



Using IBM High Performance Services for HPC for architecture design with SAS analytics and SAS Grid Manager

Gregory Rohaly, IBM
Ken Gahagan, SAS Institute
Ben Smith, IBM

April 2016

 @IBMSystemsISVs



Table of contents

Abstract	1
Introduction	1
Prerequisites	2
IBM High Performance Services for HPC	2
SAS analytics.....	2
SAS Grid Manager.....	2
Spectrum Scale file management.....	3
Architecture for SAS testing	3
General description.....	3
Test bed description	4
Data and I/O throughput.....	4
Software and configuration settings.....	5
Software	5
SAS file systems with Spectrum Scale.....	5
SAS configuration settings	5
Hardware	6
Spectrum Scale client nodes.....	6
Network	7
Data architecture.....	8
Test results	9
20-stream test results	9
40-stream test results	10
60-stream test results	11
Test results summary.....	12
User processor time, system processor time, and I/O time comparisons.....	12
Observations on improving the system performance.....	13
I/O recommendations	14
Overall conclusion.....	14
Summary	15
Resources	16
About the authors	16
Trademarks and special notices	17



Abstract

This paper discusses how the deployment of an analytic solution for a company's benefit requires time, special skills, and data center resources that might be scarce or not found at all in the company. Using IBM High Performance Services for HPC can provide the skills and resources required to save the expenses of the company with a cloud deployment and can provide performance results comparable to using onsite resources. Using SAS Institute Inc. SAS analytics with SAS Grid Manager as the example, this paper shows how you can test an analytics suite in a cloud deployment and produce results that can be similar to using local resources for the same tasks. This helps in saving time and capital for the company with a cloud solution that allows users to build, manage, support, and protect applications in a private, public, or hybrid cloud environment.

Introduction

Businesses today are finding the potential of using data analytics across all dimensions of their activities. Defining corporate strategy, managing customer relationships, and streamlining operations are a few examples of how using analytics can help to win the competitive edge that is necessary in today's world. The use of analytics in a high-performance computing (HPC) environment is necessary to achieve the best results. But how is this accomplished without incurring huge costs of hiring skilled staff, provisioning hardware, and growing the data center?

One of the answers is to use a cloud service to host analytics when the demand is there. This defers costs, reduces the need for specialized skills, and helps to plan for the required time duration when there is a necessity and then turned off when the tasks are complete.

A method for companies to deploy analytics solutions rapidly in the cloud and not hire specialized skills is to use the IBM® High Performance Services for HPC. IBM High Performance Services for HPC provides consulting that has the specialty skills to advise clients and help to build, manage, support, and protect a cloud deployment.

This paper is written to address the following goals:

- Discuss the deployment of analytics in a cloud environment using the IBM High Performance Services for HPC service for rapid delivery,
- Use the SAS analytics suite with SAS Grid Manager as an example of a complex analytics deployment.
- Test the performance of SAS analytics in the cloud and show a comparison of the results by running in a cloud environment and onsite at a company's data center.

By accomplishing these goals, it should become clear that there are benefits in being able to use an on-demand cloud service with savings in time, economics, and in-house skills.

Prerequisites

Deploying an analytics solution in a cloud environment can be a complex endeavor. To understand the services and products used in this exercise, a brief explanation is provided. Each item discussed in this section has a referenced web link for further reading. These references are provided in the Resources section of this paper.

IBM High Performance Services for HPC

IBM High Performance Services for HPC is a service component that is used during the configuration phase to help create a high-performance configuration in a cloud environment.

IBM High Performance Services for HPC delivers a versatile, high-performing cloud-based environment for organizations that:

- Need to meet variable workload demands
- Need clustered resources, but do not have the budget or in-house skills to deploy and use a technical computing infrastructure
- Require non-shared physical systems, high-speed interconnects, the latest processor technology, and a choice of SoftLayer® data center locations for optimal application performance and security

IBM High Performance Services for HPC provides a single source for end-to-end cluster support, with access to technical computing experts, to eliminate the skills barrier for using clustered resources.

SAS analytics

The testing performed for this scenario was based on SAS9.4 (maintenance release 1).

SAS analytics are used by organizations to perform deep statistical analysis in numerous industries including financial services, healthcare, manufacturing, hospitality, and others that need to analyze large volumes of data.

