Redpaper

Barry Whyte
24 July 2013

IBM SAN Volume Controller and IBM FlashSystem 820: Best Practices and Performance Capabilities

Introduction

This IBM® Redpaper™ publication provides information about the best practices and performance capabilities when implementing a storage solution using IBM System Storage® SAN Volume Controller (SVC) with IBM FlashSystem™ 820. SVC is a member of the IBM Storwize® family, and FlashSystem 820 was formerly known as the Texas Memory Systems® RAMSAN Model 820.

This paper describes a typical system configuration, and a maximum system configuration that can return over 1,400,000 input/output operations per second (IOPS) at less than 1 millisecond average response time.

This paper also provides detailed performance information for a typical system configuration with the IBM Real-time Compression™ functionality available with SVC, including the recent hardware upgrade to enhance Real-time Compression throughput. This includes response curve information, and some examples of long-running, 100% write workloads, which are usually the most difficult workload for any “not and” (NAND) flash-based storage system.

This paper aims to provide guidance for those implementing a similar system, in addition to evidence that a similarly configured system should be capable of extreme, sustainable I/O throughput.

IBM SAN Volume Controller

IBM SVC is widely regarded as an industry-leading standard when it comes to storage virtualization. SVC has been shipping for almost 10 years (at the time of writing), and provides a single point of management and control for small and large heterogeneous storage environments.

For more details about SVC, see the IBM website:

http://www.ibm.com/storage/svc

IBM FlashSystem 820

IBM acquired Texas Memory Systems (TMS) in October 2012. TMS had been producing RAM-based disk systems since 1978, later moving to use NAND flash instead of RAM in their *RAMSAN* product line.

For more details about IBM FlashSystem 820, one of the first IBM FlashSystem branded products to be released, see the IBM website:

http://www.ibm.com/systems/storage/flash/720-820/

Configured system

This paper provides information about the best practices and expected performance when SVC and FlashSystem 820 are configured together. A total of eight SVC nodes and four FlashSystem 820 devices (each 20 TB) were available for testing. Some tests used all of the hardware, while others used a subset of the hardware. The test configurations are described in the following section, and are referenced throughout this paper.

Software levels

All of the tests in this paper used the following software versions:

- SVC code version 6.4.1.3
- FlashSystem 820 version 6.3[1]
- Host OS: Suse Enterprise Server version 11 SPC2 x64 (3.0.13-0.27)
- QLogic version 8.03.07.07-k

Configuration A

Table 1 shows Configuration A elements.

Table 1 Configuration A

	Type	Count	Total FC ports	Rack U	Comments
SVC	CG8	8	32x 8 gigabits (Gb)	16	(including UPS)
FlashSystem	820	4	16x 8 Gb	4	4x 19 TB usable (R5 configured)
Hosts	X3650M4	4	24x 8 Gb (QLE2462)	8	Suse ES 11SP2
Switch	6520	1	96x 8 Gb	2	72 ports used

[1] Due to the timing of the test runs, a pre-GA candidate level of firmware was used to benchmark the results. It is expected the GA level firmware holds the same performance capabilities.

Configuration B

Table 2 shows Configuration B elements.

Table 2 Configuration B

	Type	Count	Total FC ports	Rack U	Comments
SVC	CG8	4	16x 8 Gb	8	(including UPS)
FlashSystem	820	2	8x 8 Gb	2	2x 19 TB usable (R5 configured)
Hosts	X3650M4	2	12x 8 Gb (QLE2462)	4	Suse ES 11SP2
Switch	6520	1	96x 8 Gb	2	36 ports used

Configuration C

Table 3 shows Configuration C elements.

Table 3 Configuration B elements

	Type	Count	Total FC Ports	Rack U	Comments
SVC	CG8	2	8x 8 Gb	4	(including UPS)
FlashSystem	820	1	4x 8 Gb	1	1x 19TB usable (R5 configured)
Hosts	X3650M4	1	6x 8 Gb (QLE2462)	2	Suse ES 11SP2
Switch	6520	1	96x 8 Gb	2	36 ports used

Current configuration best practice

The internal architecture of both SVC and FlashSystem 820 dictate certain optimal or minimum numbers of logical objects, as well as zoning best practices. It is assumed that you have read and understand the general SVC best practice concerning switch zoning.

In summary, these are the rules for configuration:

- Zone one port from each *pair of ports* on FlashSystem 820 to a pair of SVC ports on each node.
- Repeat for the other FlashSystem 820 and SVC node ports.

Figure 1 shows two zones, red and blue.

