



SPC BENCHMARK 1™
FULL DISCLOSURE REPORT

IBM
IBM® STORWIZE® V5030

SPC-1 V3.2

SUBMISSION IDENTIFIER: A31001

SUBMITTED FOR REVIEW: JULY 14, 2016

First Edition – July 2016

THE INFORMATION CONTAINED IN THIS DOCUMENT IS DISTRIBUTED ON AN AS IS BASIS WITHOUT ANY WARRANTY EITHER EXPRESS OR IMPLIED. The use of this information or the implementation of any of these techniques is the customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item has been reviewed by IBM, for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environment do so at their own risk.

This publication was produced in the United States. IBM may not offer the products, services, or features discussed in this document in other countries, and the information is subject to change with notice. Consult your local IBM representative for information on products and services available in your area.

© Copyright IBM 2016. All rights reserved.

Permission is hereby granted to publicly disclose and reproduce this document, in whole or in part, provided the copyright notice as printed above is set forth in full text on the title page of each item reproduced.

Trademarks

SPC Benchmark 1, SPC-1, SPC-1 IOPS, SPC-1 LRT and SPC-1 Price-Performance are trademarks of the Storage Performance Council.

IBM, the IBM logo and Storwize are trademarks or registered trademarks of IBM in the United States and other countries. All other brands, trademarks, and product names are the property of their respective owners.

Benchmark Specification and Glossary

The official SPC Benchmark 1™ (SPC-1™) specification is available on the website of the Storage Performance Council (SPC) at www.storageperformance.org.

The SPC-1™ specification contains a glossary of the SPC-1™ terms used in this publication.

Table of Contents

Audit Certification.....	iv
Letter Of Good Faith	vi
Executive Summary	7
Configuration Information	12
Benchmark Configuration and Tested Storage Configuration.....	12
Benchmark Configuration Creation Process	14
Benchmark Execution Results.....	16
Benchmark Execution Overview	16
SUSTAIN Test Phase.....	18
RAMPD_100 Test Phase.....	21
Response Time Ramp Test.....	24
Repeatability Test	26
Data Persistence Test	29
Appendix A: Supporting Files	30
Appendix B: Third Party Quotation	31
Appendix C: Tuning Parameters and Options	32
Appendix D: Storage Configuration Creation	33
Appendix E: Configuration Inventory.....	36
Appendix F: Workload Generator	37

AUDIT CERTIFICATION



Bruce McNutt
IBM ARC
IBM Corporation
650 Harry Road
San Jose, CA 95120

July 14, 2016

I verified the SPC Benchmark 1™ (SPC-1™ Revision3.2) test execution and performance results of the following Tested Storage Product:

IBM® STORWIZE® V5030

The results were:

SPC-1 IOPS™	85,020
SPC-1 Price-Performance™	\$0.47/SPC-1 IOPS™
SPC-1 IOPS™ Response Time	13.358 ms
SPC-1 Overall Response Time	5.193 ms
SPC-1 ASU Capacity	34,360GB
SPC-1 ASU Price	\$12.22/GB
SPC-1 Total System Price	\$419,607.89

In my opinion, these performance results were produced in compliance with the SPC requirements for the benchmark.

The testing was executed using the SPC-1 Toolkit Version 3.0 Build 2565. The audit process was conducted in accordance with the SPC Policies and met the requirements for the benchmark.

A Letter of Good Faith was issued by the Test Sponsor, stating the accuracy and completeness of the documentation and testing data provided in support of the audit of this result.

A Full Disclosure Report for this result was prepared by InfoSizing, reviewed and approved by the Test Sponsor, and can be found at www.storageperformance.org under the Submission Identifier **A31001**.

20 KREG LANE • MANITOU SPRINGS, CO 80829 • 719-473-7555 • WWW.SIZING.COM

The independent audit process conducted by InfoSizing included the verifications of the following items:

- The physical capacity of the data repository;
- The total capacity of the Application Storage Unit (ASU);
- The accuracy of the Benchmark Configuration diagram;
- The tuning parameters used to configure the Benchmark Configuration;
- The Workload Generator commands used to execute the testing;
- The validity and integrity of the test result files;
- The compliance of the results from each performance test;
- The compliance of the results from the persistence test;
- The compliance of the submitted pricing model; and
- The differences between the tested and the priced configuration, if any.

The Full Disclosure Report for this result was prepared in accordance with the disclosure requirements set forth in the specification for the benchmark.

The following benchmark requirements, if any, were waived according to the SPC Policies:

- The execution of the VERIFY_2 step was omitted. After examination of this deficiency, it was determined that this exception met the necessary conditions to be granted a waiver. Additional documentation can be found in the Full Disclosure Report.

Respectfully Yours,



François Raab, Certified SPC Auditor

LETTER OF GOOD FAITH



Vice President, XIV and Storwize Development

IBM Technology & Systems Group
9000 South Rita Road
Tucson, AZ 85744-0002

Phone 1-520-799-4212

June 24, 2016

Mr. Francois Raab, Certified SPC Auditor
InfoSizing, Inc.
20 Kreg Lane
Manitou Springs, CO 80829

Subject: SPC-1 Letter of Good Faith for the IBM Storwize V5030.

