Lab Review

Compression without Compromise: IBM XIV Model 314 with Real-time Compression

Date: May 2016 **Author:** Vinny Choinski, Senior Lab Analyst

Abstract: This ESG Lab Review documents hands-on testing and auditing of the IBM XIV storage solution. Testing focused on how IBM integrated Real-time Compression into the XIV grid-architecture to deliver consistent performance with compression for real-world workloads.

Background

If rapid data growth was not negatively impacting storage budgets by consuming valuable space at unanticipated rates, IT organizations would not be putting so much effort into efficient data storage practices that help with cost containment. Cost is clearly a key consideration for IT organizations, with more than half (53%) of ESG research respondents citing capital (i.e., hardware) *and/or* operational expenses—including both staff and power and cooling costs—as a storage challenge.[1] And cost management initiatives often drive demand for complex data migration, management, optimization, and automation of data placement solutions that help alleviate capacity demands on expensive primary storage.

FIGURE 1. Top Ten Storage Environment Challenges

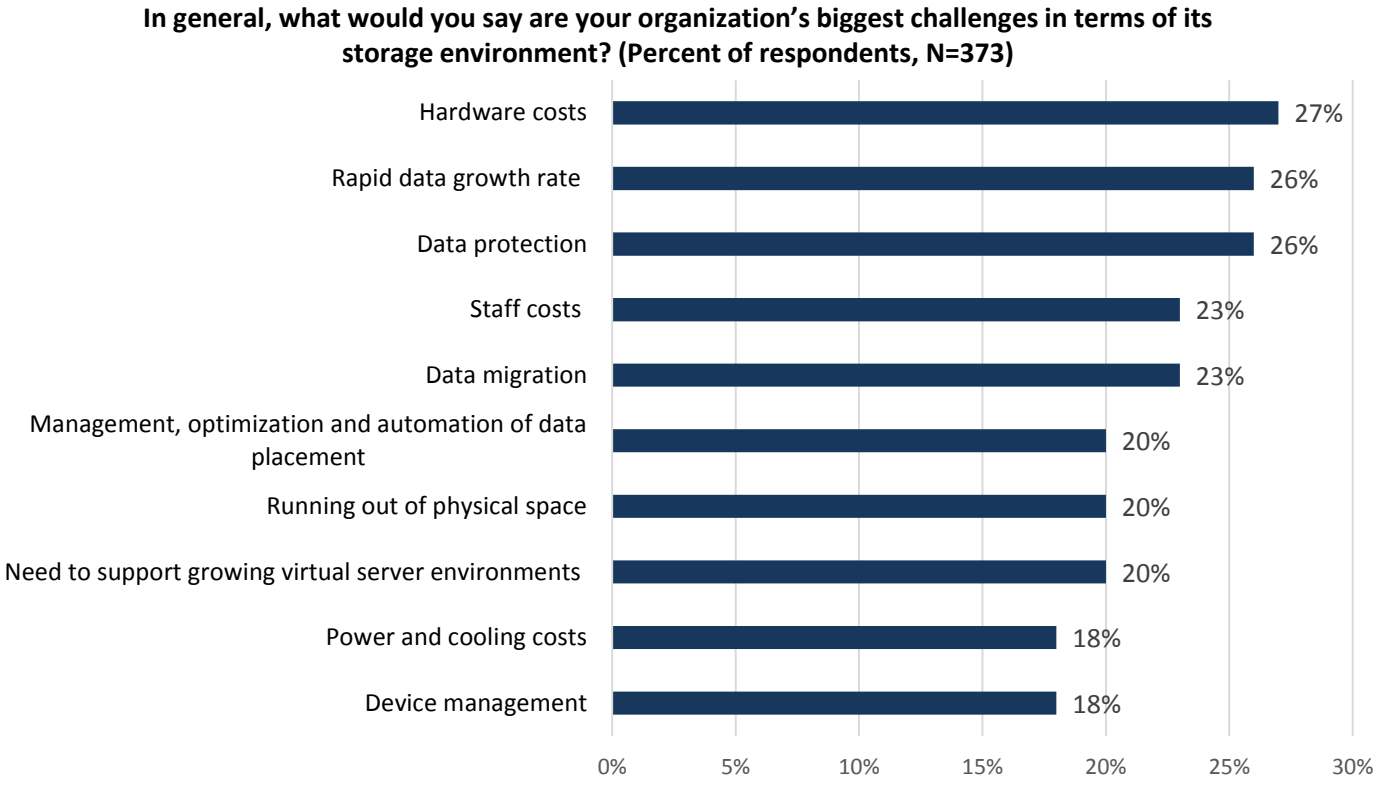

Source: Enterprise Strategy Group, 2016

[1] Source: ESG Research Report, *2015 Data Storage Market Trends*, October 2015.

It's apparent through the introduction of Real-time Compression in the latest XIV storage solution that these challenges have not been ignored by IBM. Now XIV customers can easily identify areas of capacity savings and apply compression directly on their primary storage, eliminating the need for complex bolt-on primary data management solutions.

The Solution: IBM XIV with Real-time Compression