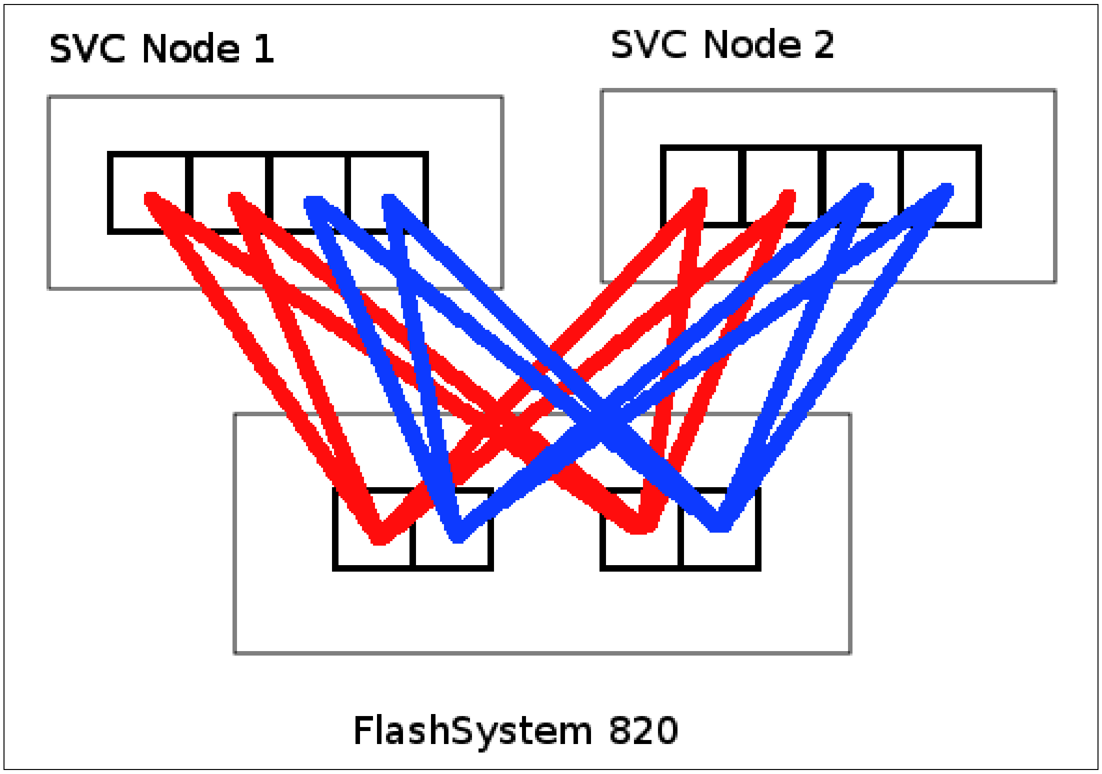

Figure 1 Zoning best practice

When configuring multiple IO group SVC clusters, you need to zone every node so that it can access FlashSystem 820. In this case, the zoning shown in Figure 1 should be repeated with each IO group (node pair). If you configure multiple FlashSystem 820 ports, repeat the zoning so that every SVC node has access to every FlashSystem 820 port.

Logical configuration on FlashSystem 820

To provide usable storage (managed disks) on SVC, you need to define some logical units (LUs) on FlashSystem 820, and map these LUs to the FlashSystem *host ports*.

Create at least four LUs on FlashSystem 820 storage, and use default masking to map all of the LUs to all of the FlashSystem 820 host ports.

When you create zoning as shown in Figure 1, each managed disk (mdisk) discovered by SVC will have eight available paths between each SVC node and the FlashSystem port.

The tested configurations in this paper used either 4 LUs or 16 LUs using varying FlashSystem 820 capacity to test 25%, 50%, and 100% allocation. That is, with 4 LUs and 50% allocation, each LU was approximately 2.5 TB in size. Similarly, with 16 LUs and 50% allocation, each LU was approximately 625 GB in size.

Note: SVC code levels before 7.1.0.2 do not support single LUs presented from FlashSystem 820 that are greater than 2TB in size. If you are attaching the FlashSystem to an older level of code you should ensure you create LUs that are smaller than 2TB. The same guidelines should be followed, and therefore the best practice would be to create 16 LUs in this case.

Logical configuration on SVC

This section describes storage pool, volume, and host configuration.

Storage pool configuration

After the four or more LUs have been discovered by SVC, the **`lscontroller`** and **`lsmdisk`** commands should show FlashSystem 820 as a controller, and the four or more LUs as unmanaged mdisks.

If you use FlashSystem 820 as the primary data storage, as with the test results here, add all of the mdisks from the controller to a single *managed disk group* (also know as a *storage pool* in the SVC GUI).

If more than one FlashSystem 820 is being presented to an SVC cluster, a best practice would be to create a single storage pool per controller.

If you use FlashSystem 820 with the SVC EasyTier function, you will likely want to create multiple LUs for each *hybrid storage pool*. Create four or more LUs per hybrid pool, with the combined capacity of these LUs matching the capacity that you want for the *ssd tier* in that pool.

Volume (vdisk) configuration

With a single volume, SVC can generally return 100% of the performance capabilities of the storage pool from which it is created. However, to make full use of all of the SVC ports, internal cores, and of course both nodes, you should create a minimum of eight volumes per FlashSystem 820. This aligns volumes with cores, and ensures that all cores on both nodes are fully used, assuming that the host multipathing software accepts the asymmetrical logical unit access (ALUA) reported by SVC volumes.