IBM Corporation is the SPC-1 Test Sponsor for the above listed product. To the best of our knowledge and belief, the required SPC-1 benchmark results and materials we have submitted for that product are complete, accurate, and in full compliance with Version 3.2 of the SPC-1 benchmark specification.

Our disclosure of the Benchmark configuration and execution of the benchmark includes all items that, to the best of our knowledge and belief, materially affect the reported results, regardless of whether such items are explicitly required to be disclosed by the SPC-1 benchmark specification.

Sincerely,

A handwritten signature in black ink that reads "Alistair L. Symon".

Alistair Symon
Vice President, XIV and Storwize Development
IBM Systems and Technology Group



SPC BENCHMARK 1™

EXECUTIVE SUMMARY

IBM IBM® STORWIZE® V5030

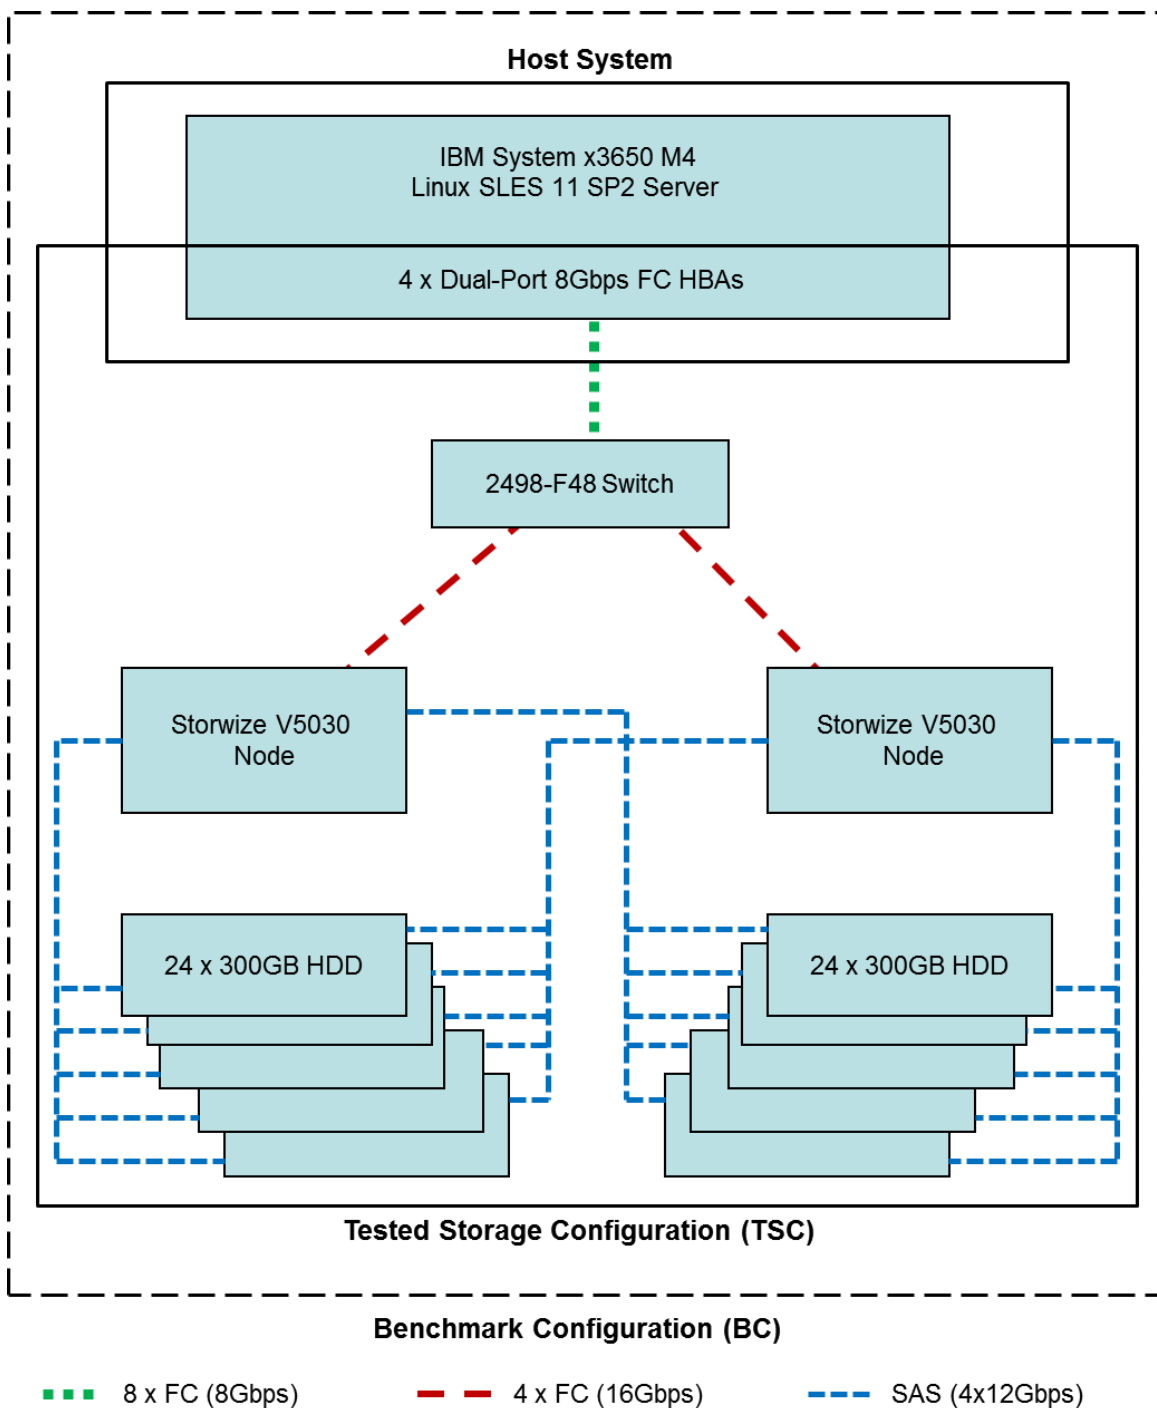
SPC-1 IOPS™	85,020
SPC-1 Price-Performance™	\$0.47/SPC-1 IOPS™
SPC-1 IOPS™ Response Time	13.358 ms
SPC-1 Overall Response Time	5.193 ms
SPC-1 ASU Capacity	34,360GB
SPC-1 ASU Price	\$12.22/GB
SPC-1 Total System Price	\$419,607.89
Data Protection Level	Protected 1 (RAID-10)
Physical Storage Capacity	71,870 GB
Pricing Currency / Target Country	U.S. Dollars / USA

