

Highlights

- IBM Real-time Compression is fully integrated in IBM XIV Storage System Gen3 software release 11.6 and, when enabled, is fully transparent to applications, such as SAS
- SAS mixed analytics workload is tested in lab environment with IBM Real-time Compression on IBM XIV Gen3, IBM Power E880 and IBM Spectrum Scale
- Test results show that up to 90% compression is achieved for SAS mixed analytics workload file systems – SASDATA, SASWORK, and SASUTIL
- Compression is uniform across the volumes allocated under each file system
- Results showed no impact on the overall performance of the SAS mixed analytics workload when compression is enabled
- Similar processor utilization and I/O access patterns are seen with Real-time Compression enabled and disabled
- Real-time Compression significantly reduces storage capacity requirements of SAS workloads
- Real-time Compression enables customers to deploy more workloads on the same XIV Storage System thus saving disk space and data center space

SAS analytics with IBM Real-time Compression on IBM XIV Storage System

The objective of this test is to demonstrate compression ratios achieved by SAS workloads with IBM Real-time Compression on IBM® XIV® Storage System, and also study the compression savings and performance implications when compression is enabled at the XIV system.

System architecture and configuration

Figure 1 illustrates the architecture and configuration used for testing SAS with IBM Real-time Compression in the lab environment.

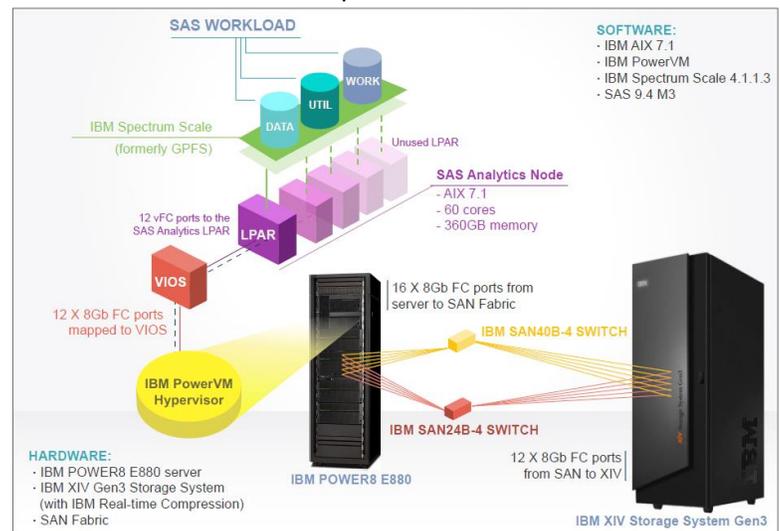


Figure 1: SAS deployment architecture on IBM Power E880 server with XIV Storage System and Spectrum Scale



Software

- SAS 9.4 M3 64-bit
- IBM AIX® 7.1 (7100-03-05-1524)
- VIOS 2.2.3.3
- IBM PowerVM® Enterprise Edition
- IBM Spectrum Scale™ (formerly IBM GPFS™) 4.1.1.3

IBM Power® System E880 server configuration

- Model: 9119-MHE
- Firmware version: IBM FW830.00 (SC830_048)
- Processor architecture: POWER8
- Clock speed: 4356 MHz
- SMT: OFF, 2, 4, 8 (SMT4 is default and it is used during the validation)
- Cores: 64 (60 cores for the logical partition (LPAR) under test, 2 cores for VIOS and 2 cores for other LPARs)
- Memory: 512 GB [360 GB for the LPAR under test, 8 GB for Virtual I/O Server (VIOS)]
- Internal drives: Twelve 600 GB (used for booting VIOS and LPARs)
- Fibre Channel (FC) connectivity: Four quad-port 8Gb FC ports (16 ports) attached to the server; used 12 ports during the testing

XIV Storage System configuration

- XIV system type: 2810
- System model: 114
- System version: 11.6.0.x (with IBM Real-time Compression)
- Drives: 180 SAS drives each with 2 TB capacity and 7200 rpm speed
- Usable capacity: 161 TB
- Modules: 15
- Cache: DDR3 360 GB
- Solid-state drive (SSD) cache: 6 TB
- Connectivity: Six 8Gb dual-port FC adapters (12 ports) connected to storage area network (SAN)
- Stripe size: 1 MB (default)
- SSD cache: Enabled (by default) for all volumes used in workload

SAN configuration

- Two FC switches connected to the XIV Gen3 and the IBM Power System; IBM System Storage® SAN24B-4 Express® (24 ports) and IBM System Storage SAN40B-4 (40 ports); both support NPIV
- Sixteen 8Gb dual-port FC ports connected from IBM Power System S222 server to the SAN fabric; eight ports connected to the first switch and eight more ports connected to the second switch
- Twelve 8Gb FC ports connected from SAN switches to XIV Gen3
- Switch zoning performed on the SAN switches such that each logical disk assigned from XIV to the LPAR has 24 paths.