SVC is a true active-active system. That is, I/O can be sent to either node within an I/O group that presents a volume. To allow load-balancing, SVC does report an ALUA preference. So each volume is assigned preferred ownership to a given node. This allows for better cache hit and destage performance.

Some operating systems do not accept the preference request, and in general this is not an issue, but with some skewed workloads, one node can end up busier on destage operations than the other. In a typically balanced system, you will not see this.

In the test configurations reported in this paper, 64 volumes per FlashSystem 820 were created. Therefore, in the cases where an eight-node SVC cluster is supporting four FlashSystem 820 systems, a total of 256 volumes were mapped from SVC to the host systems.

Host configuration

The host systems used in this test are standard IBM System x® 3650M4. Each system had three QLogic QLE8462 dual-port, 8-Gb Fibre Channel cards.

The system used a standard installation of Suse Enterprise Server version 11, service pack 2 with the standard *in distribution device driver* and *multipath driver*. The multipath configuration was modified to allow the preferred node as defined by SVC SCSI inquiry data to be honored. This was done per the guidelines in the SVC information center:

https://ibm.biz/BdxAPK

The system was further modified to enable the *noop* kernel scheduling fairness option, again as described in the normal SVC host attachment guidelines.

A single host server was zoned per SVC IO group, so for the purposes of these tests a host was generally accessing just one FlashSystem 820. As the number of SVC nodes and FlashSystem 820 storage devices was increased, so were the number of host systems, up to four.

FlashSystem preconditioning

Although this paper covers the available performance when using SVC and FlashSystem 820, it is not the purpose to provide information about general NAND flash performance itself. It is assumed that you have a basic understanding of the performance characteristics of NAND flash.

In particular, NAND flash suffers from a reduction in performance when a system becomes bottlenecked on the amount of concurrent new write/flash erase cycles that it can run. This is generally seen as an initial peak write workload followed by a sudden and dramatic falloff in sustainable write throughput.

To account for this, and to ensure actual (worst case) representative performance benchmarks in this paper, each test was performed after a sustained 8 KB random write workload had been run to 2x the total system capacity. This took approximately 2 hours to run before the measurement test workloads.

It is interesting to note the actual performance (as measured during these preconditioning cycles), and to compare them against a “typical” enterprise-class, solid-state drive (SDD) device.

Figure 2 on page 7 is a plot over time running a sustained, deeply queued (hence quite high response time) workload. This used Configuration C (described in “Configuration C” on page 3) with 100% FlashSystem 820 capacity allocation. The sample period for each point in the plot is 20 seconds. You can see that after sample 104 (therefore around 35 minutes)[2], the write loading causes the system to reduce the throughput by approximately 30% (from 250 K IOPS to 175 K IOPS)

In contrast, Figure 3 on page 7 shows a similar preconditioning workload running against a single enterprise-class SSD (drive form factor SSD). This shows that after sample 18 (therefore around 6 minutes) the write loading causes the SSD to reduce the throughput by approximately 65% (from 35K IOPS to 12K IOPS).

[2] During testing, it was noticed that a system with 50% capacity allocation took over an hour before reaching the same point. In addition, one using 25% allocation took over two hours.

The conclusion is that although all NAND flash devices have a similar workload characteristic, when compared with an industry standard enterprise-class SSD device, FlashSystem 820 suffers a much smaller reduction in performance after a much greater period of time.

In actual workloads, it is unlikely to have such a sustained workload for long periods of time. However, it is common to have a short-term (5 - 10 minute) burst of such a workload. FlashSystem 820 can easily sustain peak workloads, and prevent users from seeing a performance degradation.

Testing also showed that FlashSystem 820 recovered to maximum write performance in as little as 5 minutes after reducing the write workload below maximum, and then continued returning peak write performance for another 35 minutes before again seeing a drop in performance. Other testing has shown that some NAND flash devices take several hours to recover.

The graph in Figure 2 is for a single SVC IO group and a single FlashSystem 820, with sustained random write. Each sample as shown is 20 seconds. The write falloff occurs at sample 104.