The IBM XIV Storage System is a grid-scale enterprise data storage solution designed to deliver out-of-the-box performance predictability, high resiliency, management simplicity, and exceptional data economics. XIV uses a distributed design of interconnected modules, each with CPU, cache, optional flash cache, and high-density disk drives, working in parallel to serve data. Data is stored automatically and uniformly across all disks in mirrored 1MB partitions. A flexible cache implementation allows XIV to leverage large slots for reads while managing a smaller slot size, resulting in a superior cache hit ratio for better performance. Key features include:

- Real-time Compression (RtC), which uses the field-proven IBM Random Access Compression Engine (RACE) technology to compress active primary data. It is integrated with the existing XIV software stack and the compression workload is spread across the entire XIV grid for performance. With RtC, volumes can be non-disruptively compressed or decompressed online.
- Tuning-free high performance, which leverages XIV multi-module parallel processing, distributed data, and distributed cache to deliver tuning-free high performance, even during peak load, maintenance, and disk recovery periods.
- Exceptional security and reliability, delivered through the XIV five nines architecture design, three-way mirroring, self-healing data copies, and less than one hour drive rebuild times.
- Ease of use, through the user-friendly, intuitive GUI with easy provisioning, migration, and hyperscale management.
- Cloud agility: XIV delivers a flexible consumption model that supports hyperscale and multi-tenant environments as well as robust automation and orchestration through OpenStack, RESTful API, and VMware.

Figure 2 provides an overview of the solution. It shows a single XIV chassis fully populated with 15 modules. The blue and yellow blocks represent compressed (yellow) and uncompressed (blue) distributed data that make up the compressed and uncompressed volumes, as demonstrated by icons to the right of the XIV chassis. On the left side of Figure 2 are the host servers. The XIV supports FC and iSCSI host connectivity via ports on modules six through nine.

