

# SAS Analytics with Real-time Compression on IBM XIV Recent Test Results Summary

Testing By Narayana Pattipati

Presented at SAS Global Forum by:

Harry Seifert

Executive IT Specialist; ISV Solution Architect for SAS  
seifert@us.ibm.com

Brian Porter

Technical Solutions Architect, IBM Systems ISV Enablement  
bporter1@us.ibm.com

---

Disclaimer: Individual results may vary depending on the customers unique mix of SAS workload. The test results described in this presentation used a Mixed Analytics test suite with workload/jobs provided by SAS Institute to exercise the server and storage.

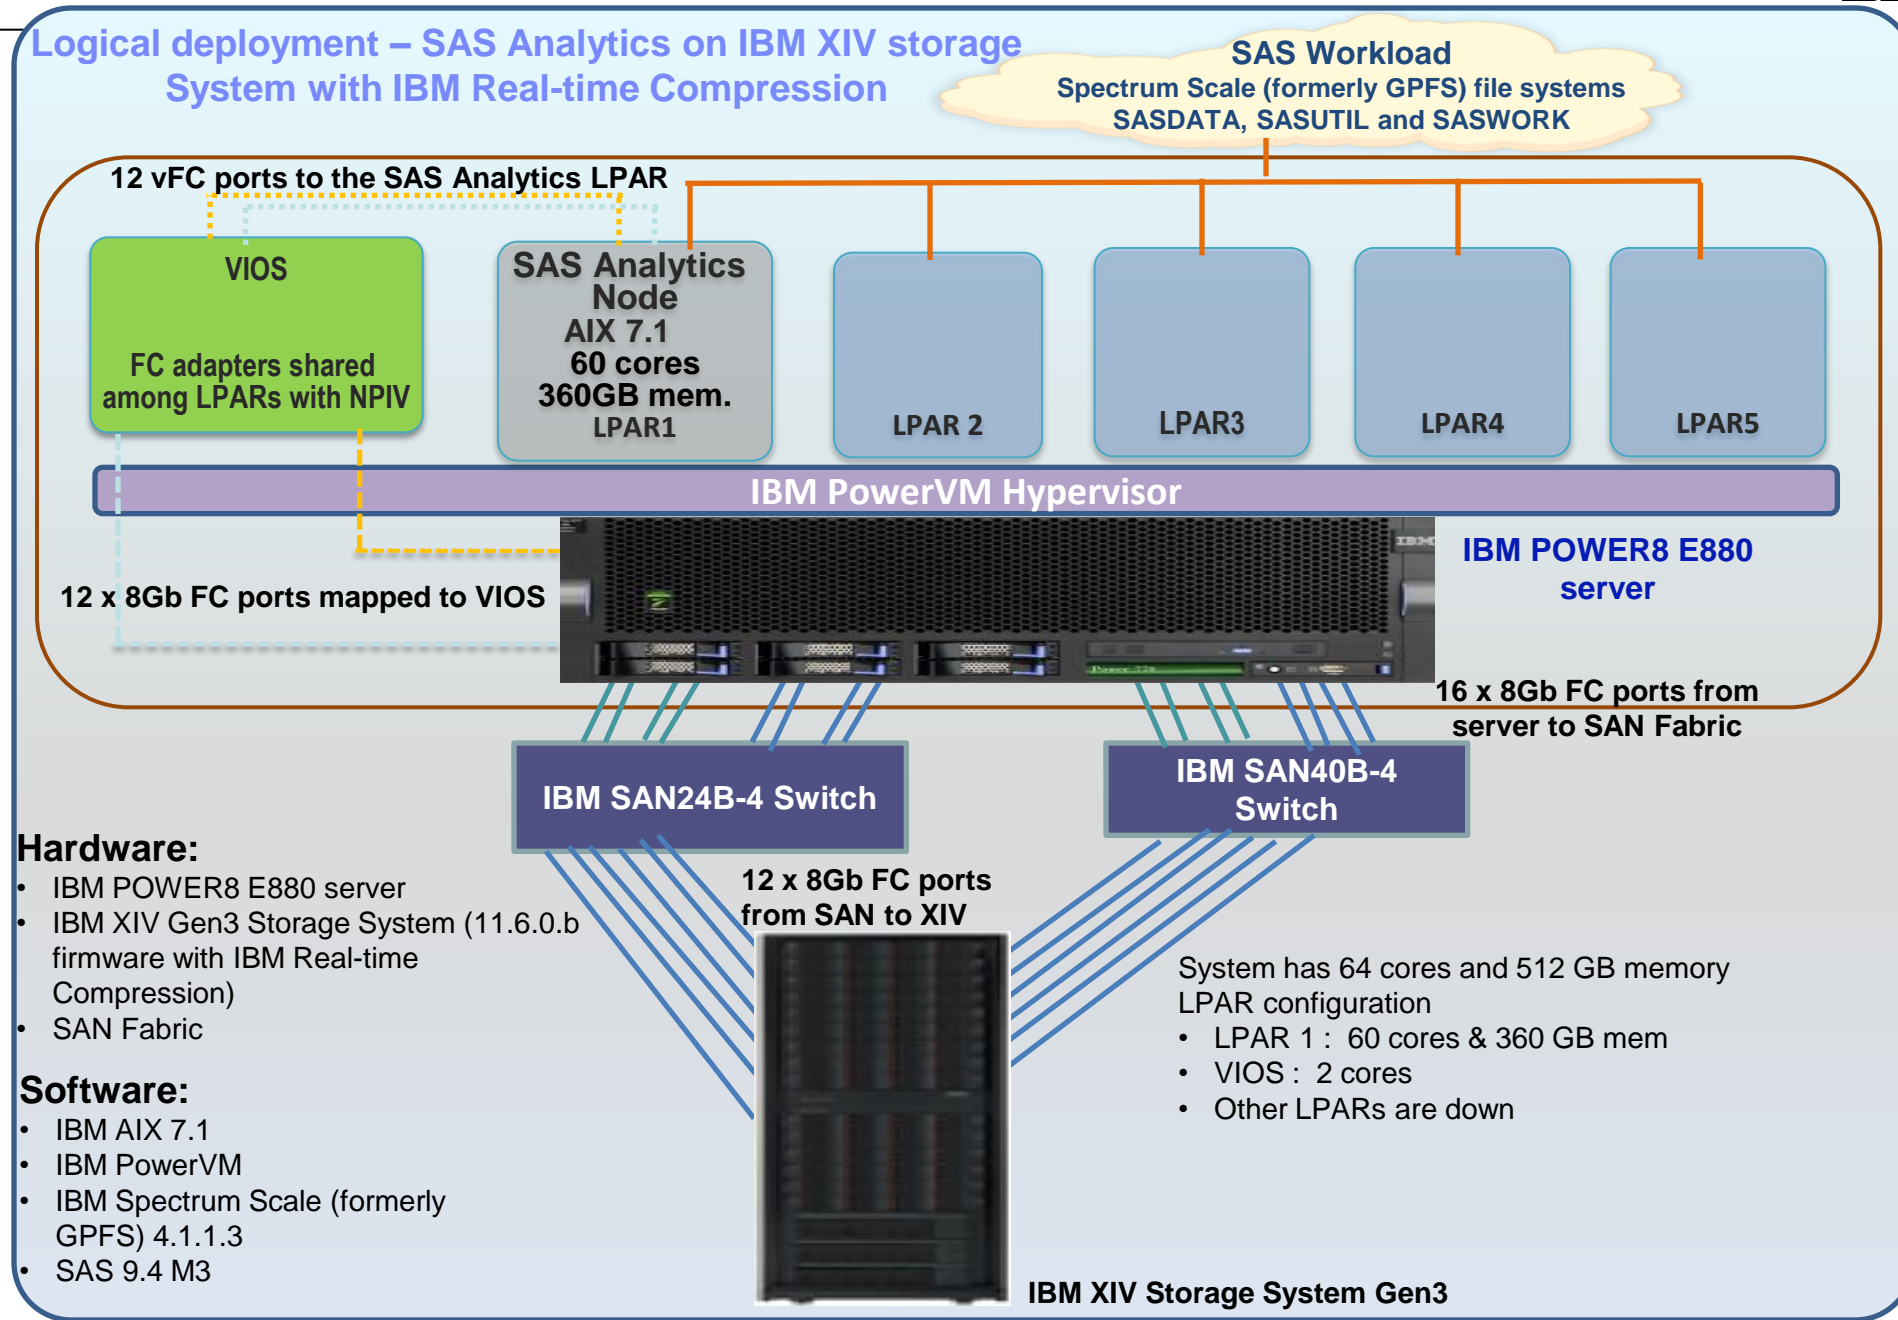
---

Who wants up to 90% data compression rates without noticeable impact on SAS mixed Analytics workload performance and response times?

IBM XIV Real-time Compression and Power8 with SAS Analytics is the answer!

- Fully integrated in IBM XIV Storage System
- Can significantly reduce storage capacity requirements of SAS Analytics workloads
- Comprestimator utility is easy and relatively accurate
- Up to 90% compression rates were realized with SAS Mixed Analytics test workloads
- No noticeable impact was observed with SAS Analytics workload performance and response times