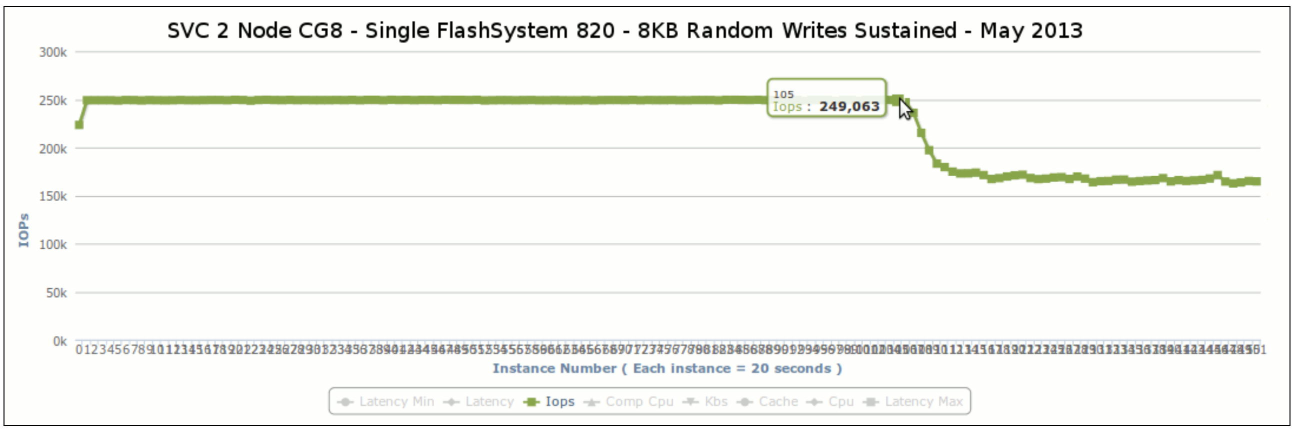

Figure 2 Single SVC IO group

The graph in Figure 3 is for a single enterprise-class SSD device with sustained random write. Each sample as shown is 20 seconds. The write falloff occurs at sample 18.

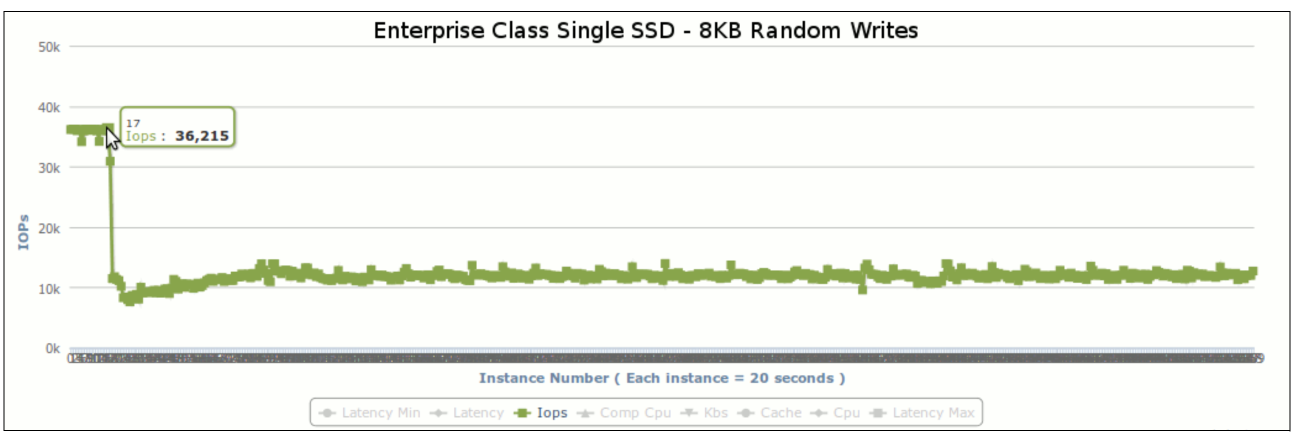

Figure 3 Single enterprise-class SSD device

Summary performance

The configurations outlined in “Configured system” on page 2 were used to benchmark the maximum throughput levels that both systems could achieve when configured as outlined in “Current configuration best practice” on page 3.

In all cases, two specific sets of tests were executed:

- SVC cache enabled
- SVC cache disabled

SVC cache considerations

Each SVC node contains a read/write cache, 24 GB per node, that can be enabled or disabled on a per-volume basis. Consider a *storage controller cache*, and the reasons why it exists.

Generally, you use a storage controller cache to help hide the latency of read and write operations from a traditional spinning magnetic medium.

It is not the intention of this paper to provide information about the performance characteristics of traditional hard disk drive (HDD) devices, but typically a storage controller cache provides user applications with the advantage of fast write performance. That is, it can write into cache (electronic speed) and complete I/O almost instantly, and handle the destage operations to the storage medium later. In addition, *prefetch (or read-ahead) algorithms* can pre-buffer potential read data from the storage medium.

With the introduction of NAND flash devices, I/O operations complete in sub-millisecond (generally 100 - 200 microsecond) latencies, rather than 10 or more milliseconds to/from magnetic media. This does suggest that in some specific workloads, the controller cache might no longer provide such a substantial benefit.

For example, you always want to create at least two copies of any cached write data, to ensure protection against the loss of one caching node. This means you need to use half of the available system resources simply to mirror the data to another cache. Generally, this accounts for write performance that is roughly half the read performance, and this is when systems become saturated.

It is conceivable, therefore, that with NAND flash devices, you might want to disable controller read/write caching. This is the case if you care most about achieving the absolute maximum number of IOPS, while still maintaining sub-millisecond response times.