FIGURE 2. Solution Overview

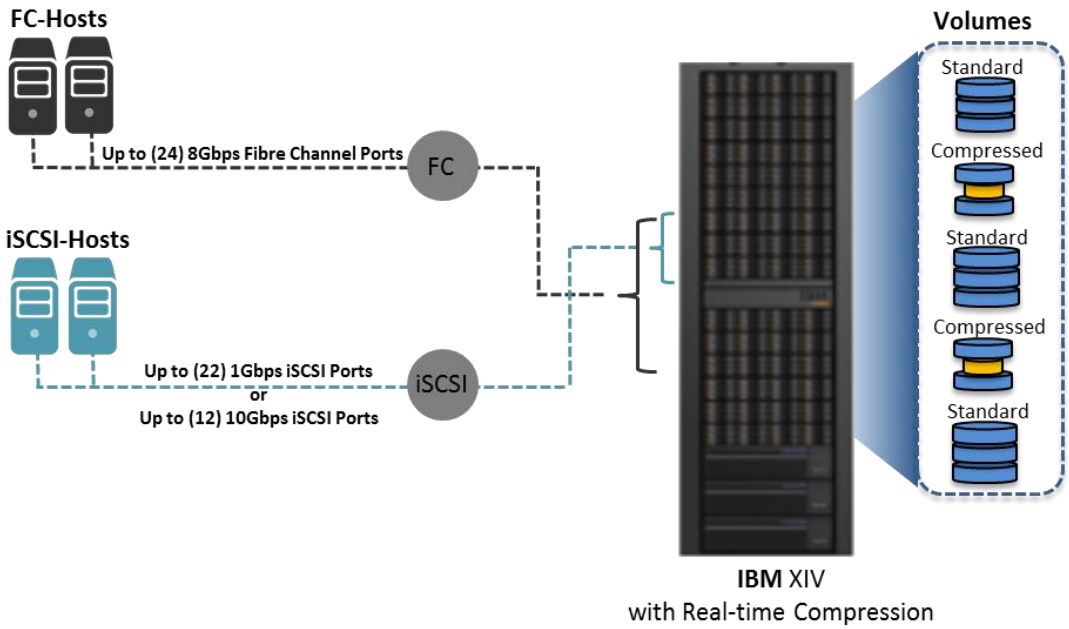

Source: Enterprise Strategy Group, 2016

ESG Lab Testing

ESG Lab performed hands-on testing and auditing of the IBM XIV storage system with Real-time Compression. Testing focused on the integration of RtC into the XIV software stack, ease of use for the new compression features, and compressed storage performance.

Compression Integration/Ease of Use

ESG Lab started our XIV testing with a review of its architecture. The upper left side of Figure 3 shows a fully populated XIV chassis with 15 interconnected nodes that create the XIV storage grid. To the right of the chassis is an expanded, block-diagram view of the nodes. Each node has its own CPU, memory, disk, and network resources. The nodes are interconnected with redundant 10Gb Ethernet switches that sit in the middle of the configuration. The nodes work in parallel to deliver enterprise-class storage to servers and applications. The three nodes just below and just above the network switches also provide host connectivity. The solution supports FC and iSCSI host connections. The latest version of the XIV storage solution includes IBM's Real-time Compression, which is distributed across the entire grid, as shown by the blue boxes in the center of each node that represent the RACE technology code inside XIV version 11.6.