SPC-1 V3.2

SUBMISSION IDENTIFIER: A31001

SUBMITTED FOR REVIEW: JULY 14, 2016

Benchmark Configuration Diagram



Tested Storage Product Description

IBM® Storwize® V5000 second-generation models deliver a range of performance, scalability,

All Storwize V5000 second-generation models include:

- I/O connectivity options for 16 Gb Fibre Channel (FC), 12 Gb SAS, 10 Gb iSCSI / Fibre Channel over Ethernet (FCoE), and 1 Gb iSCSI
- Twelve 3.5-inch large form factor (LFF) or twenty-four 2.5-inch small form factor (SFF) drive slots within the enclosure
- Support for the attachment of second-generation Storwize V5000 12 Gb SAS expansion enclosures
- Support for IBM Spectrum Virtualize™ functions, including thin provisioning, Easy Tier®, FlashCopy®, and remote mirroring
- A 2U, 19-inch rack mount enclosure with either AC or DC power
- A one- or three-year warranty.

Priced Storage Configuration Components

4 x 8Gb Dual-Port FC HBAs
1 x SAN48-5 FC switch, 24 ports, 16 SFP
2 x IBM Storwize V5030 Node, each with: 32GB Memory/Cache 4 x 16Gb FC Port with SFP (switch link) 1 x 2x12Gb SAS connection (internal/unused) 2 x 4x12Gb SAS connection (to expansion enclosures)
10 x Storwize V5030 expansion enclosure, each with: 24 x 15Krpm 300GB HDD


Storage Configuration Pricing

	Description	Qty	Unit Price	Ext. Price	Disc.	Disc. Price
	Hardware					
2078-324	V5030 base enclosure w/ 64 GB cache, 8x16 Gb FC	1	\$37,604.00	\$37,604.00	28%	\$27,074.88
2078-24F	V5030 expansion enclosure w/ 2 SAS cables	10	\$11,441.00	\$114,410.00	28%	\$82,375.20
2078-ACLB	15K RPM 300 GB drives	240	\$1,309.00	\$314,160.00	28%	\$226,195.20
2078-5305	Short wave 5m fibre channel cable	8	\$129.00	\$1,032.00	28%	\$743.04
2078-5325	Short wave 25 m fibre channel cable	8	\$189.00	\$1,512.00	28%	\$1,088.64
7014-T42	19-inch Rack	1	\$7,940.00	\$7,940.00	50%	\$3,970.00
42D0510	dual port qlogic HBA for xSeries	4	\$1,849.00	\$7,396.00		\$7,396.00
2498-F48	SAN48-5 FC switch w/24 ports active, 16 SFP	1	\$36,152.00	\$36,152.00	20%	\$28,921.60
Hardware Subtotal						\$377,764.56
	Support & Maintenance					
	V5030 Controller Enclosure (1 year, 24x7, 4h response)	3	\$2,800.00	\$8,400.00	20%	\$6,720.00
	V5030 Expansion Enclosure (1 year, 24x7, 4h response)	30	\$1,350.00	\$40,500.00	20%	\$32,400.00
	Switch (1 year, 24x7, 4h response)	3	\$1,134.72	\$3,404.16	20%	\$2,723.33
Support & Maintenance Subtotal						\$41,843.33
SPC-1 Total System Price						\$419,607.89
SPC-1 ASU Capacity						34,360GB
SPC-1 ASU Price						\$12.22/GB

Discount Details: The discounts shown are on a “field delegation” basis. A discount of 28 percent was applied for V5030 product hardware. Other product specific levels of discounting are applicable for non-V5030 hardware including the rack, switch, and HBA's.