- A newer XIV Storage System model (314) with double the RAM and CPU resources that was released after this test effort began is expected to provide additional RtC benefits to SAS workloads

# Test Environment Configuration for Non-Grid



- Tests

- Test1 - MA 30 session baseline
- Test2 - MA 30 session with XIV RTC

- SAS Analytics validated with XIV RTC Enabled

- Compression savings of 80% to 90% with the test workload
- Across SASDATA, SASWORK and SASUTIL file systems

- No performance impact seen when using RTC, with 30 sessions workload

- No impact on the performance parameters on the host side
  - MA 30 (Real Time, User time and System time)
  - CPU Utilization, IO and Disk and Memory
- No impact on the performance parameters on the storage side
- Random IO jobs finished faster due to SSD caching of random reads.

# XIV comprestimator tool estimation of space savings before RTC is enabled



Name	Size (GB)	Used (GB)	Size (Disk)	Consistency Group	Created (GUI Time)	Created On...	Compression Saving (%)	Compression Saving
						15,022 GB Hard, 3,397 GB Potential saving (88%)		
						26,104 GB Hard, 0 GB Saved (0%), 3,610 GB Potential saving (90%)		
E880Pool								
E880RTCPool								
E880_RTC_SASDATA_01	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_02	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_03	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_04	309	193	309 GB				90% Potential saving	0 GB