The key benefits of using SAS analytics include improved decision making, increased analyst productivity, improved profitability and operational efficiency, and the ability to identify opportunities for growth.

SAS Grid Manager

SAS Grid Manager workload management software was used during the test to schedule and coordinate the use of resources for running the SAS analytic jobs dispatched to the system.

SAS Grid Manager is a powerful workload management platform for demanding, distributed SAS analytics environments. It provides a comprehensive set of intelligent, policy-driven scheduling features that enable the use of all the compute infrastructure resources and ensures optimal application performance.

Spectrum Scale file management

IBM Spectrum Scale™, formerly known as IBM General Parallel File System (IBM GPFS™), is a distributed and shared file system. In this test, it was an integral component for the effective deployment of the testbed configuration. It is used in conjunction with load-balanced servers on multiple client systems, and makes effective use of distributed SAS applications.

Organizations using this service can easily meet additional resource demands without the cost of purchasing or managing in-house infrastructure. It can minimize their administrative burden and quickly address the evolving business needs.

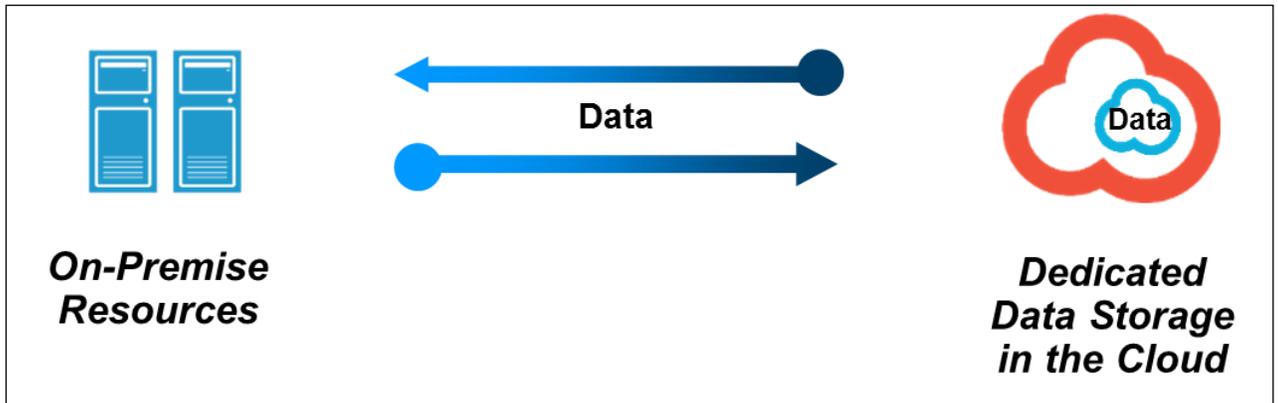


Figure 1: IBM Spectrum Scale on cloud manages efficient use of storage resources

Architecture for SAS testing

The development and design for the testing performed in the cloud environment was accomplished with the cooperation of a technical team from both SAS and IBM, along with the assistance of the IBM High Performance Services for HPC service.

The IBM High Performance Services for HPC service team worked with the SAS and IBM technical team to provide the following services:

- A secure testbed environment on physical systems in the cloud
- Local storage in the cloud
- External connectivity to the SoftLayer cloud service data center
- IBM Platform LSF® / IBM Platform Suite for SAS resources in the cloud (SAS Grid Manager)
- Support for the configuration and use of the SoftLayer cloud service

General description

Performance testing was conducted with servers and storage in a SoftLayer cloud environment. The purpose of the testing was to develop a relative measure of how well a system implemented in the cloud



and running an analytics workload (in this case, a SAS workload) might perform. In conducting such a test, it is important to keep in mind the premise that is *equal to or with better performance* should only be applied on a case-by-case basis. The test scenario described here illustrates how a particular test was developed and accomplished. It is meant to stimulate ideas of how a company interested in its own discovery might approach developing a test suitable for its situation.

The systems were configured and deployed at the SoftLayer data center in San Jose, California. The testing and reporting were conducted remotely by technical stakeholders from SAS and IBM personnel located North Carolina.