Note: In general we would recommend that you run with SVC cache enabled, especially if you are deploying any advanced functions, such as FlashCopy, Global Mirror, and so on.

It is only necessary to consider disabling the SVC cache if you require IOPS that exceed those that can be achieved with cache enabled *and* you are not using SVC advanced functions.

Peak performance measurements

As you can see in Table 4 on page 9, adding more SVC nodes and FlashSystem 820 storage causes performance, both in terms of IOPS and MBps, to scale linearly.

Table 4 Peak performance measurements

Configuration	**A**		**B**		**C**	
SVC	**8 Node**		**4 Node**		**2 Node**	
SVC Cache	**On**	**Off**	**On**	**Off**	**On**	**Off**
FlashSystem 820	**4**		**2**		**1**	
Read Miss IOPS	1,348,000	1,532,000	634,000	687,000	360,000	392,000
Write Miss IOPS	448,000	1,081,000	227,000	499,000	111,000	274,000
70/30 Miss IOPS	847,000	1,376,000	425,000	625,000	206,000	366,000
Read Miss MBps	12,660	13,000	6,350	6,590	3,050	2,930
Write Miss MBps	8,600	9,750	4,300	4,800	2,200	2,400

The results also show that disabling the SVC cache can result in more than twice the write IOPS. As covered in “SVC cache considerations” on page 8, this is due to the SVC nodes themselves having to issue fewer I/Os over the Fibre Channel interfaces (in particular, the extra I/O and control messages needed to keep the write cache mirrored between nodes).

However, as also covered in “SVC cache considerations” on page 8, this should only be considered if you require IOPS rates greater than those that can be achieved with the cache enabled. An alternative configuration is to use more SVC nodes, such as a 4-node SVC cluster with cache enabled, fronting a FlashSystem 820 system.

Each test was repeated with the 25%, 50%, and 100% FlashSystem 820 capacity utilization. In all cases, the peak performance was the same. This means users can feel confident that performance remains consistent as they fill the capacity of their FlashSystem 820.

In March 2013, IBM introduced an optional upgrade for the SVC Model CG8 node hardware. This upgrade adds an additional eight 8-Gb Fibre Channel ports per SVC node pair. When installed, and zoned correctly, this can double the available write bandwidth per node pair.

Thus, the write MBps shown in Table 4 for a cache-disabled configuration can be achieved with cache-enabled workloads. The upgrade can be ordered through the Storage Customer Opportunity REquest (SCORE)/RPQ process for nodes running 6.4.1 or later software:

`http://iprod.tucson.ibm.com/systems/support/storage/ssic/interoperability.wss`

Peak Real-time Compression performance

One of the key features enabled by SVC is Real-time Compression. You can enable and disable this function on a per-SVC volume basis. Although the performance of a compressed volume can meet the same requirements as an uncompressed volume, overall system performance becomes a limiting factor much sooner than when running non-compressed volumes.

In March 2013, IBM introduced a new upgrade feature for the SVC Model CG8 node hardware. This upgrade doubles the number of system cores, from 6 to 12, and adds an additional 24 GB of cache capacity. These extra cores and additional cache are dedicated for use by compressed volumes. This upgrade is ordered by RPQ and is only available for clusters running version 6.4.1 or later software.

Table 5 demonstrates the performance benefits of this upgrade compared with a standard CG8 node pair. In these tests, Configuration C was used, with and without the compression hardware upgrade. All of these tests were run with SVC cache enabled, because it is better not to run compressed volumes with the SVC cache disabled.

For a detailed overview of IBM Real-time Compression and best practices, including more details about best and worst case workloads, review *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859:

http://www.redbooks.ibm.com/redpapers/pdfs/redp4859.pdf

Best case denotes a workload that has good "temporal locality" and suits the compression block allocation and retrieval algorithms.

Worst case denotes a workload that has zero "temporal locality" and defeats the compression block allocation and retrieval algorithms. For this reason, IBM best practices for Real-time Compression include only selecting highly compressible volumes. IBM provides tools to determine which volumes are compressible and which volumes are not.

Table 5 Real-time Compression (RtC) CG8 compared to CG8 with compression upgrade

SVC	**2 Node**		**2 Node**	
Compression upgrade	**NO**		**YES**	
FlashSystem 820	**1**		**1**	
Read Miss IOPS	2,600	50,000	40,100	143,000
Write Miss IOPS	1,200	17,000	28,300	93,200
70/30 Miss IOPS	2,200	40,700	41,900	116,400

With the compression hardware upgrade, the best-case performance is improved by up to 5x and the worst case by up to 20x. The performance is close to the performance capability of a two-node SVC system running non-compressed volumes.

To achieve the maximum IOPS capability of a single FlashSystem 820 running 100% compressed volumes, the best practice would be to configure four SVC nodes with the compression hardware upgrade.