Availability Date: Currently Available.

Response Time and Throughput Graph



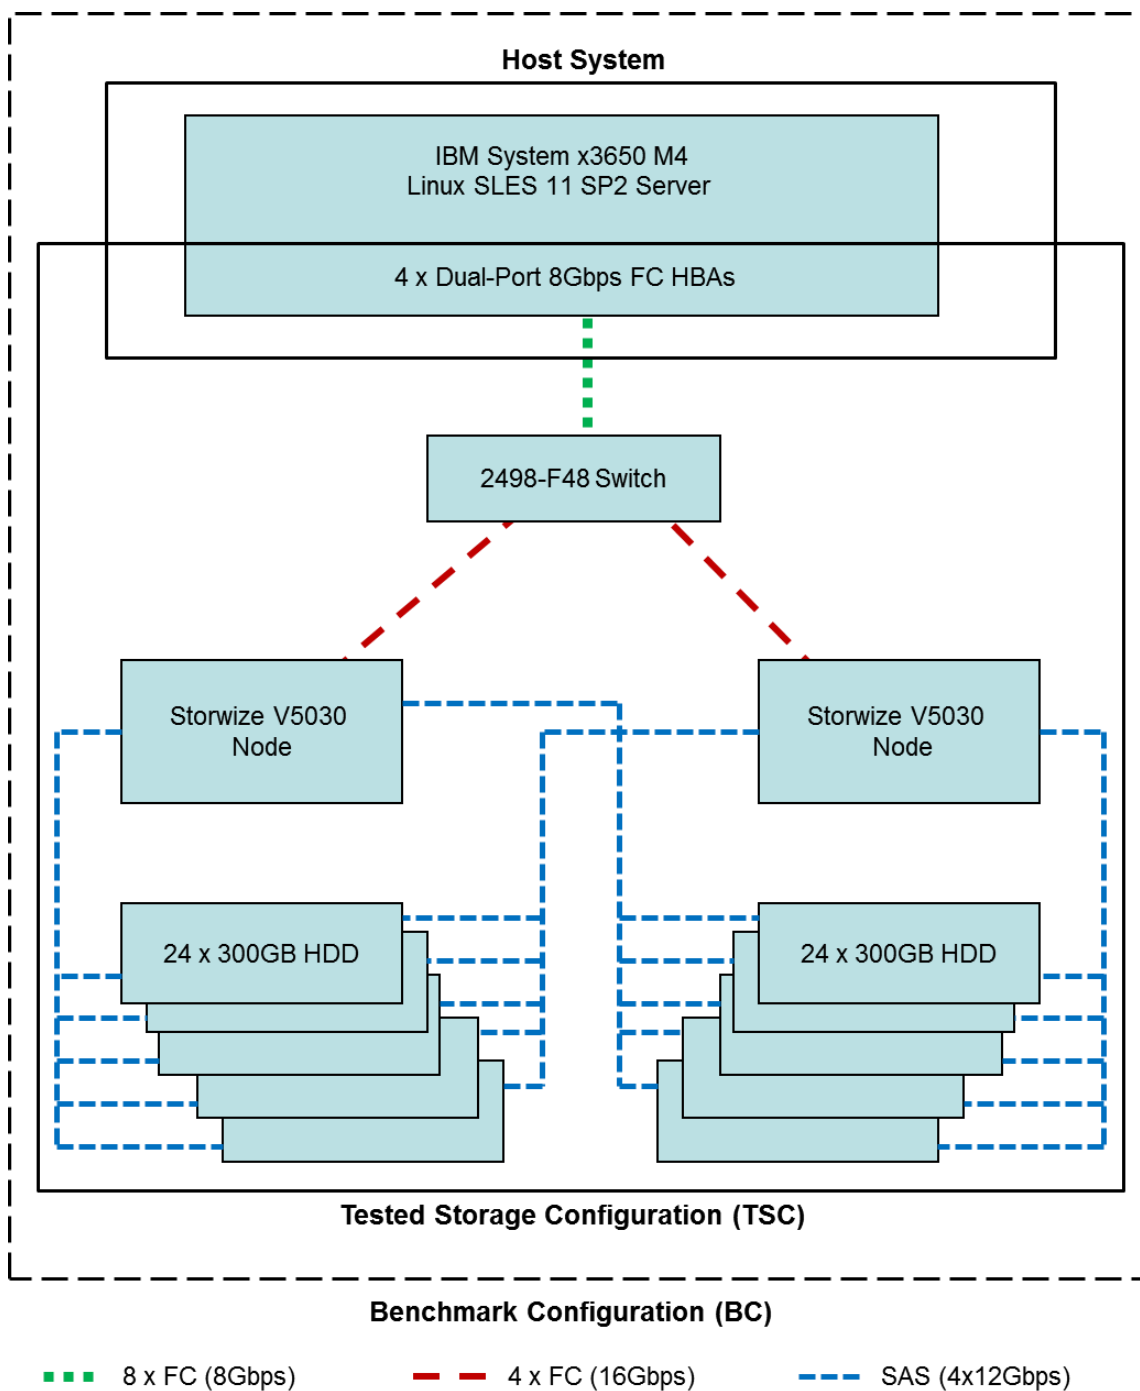
Contact Information	
Test Sponsor Primary Contact	IBM – http://www.ibm.com Bruce McNutt – bmcnutt@us.ibm.com
SPC Auditor	InfoSizing, Inc. http://www.sizing.com Francois Raab – francois@sizing.com