Test bed description

The test bed chosen for the cloud scenario was an analytics workload test suite used internally at SAS, which is described in this paper as the *SAS internal workload test*. Note that this should not be confused with the *SAS mixed analytics workload*, which is a test suite recognized by the SAS community and used to publish results by SAS to help compare performance metrics on a mixture of various architectures and operating systems.

The *SAS internal workload test* is a scaled workload of computation and I/O-oriented tests to measure concurrent and mixed job performance. The actual workload chosen consisted of 12 individual tests: four computation, four memory, and four I/O intensive tests. Each test consisted of multiple steps, some relying on existing data stores, with others (primarily computation tests) relying on generated data. The tests were chosen as a matrix of medium-running and shorter-running tests ranging in duration from approximately 20 seconds to 15 minutes. In some instances, the same test (running against replicated data streams) was run concurrently, and / or back-to-back in a serial fashion, to achieve an average of 20, 40, or 60 simultaneous streams of heavy I/O, computation (fed by significant I/O in many cases), and memory stress. In all, to achieve the 20-concurrent test matrix, 72 tests were run. For the 40-concurrent test matrix 144 tests were run and for the 60-concurrent test matrix 216 tests were run.

Data and I/O throughput

The I/O tests required an aggregate of approximately 300 GB of data, and the computation greater than 120 GB of data for a single instance of each test. Much more data is generated as a result of test-step activity, and threaded kernel procedures such as GSORT (for example, GSORT makes copies of the incoming file to be sorted).

As stated, some of the same tests run concurrently using different data, and some of the same tests are run back-to-back to garnish a total average of 20, 40, or 60 tests running concurrently. This increases the total I/O throughput of the workload significantly. Using the 20-stream example in its 27-minute span, the workload quickly jumps to 900 MBps and climbs steadily to 1.25 GBps, which is the bandwidth limit of the 10 Gb Ethernet connection for data before declining again. This creates a *SAS shop* throughput test in a grid environment in the cloud. This throughput is from all the three primary SAS file systems: SASDATA, SASWORK, and UTILLOC.

Software and configuration settings

The software used in the tests are products commercially available from SAS and IBM.

Note: In the following list, there is a reference to *IBM Platform Suite for SAS*. This is the fundamental core software for the SAS Grid Manager product but does not include the entire feature set. For all practical purposes, SAS Grid Manager was used in the testing scenarios.

Software

The products include:

- SAS analytics version 9.4 Maintenance 1
- IBM GPFS 3.5.0.17 (Spectrum Scale)
- IBM Platform Suite for SAS version 7.1 (SAS Grid Manager)

SAS file systems with Spectrum Scale

All file systems were created as Spectrum Scale distributed file systems.

For SAS applications there are three primary file systems:

- SASDATA - SAS permanent data file space
- SASWORK - SAS working data file space
- UTILLOC - SAS utility data file space

The following allocations were made for the code set, data, result space, working, and utility space:

- /gpfs/fs1/sasdata - 6.4 TB
- /gpfs/fs2/saswork - 3.2 TB
- /gpfs/fs2/util - 3.2 TB

This provides a general size of the application's on-storage footprint. It is important to note that throughput and not capacity is the key factor in configuring storage for SAS performance.

SAS configuration settings

For the testing, the following configuration changes were made from the default values:

- Changed SAS DATA from /tmp to /gpfs/fs1/sasdata
- Changed SAS WORK from /tmp to /gpfs/fs2/saswork
- Changed SAS UTILLOC from /tmp to /gpfs/fs2/util

The changes were reflected in the SAS Foundation base configuration file.

Hardware

There was a total of 13 servers allocated for the testing exercise. Of the 13 servers, five were designated as compute nodes or Spectrum Scale client nodes and eight were designated as storage servers or Spectrum Scale Network Shared Disk (NSD) servers using Spectrum Scale (GPFS).

Spectrum Scale client nodes

The five hardware servers have 10 cores and 64 GB of memory in each server. This provides 50 cores and 6.4 GB per core of available memory. Note that in the client servers, the two internal drives are for the operating system only. The cores might be described as clients or Spectrum Scale clients. The client roles are designated so that there is a management node and the remaining clients (cores) are available for job scheduling through the SAS Grid Manager application. The management node controls the workflow of the cluster and allows for job scheduling, queuing, and general workflow management.

The clients are described in the following table.