Detailed response curve comparisons

Although achieving maximum throughput and IOPS rates is interesting, the main value of IBM FlashSystem 820 to application users and developers is the greatly improved response time that the system provides.

The following pages show standard performance response curves, where the workload is increased and samples are taken to generate a "hockey stick" curve. These all demonstrate that sub-millisecond response times can generally be achieved up to 90% or more of the total IOPS throughput, and that SVC itself adds little latency.

All of the test results in this section used Configuration A (an eight-node SVC cluster and four FlashSystem 820 systems with SVC cache disabled).

Three workloads are demonstrated:

- 100% read random 4 KB
- 100% write random 4 KB
- 70% read 30% write (70/30) random 4 KB

Table 6 shows the following peak IOPS rates, while maintaining a response time of under 1 millisecond.

Table 6 Peak IOPS while maintaining under 1 ms response time

	Peak IOPS under 1 ms	**Per SVC IO group/single 820 peak IOPS under 1 ms**
Read Miss IOPS	1,400,000	350,000
Write Miss IOPS	1,000,000	250,000
70/30 Miss IOPS	1,350,000	337,500

All response time measurements are recorded at the host application, so they include the fabric round-trip latency and the SVC and FlashSystem 820 latency. The results shown in the following graphs are all with SVC cache disabled.

A 4 KB random read performance

Figure 4 shows a 4 KB random read performance.

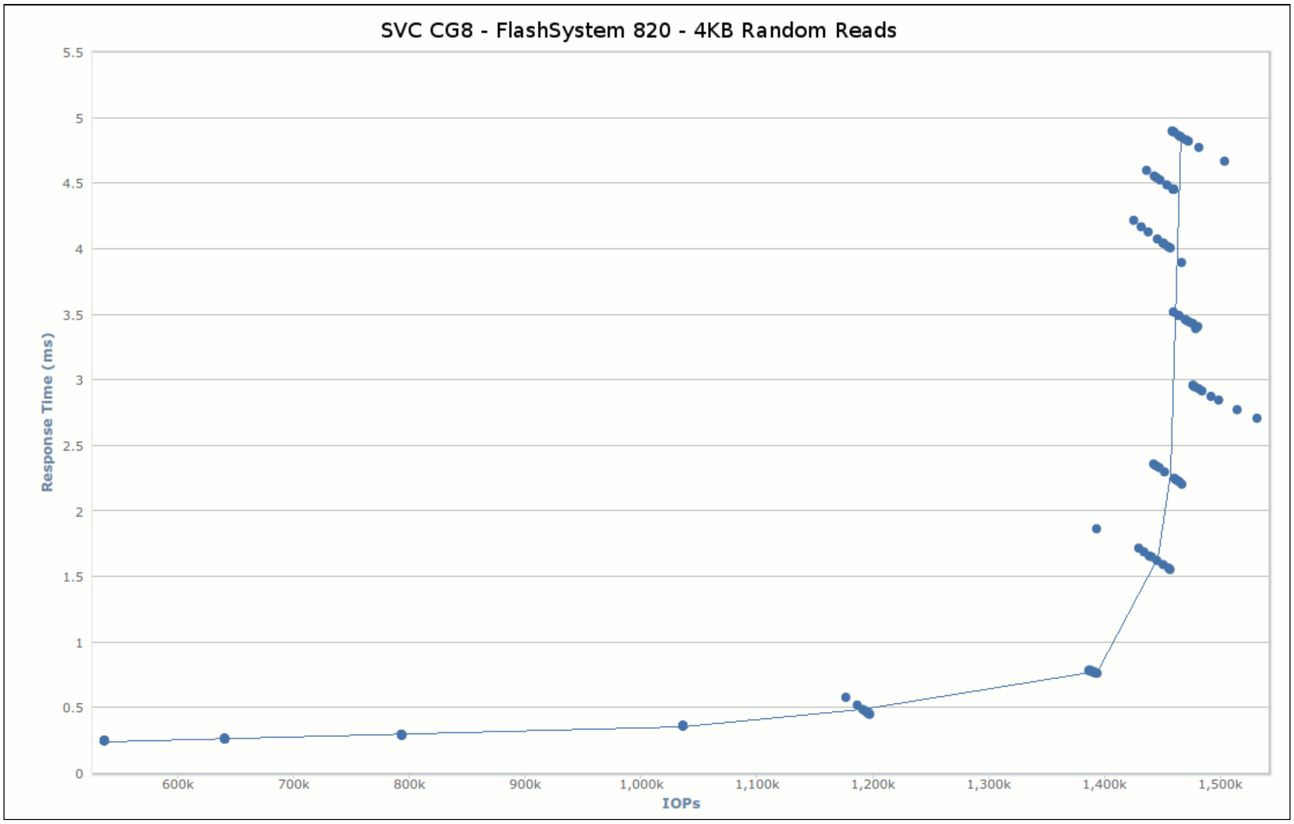

Figure 4 4 KB random read performance

A 4 KB random write performance

Figure 5 shows a 4 KB random write performance.