Revision Information	
SPC Benchmark 1™ Revision	V3.2.0
SPC-1 Workload Generator Revision	V3.0 build 2565
Publication Revision History	Initial Publication

CONFIGURATION INFORMATION

Benchmark Configuration and Tested Storage Configuration

The following diagram illustrates the Benchmark Configuration (BC), including the Tested Storage Configuration (TSC) and the Host System(s).



Storage Network Configuration

The Tested Storage Configuration (TSC) involved a 2-node storage cluster (IBM Storwize V5030) as external storage, driven by a single host system (IBM System x3650 M4). The Host system was networked with the external storage cluster via a 24-port Fibre Channel switch. The host system connected to the switch via four dual-port 8 Gbps Fibre Channel HBAs. Each storage cluster node connected to the switch via four 16 Gbps Fibre Channel small form-factor pluggable (SFP) ports.

Host System and Tested Storage Configuration Components

The following table lists the components of the Host System(s) and the Tested Storage Configuration (TSC).

Host Systems
1 x IBM System x3650 M4 1 x Intel Xeon E5-2630 / 2.3GHz / 6 Cores / 15MB L3 64GB RAM Linux SLES 11 SP2 server
Priced Storage Configuration
4 x 8Gb Dual-Port FC HBAs
1 x SAN48-5 FC switch, 24 ports, 16 SFP
2 x IBM Storwize V5030 Node, each with: 32GB Memory/Cache 4 x 16Gb FC Port with SFP (switch link) 1 x 2x12Gb SAS connection (internal/unused) 2 x 4x12Gb SAS connection (to expansion enclosures) 10 x Storwize V5030 expansion enclosure, each with: 24 x 15Krpm 300GB HDD

Differences Between Tested and Priced Storage Configurations

There were no differences between the Tested Storage Configuration and the Priced Storage Configuration.

Component Changes in Revised Full Disclosure Report

The following table outlines component changes that were made in revisions to this Full Disclosure Report.

Original Component	Revised Component	Description of Change
n/a	n/a	Initial submission

Benchmark Configuration Creation Process

Customer Tuning Parameters and Options

All the customer tuning parameters and options that have been altered from their default values for this benchmark are included in Appendix C and in the Supporting Files (see Appendix A).

Tested Storage Configuration Creation

A detailed description of how the logical representation of the TSC was created is included in Appendix D and in the Supporting Files (see Appendix A).

Tested Storage Configuration Inventory

An inventory of the layers in the logical representation of the TSC, as seen by the Benchmark Configuration, is included in Appendix E and in the Supporting Files (see Appendix A).

Workload Generator Storage Configuration

The SPC-1 Workload Generator storage configuration commands and parameters used to invoke the execution of the tests are included in Appendix F and in the Supporting Files (see Appendix A).

Logical Volume Capacity and ASU Mapping

The following table details the capacity of each ASU and how they are mapped to logical volumes (LV).

	LV per ASU	LV Capacity	Used per LV	Total per ASU	% ASU Capacity
ASU-1	36	429.50	429.50	15,461.9	45.00%
ASU-2	36	429.50	429.50	15,461.9	45.00%
ASU-3	8	429.50	429.50	3,436.0	10.00%
SPC-1 ASU Capacity				34,359.7	

Physical Storage Capacity and Utilization

The following table details the Physical Capacity of the storage devices and the Physical Capacity Utilization (percentage of Total Physical Capacity used) in support of hosting the ASUs.

Devices	Count	Physical Capacity	Total Capacity
15K RPM 300GB HDD	240	299.5	71,870.4
Total Physical Capacity			71,870.4
Physical Capacity Utilization			47.81%

Data Protection

The data protection level used for all logical volumes was **Protected 1**, which was accomplished with RAID-10 (drives are configured in ranks of 10 per RAID-10 array).