FIGURE 3. XIV with Real-time Compression Architecture Overview

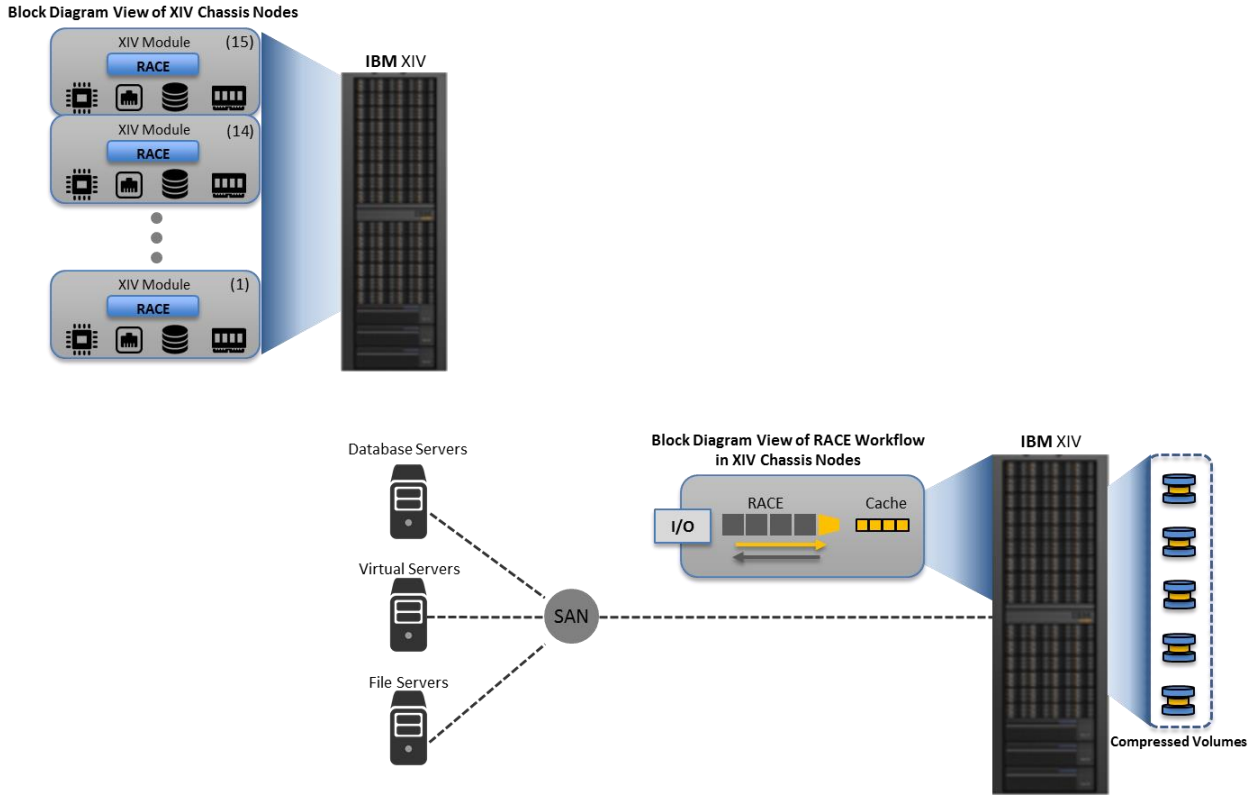

Source: Enterprise Strategy Group, 2016

The bottom of Figure 3 shows the data compression workflow. To the left are FC- and iSCSI-connected hosts. The center shows the RACE technology and cache in an expanded block-diagram view of an XIV node. To the right is the XIV chassis and its compressed volumes. Compression is done between I/O controller and the cache. The gray blocks represent decompressed data being sent to the hosts, the yellow compressed data stored on the XIV. Compressed data is first sent to cache, then to disk. This process improves cache efficiency by enabling more data to be stored in the same amount of physical memory and performance by increasing the probability of cache hits. The process also makes it more likely that related data will be stored in the same area when it is written to disk, which helps reduce latency of future reads.

Next, ESG Lab explored the integration of Real-time Compression features into the XIV GUI. We found the features to be well integrated into the already intuitive interface and, as shown in Figure 4, very easy to use. With XIV version 11.6, both RtC, and a highly accurate compression estimator utility called Comprestimator were added to the XIV software. The utility runs in the background and, as it finishes scanning volumes, presents potential compression savings throughout the GUI. Users can then decide if it makes sense to enable compression for their specific storage environment. The bottom left of Figure 4 shows typical compression rates by data types.

Finally, to take the process full cycle, we applied compression to a specific volume. The process was easy—we simply highlighted the volume we wanted to compress, clicked to display an options menu, and selected ***Compress***. The process ran in the background, compressed data to a temporary volume, and seamlessly flipped over to a compressed volume when the process completed. Compression and decompression are non-disruptive and can be done while the volume is online. It should be noted that Real-time Compression does require a license, but customers have the option to test it at no cost in their own storage environment by taking advantage of a 45-day trial program.