E880_RTC_SASDATA_05	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_06	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_07	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASDATA_08	309	193	309 GB				91% Potential saving	0 GB
E880_RTC_SASUTIL_01	137	18	137 GB				79% Potential saving	0 GB
E880_RTC_SASUTIL_02	137	18	137 GB				84% Potential saving	0 GB
E880_RTC_SASUTIL_03	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_04	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_05	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASUTIL_06	137	18	137 GB				94% Potential saving	0 GB
E880_RTC_SASUTIL_07	137	18	137 GB				94% Potential saving	0 GB
E880_RTC_SASUTIL_08	137	18	137 GB				93% Potential saving	0 GB
E880_RTC_SASWORK_01	154	25	154 GB				88% Potential saving	0 GB
E880_RTC_SASWORK_02	154	25	154 GB				88% Potential saving	0 GB
E880_RTC_SASWORK_03	154	25	154 GB				87% Potential saving	0 GB
E880_RTC_SASWORK_04	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_05	154	25	154 GB				85% Potential saving	0 GB
E880_RTC_SASWORK_06	154	25	154 GB				85% Potential saving	0 GB
E880_RTC_SASWORK_07	154	25	154 GB				86% Potential saving	0 GB
E880_RTC_SASWORK_08	154	25	154 GB				86% Potential saving	0 GB