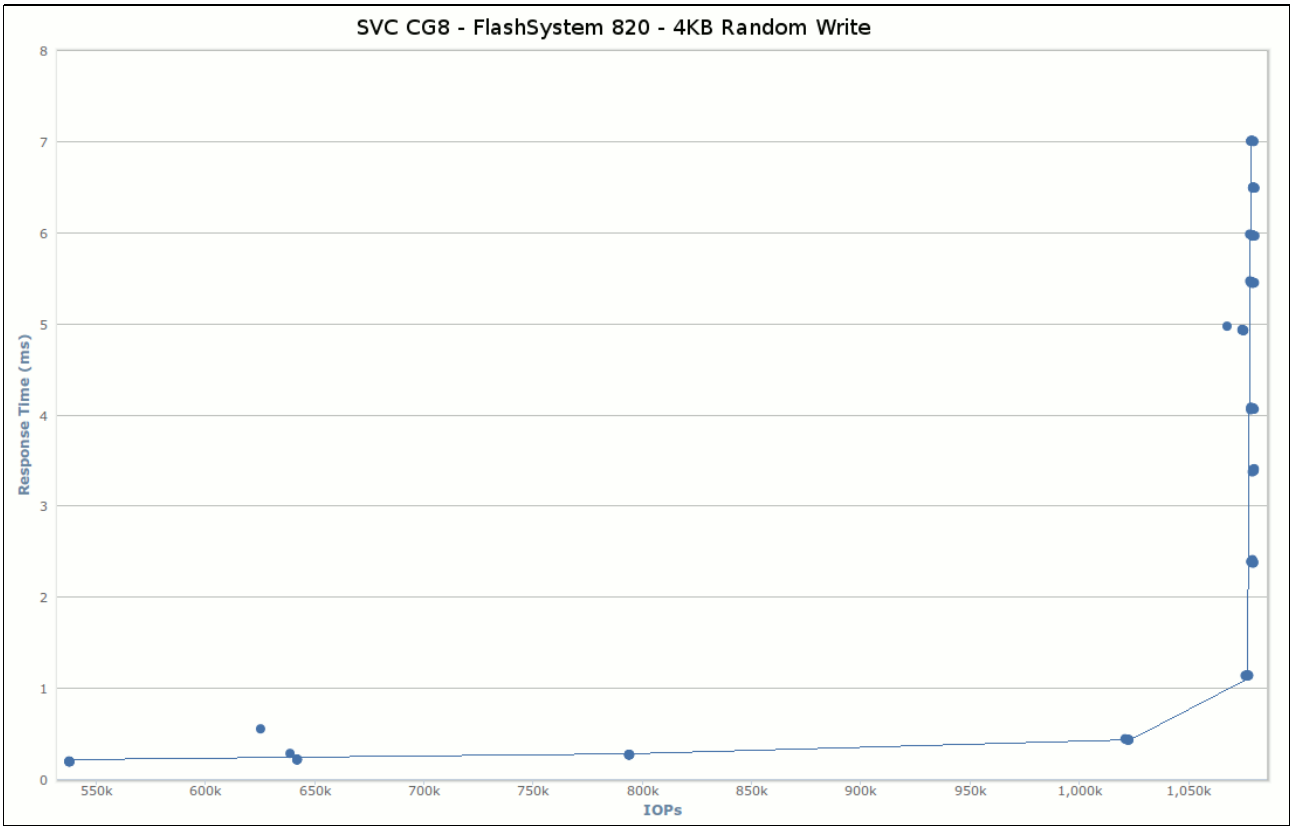

Figure 5 4 KB random write performance

A 4 KB mixed 70/30 performance

Figure 6 shows a mixed 70/30 performance.