Spectrum Scale file system

SAS workload contains three different file systems—SASWORK, SASDATA, and SASUTIL. The three file systems are created out of separate set of volumes as shown in the following list. The IBM Real-time compression is enabled or disabled on the volumes from the XIV GUI tool. Refer to IBM Real-time Compression on the XIV Storage System at <http://www.redbooks.ibm.com/abstracts/redp5215.html> for more information.

- SASWORK: 3.75 TB with 24 volumes of 150 GB each
- SASDATA: 4.8 TB with 16 volumes of 300 GB each
- SASUTIL: 2 TB with 16 volumes of 125 GB each

Recommended tuning

General SAS tuning guidelines are applied during the performance validation with IBM Real-time Compression. No specific tuning is performed for IBM Real-time Compression. Refer to the following white papers for more details on the tuning guidelines for SAS workloads on IBM Power Systems and XIV Storage System with IBM Spectrum Scale.

AIX tuning guidelines for deploying SAS

ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101529

SAS deployment on IBM POWER8 processor- based systems with IBM XIV and IBM FlashSystem

ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102515

IBM Spectrum Scale (formerly GPFS) tuning guidelines for deploying SAS on IBM Power servers

ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102255

Workload, test scenarios, and results

The workload used during the performance validation is SAS mixed analytics workload. The workload consists of a mix of analytics jobs that run concurrently. The jobs stress compute, memory and I/O capabilities of a given IT infrastructure.

The workload consists of 20 individual SAS tests: 10 compute-intensive, two memory-intensive, and eight I/O-intensive. Some of the tests rely on existing data stores and some tests rely on data generated during the test run. The tests are a mix of short-running (in minutes) and long-running (in hours) jobs. The tests are repeated to run concurrently and in a serial fashion to achieve a 30-session workload. The 30-session workload consists of 101 jobs. During the peak load, the 30-session workload runs 55 processor and I/O-intensive jobs concurrently.

The key performance metric is the compression ratio or compression savings achieved by IBM Real-time Compression for SAS mixed analytics workload file systems.

From the workload point of view, the performance metric is workload response time (in minutes), which is the cumulative real time of all the jobs in the workload. Lower response time is better. However, other performance metrics, such as processor time (user + system time), server utilization, and I/O throughput were also studied. These metrics were gathered to understand the impact on performance when compression is enabled.

The workload response time (real time) and processor time (user + system time) are captured from the log files of the SAS jobs. These statistics are logged with the SAS *fullstimer* option. IBM Power Systems™, starting with the IBM POWER7® processor architecture, use Processor Utilization Resource Register (PURR) accounting for accurate reporting of system usage. PURR factor for POWER8 processor needs to be applied to the processor time metrics described in the document. Refer to appendix B of *SAS business analytics deployment on IBM POWER8 processor- based systems with IBM XIV and IBM FlashSystem paper* referenced at the end of this paper, for more details on PURR factor.

Test scenarios

The following tests were run:

- SAS mixed analytics 30-session workload with IBM Real-time Compression disabled
- SAS mixed analytics 30-session workload with IBM Real-time Compression enabled

The testing was performed with no competing workload running on both server and storage systems.

The first test was run after tuning the stack for optimal performance. The second test was run after enabling Real-time Compression on all the volumes of the three SAS workload file systems (SASDATA, SASWORK and SASUTIL).

Compression ratios or compression savings were collected from the XIV GUI tool for the second test. The test team also collected the workload, host side and storage side performance metrics and compared them between baseline and compression tests.

Test results

In the lab setup, IBM Real-time Compression on XIV storage achieved very good compression savings for the SAS mixed analytics workload across all the three file systems. SASDATA and SASUTIL file system volumes are compressed by about 90% and WORK file system volume were compressed by 79%.