FIGURE 4. Real-time Compression XIV GUI Integration

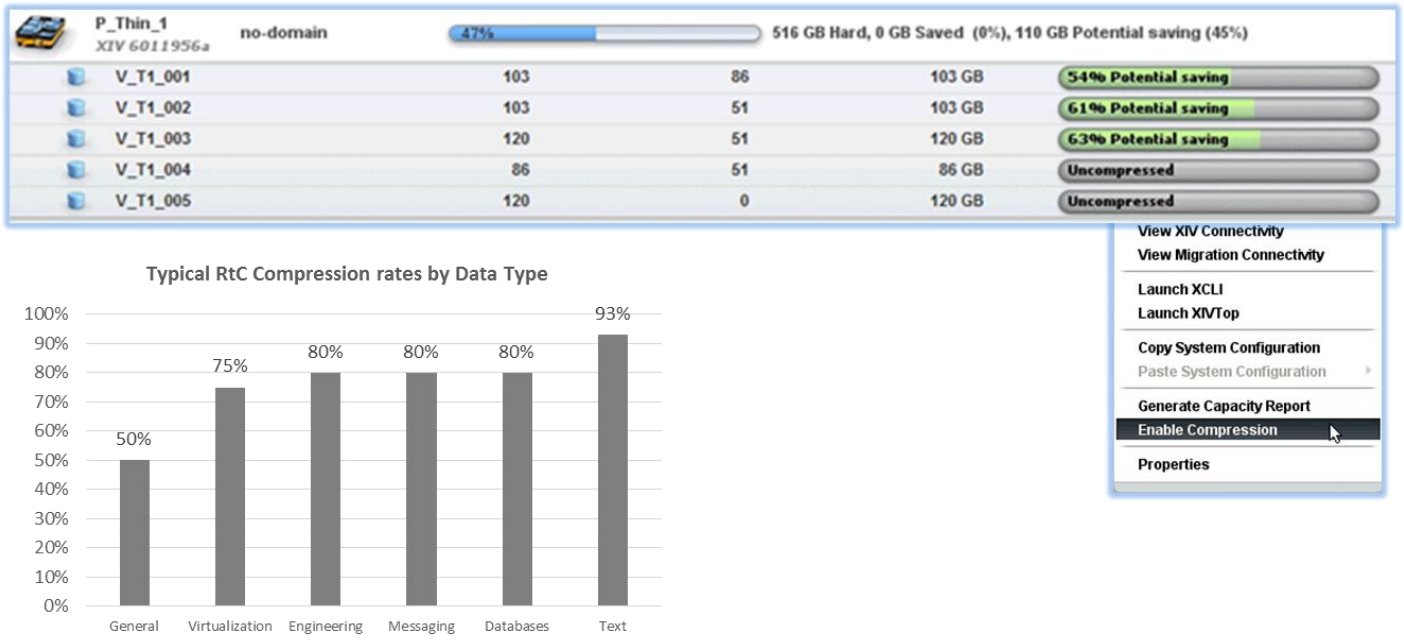

Source: Enterprise Strategy Group, 2016

Why This Matters

It not unusual to find unused captive storage in IT environments. Hasty deployments, application over-provisioning, and stale lifecycle management are a few of the culprits. One might expect that reclaiming space once it has been identified would be simple, but in actual practice it often involves lengthy migrations, upgrades, and even complex re-configurations.

ESG Lab has confirmed that, for XIV customers, IBM has made it easy not only to reclaim space, but also to actually create it. By seamlessly integrating RtC and its compression estimator into the XIV grid architecture, IBM enables customers to identify compression savings and compress data volumes online, without disruption.

Compressed Performance

To demonstrate the performance capabilities of XIV with Real-time Compression, ESG Lab, with support from the IBM XIV performance team, ran and analyzed the results of three categories of tests. The three test categories were designed to demonstrate maximum throughput, database transactions, and IOPS on compressed data sets in our XIV test environment.

We started our performance testing focused on the throughput capabilities of the solution. For this series of tests, the environment consisted of a single XIV chassis and a test volume previsioned on 180 4TB disk drives. The open-source, industry-standard vdbench utility was used to generate the workload and the host configuration included four IBM x3550 servers that were 8Gb SAN attached to the XIV storage.

The workload was configured on all four x3350 hosts with four threads per client. The vdbench utility was configured to run 3 TB of writes and 1 TB of reads with a block size of 64K. As shown in Figure 5, the test was repeated several times with different levels of compression. Figure 5 shows write and read throughput for data sets with 50%, 65%, and 91% compressed ratios. The tests were run to demonstrate solution performance with compression.