BENCHMARK EXECUTION RESULTS

This portion of the Full Disclosure Report documents the results of the various SPC-1 Tests, Test Phases, and Test Runs.

Benchmark Execution Overview

Workload Generator Input Parameters

The SPC-1 Workload Generator commands and input parameters for the Test Phases are presented in the Supporting Files (see Appendix A).


Primary Metrics Test Phases

The benchmark execution consists of the Primary Metrics Test Phases, including the Test Phases SUSTAIN, RAMPD_100 to RAMPD_10, RAMPU_50 to RAMPU_100, RAMP_0, REPEAT_1 and REPEAT_2.

Each Test Phase starts with a transition period followed by a Measurement Interval.

Measurement Intervals by Test Phase Graph

The following graph presents the average IOPS and the average Response Times measured over the Measurement Interval (MI) of each Test Phase.



Exception and Waiver

In the course of the audit of this result, it was discovered that the VERIFY_2 step called for in Clause 6.3.4.1 of the Benchmark Specification was not requested during the execution and, therefore, did not take place. After review of this exception, the requirement to execute the VERIFY_2 step was waived by the Auditor.

The SPC Policies list three criteria for waiving a requirement (see Clause 9.3.4.3). For this waiver, the criteria were evaluated as follows:

- *It has no effect, whatsoever, on the reported metrics.*

The VERIFY_2 step is called for being executed “immediately following the last performance test”. Therefore, its execution could not have any effect on the reported metrics, which were collected during the performance test.

- *It does not affect compliance with any other requirement.*

The VERIFY_2 step is only mentioned in Clause 6.3.4.1 of the SPC-1 V3.2 Specification and has no bearing on other requirements in the Benchmark Specification.

- *Its compliance would represent a significant financial or operational burden on the part of the Test Sponsor.*

The discovery of the missing VERIFY_2 step took place after the Tested Configuration had been dismantled. As a result, repeating the testing to include the VERIFY_2 step would represent a significant operational burden for the Test Sponsor.

SUSTAIN Test Phase

SUSTAIN – Results File


The results file generated during the execution of the SUSTAIN Test Phase is included in the Supporting Files (see Appendix A) as follows:

- SPC1_METRICS_0_Raw_Results.xlsx


SUSTAIN – Execution Times

Interval	Start Date	Start Time	End Time	Duration
Transition Period	9-Jun-16	01:10:49.8	01:13:48.0	02:58.2
Measurement Interval	9-Jun-16	01:13:48.0	09:13:49.0	8:00:01.0


SUSTAIN – Throughput Graph




SUSTAIN – Response Time Graph



SUSTAIN – Data Rate Graph



SUSTAIN – Response Time Frequency Graph



RAMPD_100 Test Phase

RAMPD 100 – Results File


The results file generated during the execution of the RAMPD_100 Test Phase is included in the Supporting Files (see Appendix A) as follows:

- SPC1_METRICS_0_Raw_Results.xlsx


RAMPD 100 – Execution Times

Interval	Start Date	Start Time	End Time	Duration
Transition Period	9-Jun-16	09:14:48.1	09:17:48.1	03:00.0
Measurement Interval	9-Jun-16	09:17:48.1	09:27:49.1	10:01.0


RAMPD 100 – Throughput Graph




RAMPD 100 – Response Time Graph



RAMPD 100 – Data Rate Graph



RAMPD 100 – Response Time Frequency Graph



RAMPD 100 – Intensity Multiplier

The following table lists the targeted intensity multiplier (Defined), the measured intensity multiplier (Measured) for each I/O STREAM, its coefficient of variation (Variation) and the percentage of difference (Difference) between Target and Measured.

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0018	0.0006	0.0007	0.0007	0.0023	0.0015	0.0016	0.0008
Difference	0.040%	0.003%	0.026%	0.004%	0.188%	0.003%	0.036%	0.010%

RAMPD 100 – I/O Request Summary

I/O Requests Completed in the Measurement Interval	51,006,231
I/O Requests Completed with Response Time <= 30 ms	44,913,222
I/O Requests Completed with Response Time > 30 ms	6,093,009

Response Time Ramp Test

Response Time Ramp Test – Results File


The results file generated during the execution of the Response Time Ramp Test is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_METRICS_0_Raw_Results.xlsx**


Response Time Ramp Test – Phases

The Response Time Ramp Test is comprised of 11 Test Phases, including six Ramp-Down Phases (executed at 100%, 95%, 90%, 80%, 50%, and 10% of the Business Scaling Unit) and five Ramp-Up Phases (executed at 50%, 80%, 90%, 95%, and 100% of the Business Scaling Unit).


Response Time Ramp Test – Average Throughput Graph



Response Time Ramp Test – Average Response Time Graph



Response Time Ramp Test – RAMPD 10 Response Time Graph



Repeatability Test

Repeatability Test Results File

The results file generated during the execution of the Repeatability Test is included in the Supporting Files (see Appendix A) as follows:


- SPC1_METRICS_0_Raw_Results.xlsx

Repeatability Test Results


The throughput measurements for the Response Time Ramp Test (RAMPD) and the Repeatability Test Phases (REPEAT_1 and REPEAT_2) are listed in the tables below.