The following figures (Figure 2, Figure 3, and Figure 4) show the compression savings for SAS mixed analytics workload file systems SASDATA, SASWORK, and SASUTIL (respectively). In the lab setup, SASDATA, SASWORK, and SASUTIL are created on 16, 24, and 16 volumes respectively. All the volumes in a file system achieved uniform compression savings, as shown in the following figures.

Compression savings percentage is a simple calculation involving the size used by volumes before and after compression. For example, for SASDATA file system, as shown in Figure 2, each volume used 193 GB before the compression. And after the compression, each volume used only 18 GB, resulting in the capacity savings of 175 GB (193 – 18 = 175) per volume.

$$\text{Compression saving for SASDATA volumes} = \frac{(175/193) * 100}{100} = 90\%$$

Name	Size (GB)	Used (GB)	Size (Disk)	Consistency Group	Created (GUI Time)	Created On...	Compression Saving (%)	Compression Saving
E880_RTC_SASDATA_01	309	18	309 GB				90%	175 GB
E880_RTC_SASDATA_01.original	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_02	309	18	309 GB	The figure shows snippet of SASDATA volumes after compression. SASDATA filesystem has 16 volumes. The snippet is showing only 8 volumes. All the 16 volumes achieved uniform compression savings of 90% in the lab setup.			90%	175 GB
E880_RTC_SASDATA_02.original	309	193	309 GB		91% Potential saving	0 GB		
E880_RTC_SASDATA_03	309	18	309 GB		90%	175 GB		
E880_RTC_SASDATA_03.original	309	193	309 GB		91% Potential saving	0 GB		
E880_RTC_SASDATA_04	309	18	309 GB		90%	175 GB		
E880_RTC_SASDATA_04.original	309	193	309 GB		91% Potential saving	0 GB		
E880_RTC_SASDATA_05	309	18	309 GB		90%	175 GB		
E880_RTC_SASDATA_05.original	309	193	309 GB		91% Potential saving	0 GB		
E880_RTC_SASDATA_06	309	18	309 GB	XIV provides an option to save original volumes before compressing them. The snippet also shows corresponding original volumes.			90%	175 GB
E880_RTC_SASDATA_06.original	309	193	309 GB		90% Potential saving	0 GB		
E880_RTC_SASDATA_07	309	18	309 GB		90%	175 GB		
E880_RTC_SASDATA_07.original	309	193	309 GB		91% Potential saving	0 GB		
E880_RTC_SASDATA_08	309	18	309 GB				90%	175 GB
F880_RTC_SASDATA_08.original	309	193	309 GB				91% Potential saving	0 GB

Figure 2: IBM Real-time Compression savings for SASDATA file system volumes

Name	Size (GB)	Used (GB)	Size (Disk)	Consistency Group	Created (GUI Time)	Created On...	Compression Saving (%)	Compression Saving
E880_RTC_SASUTIL_01	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_01.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_02	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_02.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_03	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_03.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_04	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_04.original	137	18	137 GB				94% Potential saving	0 GB
E880_RTC_SASUTIL_05	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_05.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_06	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_06.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_07	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_07.original	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_08	137	2	137 GB				89%	17 GB
E880_RTC_SASUTIL_08.original	137	18	137 GB				94% Potential saving	0 GB

Figure 3: IBM Real-time Compression savings for SASUTIL file system volumes

Name	Size (GB)	Used (GB)	Size (Disk)	Consistency Group	Created (GUI Time)	Created On...	Compression Saving (%)	Compression Saving
E880_RTC_SASWORK_01	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_01.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_02	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_02.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_03	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_03.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_04	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_04.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_05	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_05.original	154	25	154 GB				85% Potential saving	0 GB
E880_RTC_SASWORK_06	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_06.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_07	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_07.original	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_08	154	5	154 GB				79%	20 GB
E880_RTC_SASWORK_08.original	154	25	154 GB				86% Potential saving	0 GB

Figure 4: IBM Real-time Compression savings for SASWORK file system volumes

The compression savings seen in the lab setup are specific to the test workload. The actual compression savings depend on the workload, data, and customer IT infrastructure setup at customer data centers.

Performance implications due to compression

During the performance validation by enabling IBM Real-time Compression on XIV Storage System in the lab setup, the overall SAS mixed analytics workload did not see any performance impact. The real time, user time, and system time metrics of the workload were similar for the tests with and without compression.