Physical servers	5 each / 10 cores each / 50 cores total
Processor	Intel® Xeon® processor E5-2690 2.9 GHz - two 20 MB cache, 1 socket, 10 core
OS	CentOS 6.5 kernel version 2.6.32-431.20.3.el6.x86_64
Memory	64 GB RAM, 16 GB DDR3 DIMMs, Registered, 1333 MHz
Internal storage	Two 500 GB SATA II disk drives, RAID 1 for OS
Network	10 Gbps private network uplink

Table 1: Spectrum Scale client nodes

The clients are designated as shown in the following table.

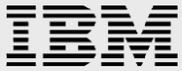
Node	Daemon node	Designation
1	compute050.sas....	Grid control
2	compute051.sas....	Client
3	compute052.sas....	Client
4	compute053.sas....	Client
5	compute054.sas....	Client

Table 2: Spectrum Scale client designations

Eight servers are configured and deployed as Spectrum Scale (GPFS) NSD servers. Each server is equipped with two 800 GB solid-state drives (SSDs) for data storage and the drives are referenced as SSD1 and SSD2. The SSDs in each server can be used to create two GPFS file systems. SSD1 is for /gpfs/fs1 and SSD2 is for /gpfs/fs2.

Each of the eight SSD1 drives is combined to create the /gpfs/fs1 file system and the other eight SSD2 drives are used to create the /gpfs/fs2 file system. The eight Spectrum Scale servers and the five Spectrum Scale clients share the 10 Gb network in order to access data.

The storage is configured using the Spectrum Scale (GPFS) NSD design. When a GPFS file system is created, a list of raw devices (that is, SSD1 are SSD2) is assigned to GPFS as NSDs. After an NSD is



defined, all of the nodes in the GPFS cluster can access the disk, using the GPFS NSD network protocol for shipping data over a TCP/IP connection.

Note that each server uses the 10 Gb network for client access. All 13 servers are joined on the 10 Gb network for storage access.

The following table provides details about the Spectrum Scale server nodes.

Physical servers	8 each / 2 SSDs each / 16 SSDs total
Spectrum Scale (GPFS) version	3.5.0.17
Number of SSDs per server	2
SSD size	800 GB each
Storage total (raw)	12.8 TB (8 servers, 2 SSDs, 800 GB)
File system 1	/gpfs/fs1/
File system 2	/gpfs/fs2/
File access method	NSD
Network	10 Gbps TCP/IP

Table 3: Spectrum Scale servers configuration

The Spectrum Scale (GPFS) cluster server configuration was allocated as shown in the following table:

Node	Daemon node	Designation	Function
1	gpfs-dmss1.sas	Quorum-manager	Primary manager
2	gpfs-dmss2.sas	Quorum-manager	Secondary manager
3	gpfs-dmss3.sas	Quorum node	NSD server
4	gpfs-dmss4.sas	Quorum node	NSD server
5	gpfs-dmss5.sas	Quorum node	NSD server
6	gpfs-dmss6.sas	Quorum node	NSD server
7	gpfs-dmss7.sas	Quorum node	NSD server
8	gpfs-dmss8.sas	Quorum node	NSD server

Table 4: Designation of Spectrum Scale servers

Network

When the testing is performed the jobs are submitted and the client nodes access data through the 10 Gb Internet Protocol network. The data will be provided from the Spectrum Scale servers using the GPFS NSD protocol.

The 10 Gb virtual local area network (VLAN) and the servers are zoned off from the rest of the data center traffic thus providing a quiet setting with no competing activity on the network, servers, or storage.

Data architecture

The flexibility of the configuration allows for different scenarios to be tested (which is discussed in later sections). A high-level illustration of the testbed configuration is shown in the following figure.