FIGURE 5. Compressed Throughput Performance

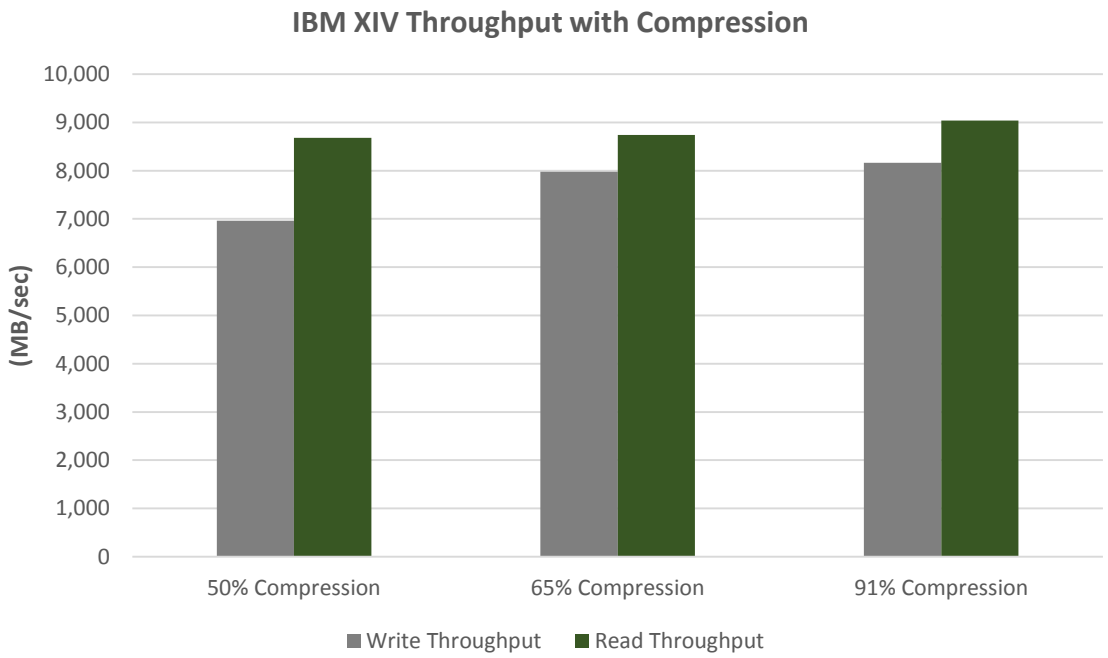

Source: Enterprise Strategy Group, 2016

What the Numbers Mean

- Performance actually increased as compression was added to the XIV test environment.
- Write throughput increased by 17% as compression was increased from 50% to 91%.
- Read throughput increased by 4% as compression was increased from 50% to 91%.

Next, ESG Lab explored the ability of the solution to deliver high performance with compression in an online transaction processing (OLTP) database environment. Testing was conducted using an industry-standard order entry test harness. Testing simulated a number of OLTP tasks, including new order, stock level, order status, delivery, and payment processing. To make testing more real-world and to match field-verified compression ratios, IBM conducted data betterment by loading the database with simulated, random, but real data such as first names, last names, and street names.

The OLTP database test environment consisted of a single XIV chassis that was 8Gb FC connected to a host server running an Oracle database. XIV storage was provisioned to support a 1.2TB database on a volume that delivered an 80% data compression ratio for the test data set via IBM Real-time Compression. The test database was Oracle version 11.2.03 configured with 10 GB of SGA.

The OLTP testing started with 1,000 users and scaled to over 32,000 users in our environment under test. As shown in Figure 6, the solution scaled in a near-linear fashion as we reached the maximum number of transactions with 32,000 simulated database users.