SAS mixed analytics 30-session workload performance		
	IBM Real-time Compression disabled	IBM Real-time Compression enabled
Real time (min)	1715	1510
User time (min)	1394	1397
System time (min)	50	50

Table 1: SAS mixed analytics workload performance comparison with and without IBM Real-time Compression

As shown in Table 1, the user and system time were similar for the workload with and without compression. The real time actually reduced with compression enabled because the random I/O test, ranrw, was serviced from solid-state drives (SSDs) within the XIV system, hence it finished much faster.

The server-side metrics also remained similar between the two tests. Because the compression and decompression is handled by the XIV Storage System completely, the server side metrics remained same, except for slight changes in the I/O access patterns. The processor utilization, memory utilization, and I/O metrics remained similar between the two tests, as shown in Figure 5 and Figure 6.

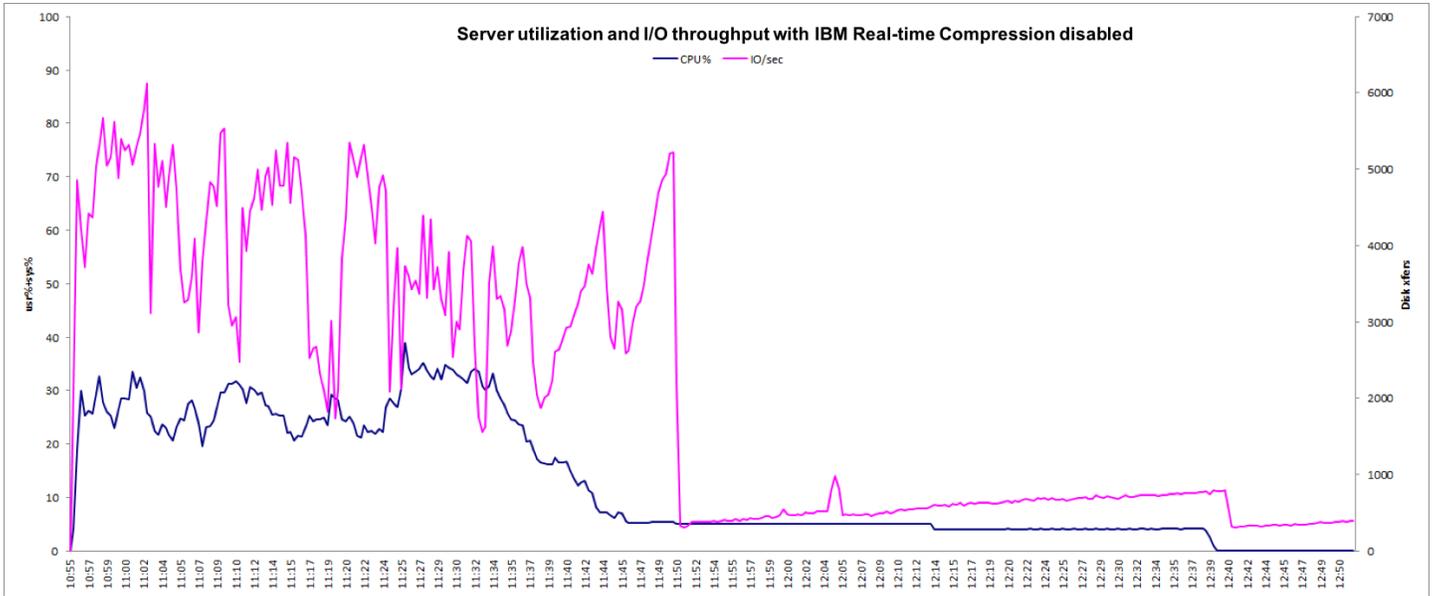


Figure 5: Processor utilization and I/O throughput at server with IBM Real-time Compression disabled

