** Real-time Compression provides a high performance engine to support workloads that other solutions cannot, significantly expanding the amount and type of data that is available to be compressed. RACE takes a variable-length data stream and produces a fixed output, significantly improving storage efficiency. (See Exhibit 4, at the bottom of the previous page.)

RACE compressed data has a consistent, logical layout, with data that's accessed together

being compressed together. With fixed-sized chunks to store, Storwize V7000 Unified can get better compression, with less data being compressed and decompressed when accessing a specific piece. As a result, there are fewer disk I/Os and, thus, better performance. (See Exhibit 5, above.)

Real-time Compression eliminates the requirement to reserve additional buffers for uncompressed data in transit. It supports all Storwize V7000 and Storwize V7000 Unified devices, including internal or external virtualized storage, enhancing the value of the data center's existing storage assets. Real-time Compression has been integrated nicely into the Storwize V7000 GUI to enhance its usability, giving IT staff more information about the compression performance and savings being achieved.

Existing Storwize V7000 volumes can be compressed easily with volume mirroring, eliminating unused space during the conversion. In terms of scalability, Real-time Compression currently supports up to 200 compressed volumes for each Storwize V7000 control enclosure.[7]

---

[7] IBM has indicated that the Storwize V7000 will support more than 200 volumes in the near future.

Conclusion

Why should your data center deploy Storwize V7000 or Storwize V7000 Unified with Real-time Compression? From an economic standpoint, these systems provide the enterprise with the financial benefits it needs to meet the growing storage requirements within the limits of a constrained budget. **If you already own a block-only Storwize V7000 or Storwize V7000 Unified, all you need is a license for Real-time Compression and you can take advantage of more physical storage without significant effect on your budget. If you own other block storage, you can put that storage behind a new Storwize V7000 Unified system and receive the same, space-saving benefits.**

Real-time Compression provides the data center with the flexibility to increase capacity and maintain performance, for all kinds of data, including your most active data. *With the same storage budget*, the data center can expect to see a reduction in per usable TB storage costs for most common configurations in the order of 30-to-40%, according to IBM. *To maintain existing usable capacity*, the data center can expect to see a reduction in cost per usable TB of approximately 25-to-35 percent. By applying compression to existing storage, the data center might be able to double usable capacity for a very modest incremental investment.

If your data center needs to compress both block and NAS storage, Storwize V7000 Unified provides the single platform that you need to simplify deployment. If you are suffering from too much data with an insufficient budget, Storwize V7000 Unified with Real-time Compression may well be the solution you need. It will be like getting more shelf space without enlarging the grocery store and without paying the going price for additional shelves. Isn't that the kind of surprise you want to pull out of your bag of tricks the next time you are asked? Check it out!

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twentieth 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 25 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, 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.