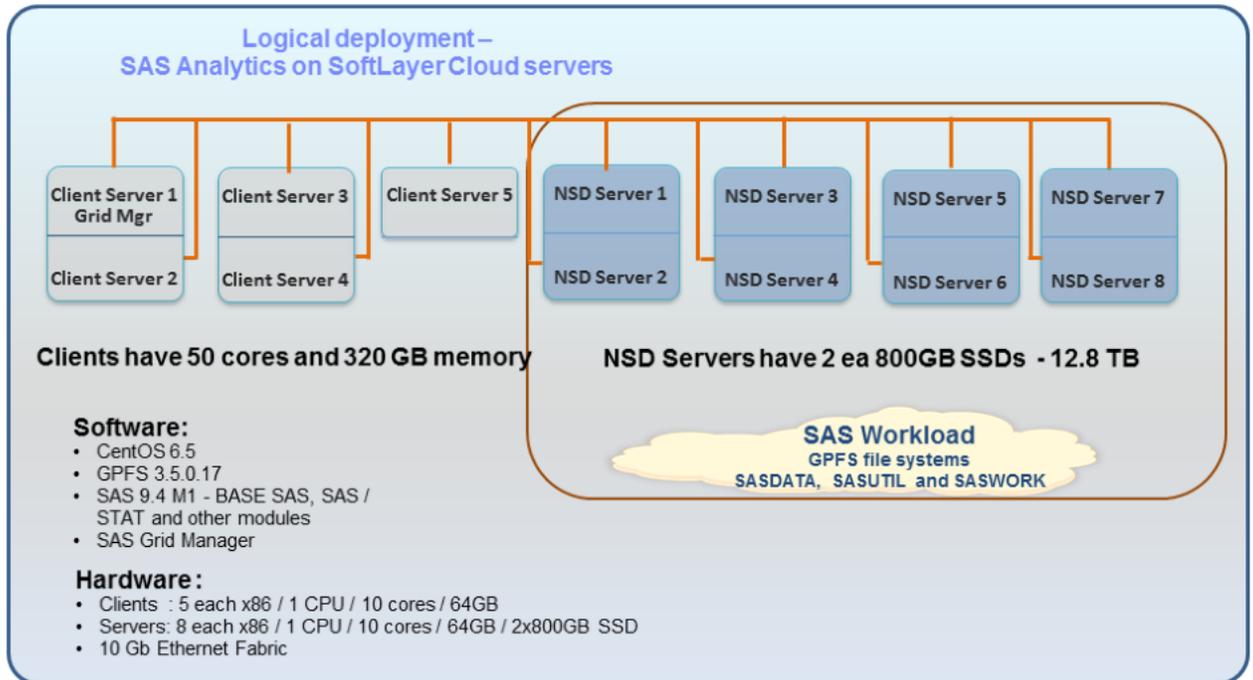


Figure 2: SoftLayer cloud test bed used for testing SAS 9.4 maintenance 1

Test results

This section provides the results of running the SAS internal workload test in the cloud-defined architecture. As the results are presented, keep in mind that the purpose of the tests is to show the feasibility of using cloud services. It is not intended to show the merits of a particular system architecture such as the servers, storage, or networks. Rather, the intention is to demonstrate that under a variety of workloads, it is possible to achieve reasonable performance using cloud resources.

The SAS internal workload test is a mix of processor and I/O-intensive SAS applications. It is intended to simulate SAS users in which various-sized mixed analytics workloads are generated and utilized processor, RAM and I/O resources. Of the three metrics just mentioned, I/O is usually a focus of attention and one of the critical aspects of the workload.

The test was set up as three different runs and executed using SAS 9.4 to simulate three situations of 20 streams, 40 streams, and 60 streams of simultaneous mixed analytics user sessions. Within each of the test blocks, some of the same tests run concurrently using different data, and some of the tests are run back-to-back, to produce a total average of the 20, 40, and 60 stream tests. Running the various tests in this fashion raises the total processor and I/O throughput of the workload significantly and provides the system under test with a heavy enough load to analyze where potential performance issues exist. The throughput is from all the three primary SAS file systems: SASDATA, SASWORK, and UTILLOC. The results of these tests are the performance measurements presented in the following tables for each test stream.

As stated earlier, the testing was performed in a quiet setup with no other activity being performed on the servers, storage, and network devices during the runs. Performance measurements for the tests begin at the start time of the execution and conclude at the end time of the job or jobs. The elapsed real time (wall clock) of all the job run times as well as the processor, system, and I/O times are recorded.

20-stream test results

The 20-stream test is presumably the easiest load that can test the system during these runs and gives an opportunity to form a baseline and might be used to compare and contrast the 40-stream and 60-stream runs performed later. The output in the following table gives the first look at the system performance for a variety of workloads.