Before RTC is enabled, comprestimator estimated ~ 80% - 90% potential saving in space for all three SAS workload file systems

# XIV RTC compression of SASDATA, WORK and UTIL volumes after RTC enabled

SASDATA original volumes Used Space = 193 GB








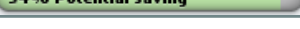
SASDATA compressed volumes used space = 18 GB

SASDATA volumes compression saving = 175 GB

**Compression Saving (%) for SASDATA volumes =  $(175/193)*100 = 91\%$**

	E880_RTC_SASDATA_01	309	18	309 GB		175 GB
	E880_RTC_SASDATA_01.original	309	193	309 GB		0 GB
	E880_RTC_SASDATA_02	309	18	309 GB		175 GB
	E880_RTC_SASDATA_02.original	309	193	309 GB		0 GB
	E880_RTC_SASDATA_03	309	18	309 GB		175 GB
	E880_RTC_SASDATA_03.original	309	193	309 GB		0 GB
	E880_RTC_SASDATA_04	309	18	309 GB		175 GB
	E880_RTC_SASDATA_04.original	309	193	309 GB		0 GB
	E880_RTC_SASDATA_05	309	18	309 GB		175 GB

**Compression Saving (%) for SASUTIL volumes = 89%**

	E880_RTC_SASUTIL_01	137	2	137 GB		17 GB
	E880_RTC_SASUTIL_01.original	137	18	137 GB		0 GB
	E880_RTC_SASUTIL_02	137	2	137 GB		17 GB
	E880_RTC_SASUTIL_02.original	137	18	137 GB		0 GB
	E880_RTC_SASUTIL_03	137	2	137 GB		17 GB
	E880_RTC_SASUTIL_03.original	137	18	137 GB		0 GB
	E880_RTC_SASUTIL_04	137	2	137 GB		17 GB
	E880_RTC_SASUTIL_04.original	137	18	137 GB		0 GB


**Compression Saving (%) for SASWORK volumes = 79%**

	E880_RTC_SASWORK_01	154	5	154 GB		20 GB
	E880_RTC_SASWORK_01.original	154	25	154 GB		0 GB
	E880_RTC_SASWORK_02	154	5	154 GB		20 GB
	E880_RTC_SASWORK_02.original	154	25	154 GB		0 GB
	E880_RTC_SASWORK_03	154	5	154 GB		20 GB
	E880_RTC_SASWORK_03.original	154	25	154 GB		0 GB
	E880_RTC_SASWORK_04	154	5	154 GB		20 GB
	E880_RTC_SASWORK_04.original	154	25	154 GB		0 GB


# CPU usage between baseline and XIV RTC



## System Summary brazos12 2/9/2016





## System Summary brazos12 2/9/2016





# Memory usage is same between baseline, and RTC



## What we learned:

- SAS Analytics workloads can perform well with IBM XIV Storage System Real-time Compression and IBM POWER8 and IBM Spectrum Scale (formerly IBM GPFS)
- Customers can realize positive performance results with their SAS Analytics workloads even if compression is enabled on the IBM XIV Storage system
- Real-time Compression is Native on the IBM XIV Storage System
- IBM Real-time Compression significantly decreases the storage capacity requirement of SAS workloads
- IBM Real-time Compression can drastically reduce the storage footprint in enterprise scale data centers that deploy SAS Analytics solutions
- Even better future test results anticipated with the latest version of IBM XIV Storage System
- See results whitepaper through SAS Global Forum app or at <http://ibm.co/1Nz3K22>
- Visit the IBM booths

Thanks for reading!