Test Phase	100% IOPS	10% IOPS
RAMPD	85,011.0	8,496.5
REPEAT_1	85,011.8	8,494.5
REPEAT_2	85,025.2	8,497.0


REPEAT 1 100 – Throughput Graph




REPEAT 1 100 – Response Time Graph



REPEAT 2 100 – Throughput Graph



REPEAT 2 100 – Response Time Graph



Repeatability Test – Intensity Multiplier

The following tables lists the targeted intensity multiplier (Defined), the measured intensity multiplier (Measured) for each I/O STREAM, its coefficient of variation (Variation) and the percentage of difference (Difference) between Target and Measured.

REPEAT_1_100 Test Phase

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2809	0.0700	0.2100	0.0180	0.0701	0.0350	0.2810
Variation	0.0022	0.0007	0.0016	0.0006	0.0039	0.0015	0.0018	0.0006
Difference	0.010%	0.019%	0.038%	0.019%	0.032%	0.111%	0.042%	0.006%

REPEAT_2_100 Test Phase

	ASU1-1	ASU1-2	ASU1-3	ASU1-4	ASU2-1	ASU2-2	ASU2-3	ASU3-1
Defined	0.0350	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Measured	0.0350	0.2810	0.0700	0.2100	0.0180	0.0699	0.0350	0.2810
Variation	0.0018	0.0005	0.0014	0.0007	0.0035	0.0015	0.0030	0.0009
Difference	0.081%	0.008%	0.014%	0.016%	0.005%	0.113%	0.028%	0.005%

Data Persistence Test

Data Persistence Test Results file

The results files generated during the execution of the Data Persistence Test is included in the Supporting Files (see Appendix A) as follows:

- **SPC1_PERSIST_1_0_Raw_Results.xlsx**
- **SPC1_PERSIST_2_0_Raw_Results.xlsx**

Data Persistence Test Execution

The Data Persistence Test was executed using the following sequence of steps:

- The PERSIST_1_0 Test Phase was executed to completion.
- The Benchmark Configuration was taken through an orderly shutdown process and powered off.
- The Benchmark Configuration was powered on and taken through an orderly startup process.
- The PERSIST_2_0 Test Phase was executed to completion.

Data Persistence Test Results

Data Persistence Test Phase: Persist1	
Total Number of Logical Blocks Written	12,372,989
Total Number of Logical Blocks Verified	12,372,989
Total Number of Logical Blocks that Failed Verification	0
Time Duration for Writing Test Logical Blocks (sec.)	300
Size in bytes of each Logical Block	8,192
Number of Failed I/O Requests in the process of the Test	0

APPENDIX A: SUPPORTING FILES

The following table details the content of the Supporting Files provided as part of this Full Disclosure Report.

File Name	Description	Location
/SPC1_RESULTS	Data reduction worksheets	root
SPC1_INIT_0_Raw_Results.xlsx	Raw results for INIT Test Phase	/SPC1_RESULTS
SPC1_METRICS_0_Quick_Look.xlsx	Quick Look Test Run Overview	/SPC1_RESULTS
SPC1_METRICS_0_Raw_Results.xlsx	Raw results for Primary Metrics Test	/SPC1_RESULTS
SPC1_METRICS_0_Summary_Results.xlsx	Primary Metrics Summary	/SPC1_RESULTS
SPC1_PERSIST_1_0_Raw_Results.xlsx	Raw results for PERSIST1 Test Phase	/SPC1_RESULTS
SPC1_PERSIST_2_0_Raw_Results.xlsx	Raw results for PERSIST2 Test Phase	/SPC1_RESULTS
SPC1_Run_Set_Overview.xlsx	Run Set Overview Worksheet	/SPC1_RESULTS
SPC1_VERIFY_0_Raw_Results.xlsx	Raw results for VERIFY Test Phase	/SPC1_RESULTS
/C_Tuning	Tuning parameters and options	root
step5_linux_scheduler.sh	Setup of LUN access	/C_Tuning
/D_Creation	Storage configuration creation	root
step1_mkranks.cyg	Creating RAID-10 arrays of disks	/D_Creation
step2_mkvdisk	Creating VDIs on arrays	/D_Creation
step3_mkhost	Defining list of host Fibre Channel paths	/D_Creation
step4_mapfc_any	Mapping VDisk to Fibre Channel paths	/D_Creation
/E_Inventory	Configuration inventory	root
lsarray.txt	List of array MDIs	/E_Inventory
lsdrive.txt	list of configured storage devices	/E_Inventory
lsmdisk.txt	List of managed disks (MDIs)	/E_Inventory
lsvdisk.txt	View of volumes (VDIs)	/E_Inventory
lsportfc.txt	List of FibreChannel ports	/E_Inventory
lsportsas.txt	List of SAS ports	/E_Inventory
/F_Generator	Workload generator	root
SPC1.asu	Defining LUNs hosting the ASUs	/F_generator
overnight.sh	Executing all test phases until shutdown	/F_generator
runpersist2.sh	Executing PERSIST_2 phase after restart	/F_generator