Remember that during the stream test not all jobs are run for an equal number of times. Also depending on the job's purpose there are some jobs run in parallel and some run in sequence. So, it is with this in mind that the detail in the following table reveals the mix of loads and the response of the system to handle the work.

In the following table, the *elapsed real time* is the wall clock time consumed by the jobs that were run. As a result, adding the user + system + I/O times will be equal to the *elapsed real time*.

It was mentioned earlier that I/O activity is one of the most important aspects of performance. By using the formula: $(\text{user} + \text{system}) / \text{IO}$, there is a ratio calculated that indicates the efficiency of the I/O system to process read and write requests. When the efficiency is 100% (or better), it indicates the I/O system is able to meet the demands of the user and system processor requests. If the ratio is less than 100%, then it



means that the system is waiting for the I/O requests to be fulfilled. The I/O ratio can be less than 100% and still be considered usable for work considerations. However, when it falls to low levels, it is important to understand the cause for the issue and the solution for improvement. How low is the threshold? It is a judgment call by the performance team. Improving the I/O ratio might sometimes be simple and sometimes be very expensive. The following table shows the results of the 20-stream run. Note that the I/O efficiency ranges from 95% to 51%.

20-stream run Job type	Number of jobs run: 64	Elapsed real time in seconds	User processor time in seconds	System processor time in seconds	I/O time in seconds	20-stream processor and I/O efficiency
Codegen	8	1222	1163	2.24	56.76	95.36%
comp_test2	28	2628	2415	6.65	206.35	92.15%
comp_test5	2	35	27	4.88	3.12	91.09%
Customer	4	346	229	62.69	54.31	84.30%
dim_50mil	2	410	251	67.99	91.01	77.80%
io1_ca	2	291	210	4.30	76.70	73.64%
comp_test1	14	301	196	2.00	103.00	65.78%
hist_clm	2	1788	1027	0.00	761.00	57.44%
mixed4	2	43	22	0.16	20.84	51.53%
Total (sec)		7064	5540	151	1373	
Total (min)		118	92	3	23	

Table 5: 20-stream test run results

40-stream test results

The 40-stream test approximately doubles the work in the system. At times there might be twice the number of jobs that are run but it is not an exact doubling of the work of the 20-stream test. This is because of the way the scripts are written to keep the streams active. Because of how the job scheduling occurs it is not necessarily viable to do a direct comparison against the 20-stream results. The results are shown in the following table.

Some differences can be noted, especially in the processor and I/O efficiency. For example, **dim_50mil** went up from 77.8% in the 20-stream run to 51.99% in the 40-stream test indicating that the I/O activity in the system is not keeping up with the processor and system requests as well. You can find more details in a later topic.

40-stream run Job type	Number of jobs run: 128	Elapsed real time in seconds	User processor time in seconds	System processor time in seconds	I/O time in seconds	40-stream processor and I/O efficiency
Codegen	16	2405	2328	8.84	68.16	97.17%
comp_test2	56	5181	4827	16.2	337.8	93.48%
comp_test5	4	63	50	7.57	5.43	91.38%
Customer	8	940	458	133.89	348.11	62.97%
dim_50mil	4	1237	506	137.08	593.92	51.99%
io1_ca	4	478	419	8.55	50.45	89.45%
comp_test1	28	604	392	4.83	207.17	65.70%
hist_clm	4	4553	2061	636.6	1855.4	59.25%
mixed4	4	72	44	0.37	27.63	61.63%
Total (sec)		15533	11085	954	3494	74.78%
Total (min)		259	185	16	58	

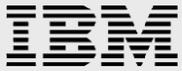
Table 6: 40-stream test run results

60-stream test results

The 60-stream test exercised the system capabilities in the most demanding manner of the three test scenarios. The 60-stream test had the longest run times. Again the processor and I/O efficiency gives an indication of how the I/O activity is being used in the system. Refer to the results in the following table.