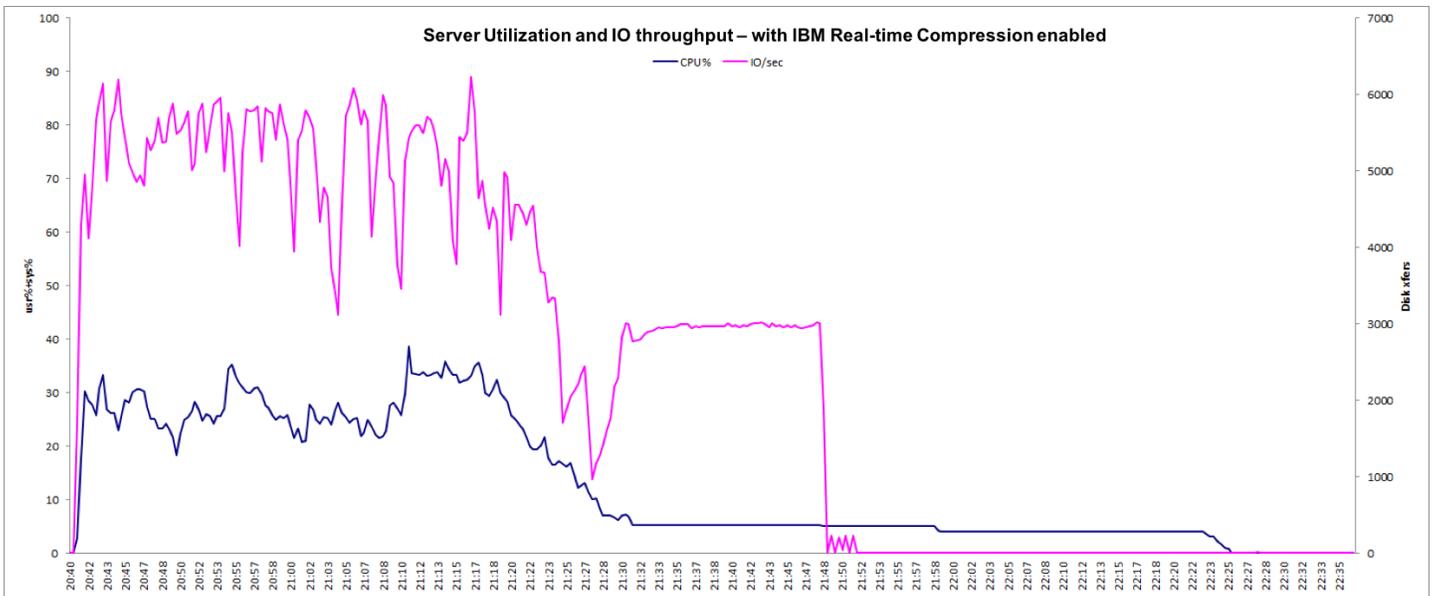


Figure 6: Processor utilization and I/O throughput at server with IBM Real-time Compression enabled

To summarize, during the performance testing with SAS mixed analytics workloads in the lab setup, IBM Real-time Compression achieved up to 90% compression rates, with no performance impact on the SAS jobs. The compression rates significantly reduce the storage capacity requirements for deploying SAS file systems - SASDATA, SASWORK, and SASUTIL

Formoreinformation

To learn more about the IBM Real-time Compression, IBM XIV Storage System or IBM Power Systems, contact your IBM representative or IBM Business Partner, or visit the following websites:

- IBM Real-time Compression on the IBM XIV Storage System ibm.com/redbooks/abstracts/redp5215.html
- IBM Real-time Compression with IBM XIV Storage System Model 314 ibm.com/redbooks/abstracts/redp5306.html
- AIX tuning guidelines for deploying SAS ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP101529
- SAS business analytics deployment on IBM POWER8 processor- based systems with IBM XIV and IBM FlashSystem ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102515
- IBM Spectrum Scale (formerly GPFS) tuning guidelines for deploying SAS on IBM Power servers ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102255

About the authors

Beth L. Hoffman is an IBM Executive IT Specialist and Big Data Solution Architect. She leads work with big data and analytics ISV solutions to enable, optimize, and differentiate them on IBM Power Systems and IBM storage systems. You can reach Beth at bethvh@us.ibm.com or connect with her on LinkedIn at: <http://www.linkedin.com/in/bethhoffmanibm>

Narayana Pattipati is an IBM Certified Expert IT Specialist and Big Data and Analytics technical consultant. He leads the technical enablement and optimization of big data and analytics ISV solutions on IBM Power Systems and IBM storage systems. You can reach Narayana npattipa@in.ibm.com or connect with him on LinkedIn at: www.linkedin.com/in/narayana-pattipati-0146736



© Copyright IBM Corporation 2016

IBM Corporation
IBM Systems
3039 Cornwallis Road
RTP, NC 27709

Produced in the United States of America
April 2016
All Rights Reserved

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 ibm.com/legal/copytrade.shtml

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



Please Recycle