FIGURE 6. Compressed Database Transaction Performance

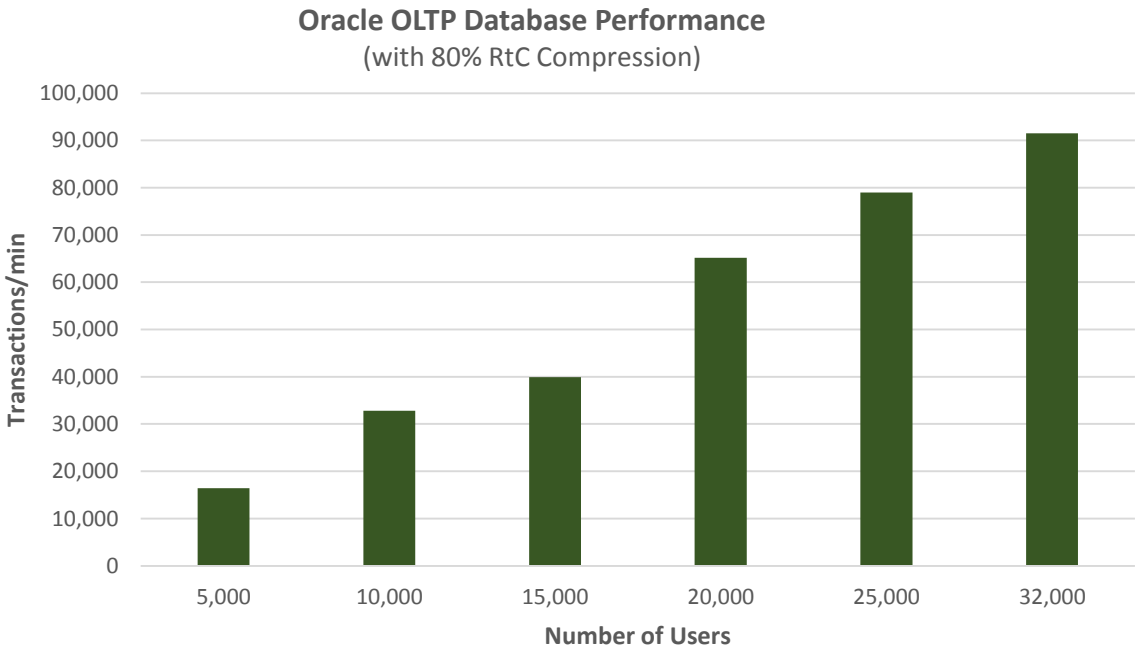

Source: Enterprise Strategy Group, 2016

What the Numbers Mean

- The solution scaled in a near-linear fashion as we increased the number of database users.
- We were able to scale from 1,000 to 32,000 before we reached the transaction limit.
- The solution scaled transactions from 3,254 transactions per minute with 1,000 users to a max of 91,521 transactions per minute with 32,000 users.
- XIV disk latency was only 11ms when we reached the maximum number of transactions.

Finally, ESG Lab reviewed the compressed IOPS performance capabilities of the solution. For these tests, we leveraged the same client configuration used during the vdbench testing. We used an internal IBM tool similar to the industry-standard iometer utility to demonstrate the variability between workloads that demonstrate zero temporal locality and those that demonstrate 100%. For IOPS testing, temporal locality refers to data that is written at approximately the same time and might be part of the same compressed stream. This is the kind of data that benefits significantly from the XIV with Real-time Compression schema.

Figure 7 shows the results of our IOPS performance testing. The right side of Figure 7 (4K Random Read) shows the biggest improvement in IOPS between zero and 100% locality—145,000 versus 565,000, respectively. Temporal locality of 100% means that data written at the same time is always part of the same data stream. Temporal locality of 0% means that data written at the same time is never part of the same data stream. It should be noted that even the worst case scenario was able to deliver 145,000 IOPS.