60 stream run Job type	Number of jobs run: 192	Elapsed real time seconds	User processor time in seconds	System processor time in seconds	I/O time in seconds	60-stream processor and I/O efficiency
Codegen	24	3684	3496	9.93	178.07	95.17%
comp_test2	84	7982	7269	79.8	633.2	92.07%
comp_test5	6	92	81	11.88	-0.88	100.96%
Customer	12	1756	688	194.94	873.06	50.28%
dim_50mil	6	2389	755	206.28	1427.72	40.24%
io1_ca	6	703	628	12.34	62.66	91.09%
comp_test1	42	944	588	8.4	347.6	63.18%
hist_clm	6	8515	3110	975	4430	47.97%
mixed4	6	99	66	0.55	32.45	67.22%
Total (sec)		26164	16681	1499	7984	72.02%
Total (min)		436	278	25	133	

Table 7: 60-stream test run results



At the end of the 60-stream run, you can analyze the results of all the tests to determine if the cloud-based system is adequate and identify the potential bottlenecks.

Test results summary

For the three test runs on the cloud designed system, there can be some observations and analysis performed. For each aspect of the system performance, (user processor, system processor, and I/O) what were the areas that showed the most demands under increased load conditions? The following section shows some of the factors to consider when evaluating a system.

User processor time, system processor time, and I/O time comparisons

You can start with observing the overall scores for the tests. The following table provides a summary for the three stream tests. The elapsed time, user CPU, system CPU, and I/O times are shown in the following table. The percentages are calculated by dividing the respective time duration by the elapsed real time. This ratio gives a percentage of the time consumed for that activity.

20 / 40 / 60 Stream comparison	Elapsed real time in seconds	User processor time in seconds	% processor to total time	System processor time in seconds	% System to total time	I/O time in seconds	% I/O to total time
20-stream	7064	5540	78%	151	2%	1373	19%
40-stream	15533	11085	71%	954	6%	3494	22%
60-stream	26164	16681	64%	1499	6%	7984	31%

Table 8: Summary of the time consumed and percentages

In the 20-stream test, the processor efficiency was 78%, system time was 2%, and the I/O time was at 19%. This indicates that most of the time was spent in the processor running the code, which is the required outcome. However, note that as the workload increases through the other tests, the I/O efficiency decreases in an overall manner. This indicates that if the system were to be operating a workload in the range of the 60-stream category, then the I/O system would need to be studied and improved.

Diving further would allow to make some observations about the individual jobs themselves and their performance. If you extract the processor to I/O efficiency from each stream test, there could be more insight gained about the characteristics of these runs. This is presented in the following table.

Stream run Job type	20-stream processor / I/O efficiency	40-stream processor / I/O efficiency	60-stream processor / I/O efficiency
codegen	95.36%	97.17%	95.17%
comp_test2	92.15%	93.48%	92.07%
comp_test5	91.09%	91.38%	100.96%
customer	84.30%	62.97%	50.28%
dim_50mil	77.80%	51.99%	40.24%
io1_ca	73.64%	89.45%	91.09%
comp_test1	65.78%	65.70%	63.18%
hist_clm	57.44%	59.25%	47.97%
mixed4	51.53%	61.63%	67.22%

Table 9: Summary of job types with I/O efficiencies

Observe that the job types: codegen, comp_test2, and comp_test5 are very processor intensive. Their ratios are in the 90% to 100% range regardless of the 20-, 40-, or 60-stream runs. If the majority of the expected workload for this cluster is similar to these tests then the system would be ideal for running such tasks.

Now observe the job types: customer, dim_mil50, comp_test1, and hist_clm. Their scores indicate that the processor is waiting for an abnormal time due to I/O traffic. As the workload increases, the efficiency decreases and this results in further wait states for processing. If the majority of the expected workload were to be jobs like these, then the I/O system would need to be improved.

Remember, to determine if something is acceptable is a judgment call. There are no absolute guidelines. For the performance team to make those determinations, they would have to weigh the proposed changes against the cost benefits for improvements.

Observations on improving the system performance

Given the analysis of user, system, and I/O time durations the following observations are made. First, the user processor time durations are efficient for jobs that tend to be processor-intensive. Second, the I/O system appears to be a bottleneck for jobs that are I/O intensive and the condition increases as more work is applied to the system.

Because of these findings, there should be some time spent on exploring the I/O system and how to improve the performance. Given the inefficiencies, when under a heavy load this would be an area of interest and would need to have some recommendations for improvement. To do this, you need to refer back to the I/O system design.