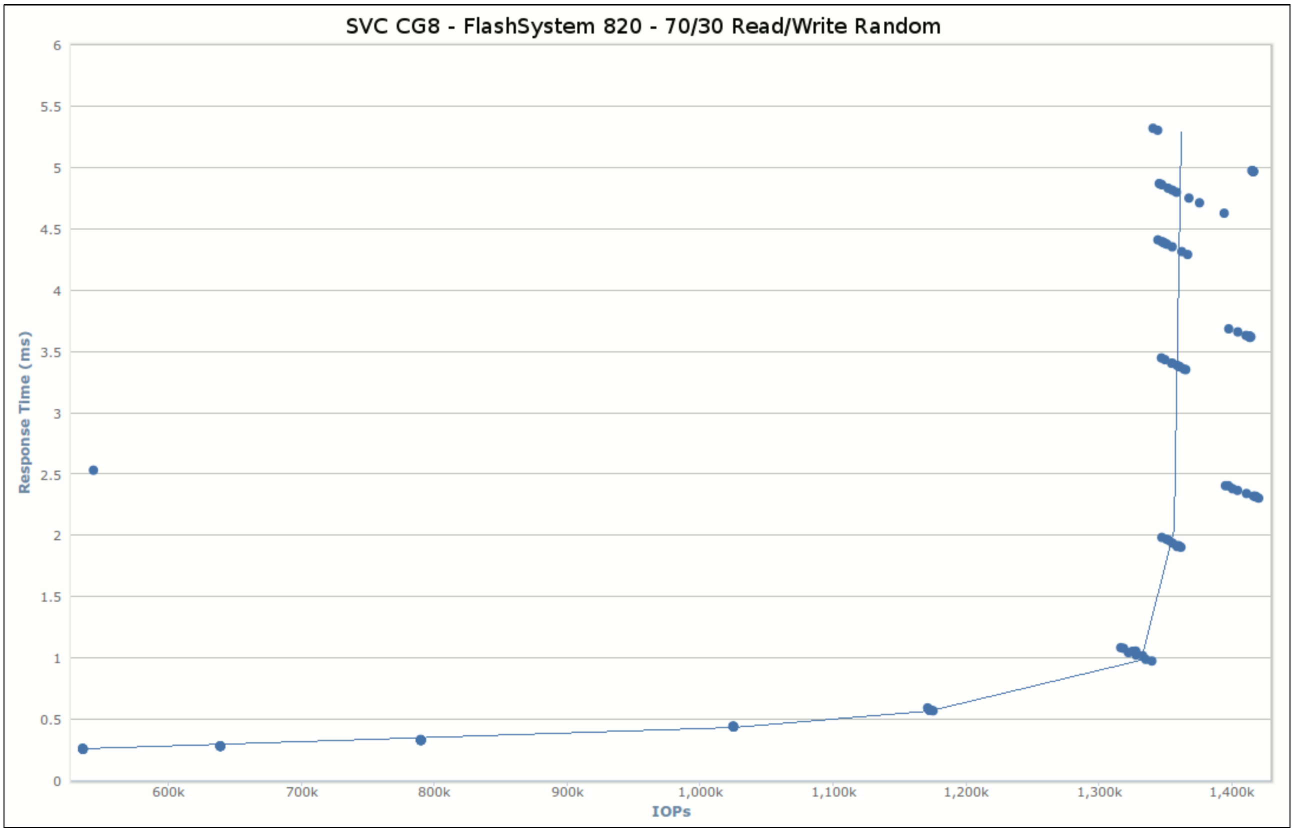

Figure 6 Mixed 70/30 performance

Conclusions

This paper documents best practice configurations for systems configured to use IBM SAN Volume Controller and IBM FlashSystem.

Although the systems under test use the SVC CG8 nodes and FlashSystem 820 storage, the same guidelines should be followed if using other SVC or FlashSystem models.

This document also details the potential performance that such systems can provide, over 1 million real-life disk IOPS on a typical online transaction processing (OLTP)-type workload. The paper also contains information about the options for SVC cache usage and Real-time Compression.

Finally, detailed response curve graphs show that extreme low latency can be achieved by such systems up to close to the maximum performance limits.

The added value of the features and functions provided by SVC, combined with the low latency and high IOPS potential of IBM FlashSystem, make configurations deploying both technologies compelling for enterprise-class NAND flash requirements.

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Hursley, UK Center.

Barry Whyte is a Master Inventor working in the Systems and Technology Group based in IBM Hursley, UK. Barry primarily works on the IBM SAN Volume Controller and Storwize V7000 virtual disk systems. Barry graduated from the University of Glasgow in 1996 with a B.Sc (Hons) in Computing Science. In his 16 years at IBM, he has worked on the successful Serial Storage Architecture (SSA) and the IBM DS8000® range. Barry joined the SVC development team soon after its inception, and has held many positions before taking on his current role as a Performance Architect. Read his blog at: https://ibm.biz/BdxAym

The project manager:

Jon Tate is a Project Manager for IBM System Storage SAN Solutions at the International Technical Support Organization, San Jose Center. Before joining the ITSO in 1999, he worked in the IBM Technical Support Center, providing Level 2/3 support for IBM storage products. Jon has over 27 years of experience in storage software and management, services, and support. He is both an IBM Certified Consulting IT Specialist and an IBM SAN Certified Specialist. He is also the UK Chairman of the Storage Networking Industry Association.

Thanks to the following people for their contributions to this project:

Carlos Fuente
Bill Scales
IBM Hursley

Chris Saul
IBM San Jose

Olga Yiparaki
IBM Tucson

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

http://**ibm.com**/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

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.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-5027-00 was created or updated on July 24, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  http://**ibm.com**/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD, Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

Redpaper™

Trademarks

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. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

DS8000®
FlashSystem™
IBM®
IBM FlashSystem™
Real-time Compression™
Redbooks®
Redpaper™
Redbooks (logo) ®
Storwize®
System Storage®
System x®

The following terms are trademarks of other companies:

Texas Memory Systems, and the Texas Memory Systems logo are trademarks or registered trademarks of Texas Memory Systems, an IBM Company.

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