FIGURE 7. Compressed IOPS Performance

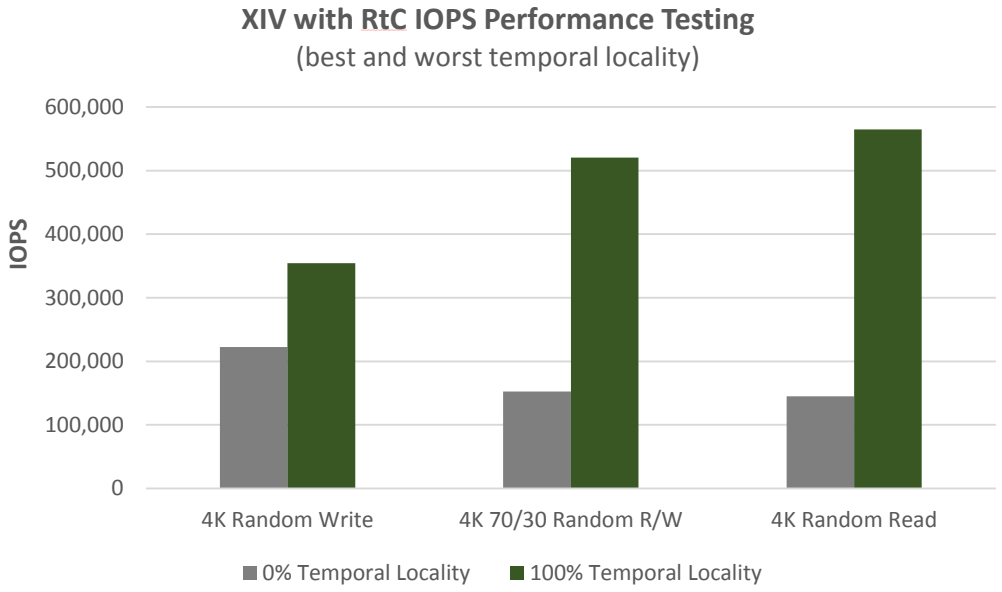

Source: Enterprise Strategy Group, 2016

Why This Matters

"Cost efficiency" and "high performance" are attributes that don't always go hand in hand when describing a single storage solution. Typically there is a tradeoff between performance and savings: Organizations expect to have to choose between expensive, high-performance, production storage and cost-effective, high-density storage for less active data. IBM is working hard to change this paradigm.

ESG Lab validated that the IBM XIV storage solution with grid-optimized RtC was able to deliver performance on compressed volumes. We tested a number of workloads with different levels of compression and, depending on the data type, measured IOPS, throughput, and number of transactions. We found that the solution was able to sustain, and in some cases, even increase, performance as compression was introduced. With XIV, customers get high performance that is also cost-efficient.

The Bigger Truth

ESG recently asked survey respondents to identify their organizations' biggest challenges pertaining to their storage environments. With more than half (53%) of respondents citing capital and/or operational expenses as storage challenges, it clear cost is a key consideration. Additionally, 19% of organizations cited cost as their *primary* challenge, whether related to CapEx or OpEx. What's driving these spending trends? After hardware costs, the rapid growth of data was the most commonly cited storage challenge.[2]

IBM appears to be keenly aware of these challenges, and as a result, the company continues to improve efficiency across its storage portfolio. In fact, in the latest version of its XIV storage solution, grid-optimized, field-proven, Real-time Compression has been introduced to improve capacity efficiency. The solution delivers high-performance storage with predictable, consistent capacity savings that can reduce the active data footprint by up to 80% depending on data type.

ESG Lab validated that Real-time Compression is fully integrated into the XIV software stack and delivers compressed performance for real-world workloads. Leveraging XIV's inherent parallel processing design philosophy, RtC is distributed across the entire grid at the node level. Compression is implemented between the I/O interface and cache, with each node handling its portion of the compression process for a particular volume. This approach helps maximize performance by distributing the workload and by maximizing the potential for cache hits. As a result, for the vdbench performance testing conducted as part of this validation, we observed a 17% increase in write throughput and a 4% increase in read throughput as we increased compression of the test data set from 50% to 90%. Traditionally, data compression has been relegated to secondary or even tertiary storage environments because performance tends to suffer as compression is introduced for capacity efficiency. This is not that case with Real-time Compression on XIV.

Compression alone won't solve every storage and data-lifecycle management challenge. It does not eliminate the need for sound retention, protection, and deep archiving policies and practices. However, it's highly valuable when you can simply activate a feature on your existing storage that delivers significant capacity savings without a performance penalty. ESG confirmed that IBM makes it easy to deploy and use Real-time Compression in XIV storage environments. The built-in compression estimator tool delivers reliable visibility of compression savings across the environment, and compression can be non-disruptively added and removed on any volume with a couple of mouse clicks. If you would like to test it in your own XIV environment, ESG highly recommends taking advantage of the 45-day no-cost trial program.

[2] Source: ESG Research Report, *2015 Data Storage Market Trends*, October 2015.