I/O recommendations

For this discussion refer to Figure 2, which details the system architecture. The I/O system consists of eight servers and each server has two SSDs. All are connected on a 10 Gb network fabric. After studying the figure, it is noted that the three file systems were allocated in such a way that eight SSDs are /fs1/sasdata (SASDATA) and eight SSDs are /fs2/saswork (SASWORK) and /fs2/util (UTILLOC). While this seems to be an equal allocation of drives for the file systems, given their sizes of 6.4 TB, 3.2 TB, and 3.2 TB respectively, it might not be indicative of the amount of activity.

It was shown during the test runs that there was much more significant activity in /fs2/saswork than any other file system. It is therefore recommended that the file systems be reallocated in the 3:5 ratio and give more SSDs to /fs2/saswork to increase throughput. This would lead to a more balanced I/O system and should result in improvement for those I/O-intensive jobs. This would result in a zero increase in cost and should yield improvement in I/O performance.

Even after testing with the new allocations, if there is still an I/O performance problem, the next potential recommendation is to add an extra SSD to each server. This would increase the I/O throughput and decrease the amount of I/O time. There would also presumably be a small incremental cost to implement.

Only by trying the recommendations and repeating the tests it would be possible to confirm whether these solutions are viable. Because of time and system constraints, the team was not able to implement and re-test the system. But the confidence was high that these recommendations would bring improvements and possibly have the system operating at an acceptable performance level for I/O activity.

Overall conclusion

The concept of using a cloud-designed system for running analytics has been put to test. The results indicate that the cloud design performed well for processor-intensive jobs. It also brought to light that the configured I/O system would be suitable for processor-intensive workloads but would not be optimal for workloads that were I/O demanding. To that point some analysis and recommendations were made for resolution.

The results of the three tests were reviewed by the IBM and SAS teams for comment. Based on their observations the following statement is made:

After review by the IBM and SAS team, the results were deemed satisfactory for these tests.

Summary

The deployment of an analytic solution for a company's benefit requires time, special skills, and data center resources that might be scarce or not found at all in the company. The choices to consider are:

- Try to develop the IT resources organically and invest in data center equipment
- Use a consultant service and cloud vendor to provide these skills and services.

If a company occasionally makes use of analytics and does not need the services all the time, then using consultants and cloud resources is a good way to optimize costs and space.

By using the IBM High Performance Services for HPC as a consultant service along with a cloud provider, a company can make use of all the skills and resources needed to deploy a cloud environment and still achieve results comparable to using onsite resources.

By using SAS Institute Inc. SAS analytics with SAS Grid Manager as an example, it was shown how an analytics solution in a cloud deployment can be tested and produce results that are similar to using local resources for the same tasks. The three tests conducted and discussed in this paper not only demonstrated the viability of a cloud-based solution but also allowed the reader to gain appreciation of taking measurements, understanding the architecture, and knowing what changes in the configuration can provide the best value. These are the procedures and methods that might not always be practiced to their fullest in IT departments and hence the benefit of using an outside service with proven skills becomes an attractive alternative.

To learn more about this topic, refer to the websites in the following section.

Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
ibm.com/partnerworld/systems
- IBM High Performance Services For HPC
ibm.com/systems/services/platformcomputing
- SAS analytics
www.sas.com
- SAS Grid Manager
www.sas.com/en_us/software/foundation/grid-manager.html
- IBM Spectrum Scale
ibm.com/systems/storage/spectrum/scale
- IBM SoftLayer Cloud Services
www.softlayer.com
- IBM Publications Center
www.elink.ibm.link.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

About the authors

Gregg Rohaly is a consultant in IBM Systems Group. He has more than 10 years' experience working with IBM Systems. You can reach Gregg at grohaly@us.ibm.com.

Ken Gahagan is a director in the Research and Development center at SAS Institute in Cary, NC. You can reach Ken at ken.gahagan@sas.com

Ben Smith is a solutions architect in IBM Systems Group. He has more than 10 years of experience working with IBM Systems. You can reach Ben at smithbe1@us.ibm.com



Trademarks and special notices

© Copyright IBM Corporation 2016.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

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 www.ibm.com/legal/copytrade.shtml.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Automagic™, CDNLayer®, Challenging But Not Overwhelming™, CloudLayer®, Flex Images®, KnowledgeLayer®, RescueLayer®, SecurityLayer®, SoftLayer®, SoftLayer® device, StorageLayer®, and The Planet® are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the IO configuration, the



storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.