APPENDIX B: THIRD PARTY QUOTATION

None

APPENDIX C: TUNING PARAMETERS AND OPTIONS

LUN Access Setup

The script ***step5_linux_scheduler.sh*** was executed to setup the LUN access. That script performs the following actions:

- The maximum number of concurrent system AIO requests is set to 1048576
- The driver selection for all V5030 device objects is set to “noop”
- The associated device queue depths are set to 1024

The above script is included in the Supporting Files (see Appendix A) and listed below.

step5_linux_scheduler.sh

```
echo 1048576 > /proc/sys/fs/aio-max-nr
for dmdev in `ls /dev/dm* | cut -d/ -f3`; do echo "noop" >
    /sys/block/$dmdev/queue/scheduler; done
for dmdev in `ls /dev/dm* | cut -d/ -f3`; do echo 1024 >
    /sys/block/$dmdev/queue/nr_requests; done
for sddev in `ls /dev/sd* | grep "IBM *2145" | cut -d/ -f3`
do
model=`cat /sys/block/$sddev/device/model`
smodel=`echo $model`;
if [[ $smodel == 2145 ]]; then
echo 1024 > /sys/block/$sddev/queue/nr_requests
echo "noop" > /sys/block/$sddev/queue/scheduler
fi
done
```


APPENDIX D: STORAGE CONFIGURATION CREATION

Environment

The scripts listed below with the “**.cyg**” extension require the installation of the freeware package Cygwin (www.cygwin.com). They were submitted to the Storwize V5030 using the freeware package PuTTY (www.putty.org) and executed as standard shell scripts in a Cygwin command window on the host system.

In each of these scripts, “**\$plink**” is substituted with the command “**plink name_of_cluster**”, where name_of_cluster is a saved network location for the TSC.

Creating RAID-10 Arrays and MDisks

The disks are organized into RAID-10 arrays (10 disks per array) using the script **step1_mkranks.cyg**. This results in 24 MDisks as seen by the V5030. A single storage pool (“thebiggroup”) contains all of the mdisks.

Creating the VDisks

A set of 80 VDisks is created using the capacity contained in the storage pool, using the script **step2_mkvdisk.cyg**.

Creating the host paths

Execution of the **step3_mkhost.cyg** script defines a list of host Fibre Channel paths, referred to collectively as “thea”.

Assigning the primary and alternate host paths

Each VDisk is mapped to all available Fibre Channel paths, using the script **step4_mapfc_any.cyg**.

Linux SUSE Configuration

In the linux host, the vdisks and associated Fibre Channel paths are discovered using the command “multipath”. This results in a set of LUNs with the names dm-0, dm-1, ..., dm-79, each having a capacity of 400 GiB.

The above scripts are included in the Supporting Files (see Appendix A) and listed below.

step1_mkranks.cyg

```
#!/usr/bin/bash
# run in cygwin command line
# Creates 24 RAID-10 arrays of 10 disks each, using equal numbers of disks
# from each chain
$plink svctask chnode -name lode1 1
$plink svctask chnode -name lode2 2
$plink svctask mkmdiskgrp -name thebiggroup -ext 256
drives=`$plink svcinfo lsdrive -nohdr | awk '{ print $1 }'`
```

```

for d in $drives
do
svctask chdrive -use candidate $d
done
c_enc=( -1 -1 -1 -1 -1    -2 -2 -2 -2 -2 )
#first five members are enclosure numbers of chain 1; second five of chain
2
n=0
for cnum in 2 3
do
chain=`$plink svcinfo lssasfabric -nohdr -delim : | \
grep
"^[^:]*:[^:]*:[^:]*:[^:]*:[^:]*:$cnum:[^:]*:[^:]*:[^:]*:[^:]*:lo1"
| cut -d: -f1 - | sort -n -`
for i in $chain
do
c_enc[$n]=$i
let n="n+1"
done
done
arrcount=0
s0=0
e0=0
while [[ $arrcount -le 23 ]]
do
devlist0=`for d in 0 1 2 3 4; do let s="(s0+d)%24 + 1"; let
e="e0+(s0+d)/24"; \
$plink svcinfo lsenclosureslot -slot $s ${c_enc[$e]} 2>/dev/null | \
awk '(FNR==8) { print $2 }'; done | awk -v ORS="" '{ print
(FNR==1?"":".") $1 }' `
echo $devlist0
devlist1=`for d in 0 1 2 3 4; do let s="(s0+d)%24 + 1"; let
e="e0+(s0+d)/24"; \
$plink svcinfo lsenclosureslot -slot $s ${c_enc[5+$e]} 2>/dev/null |
\
awk '(FNR==8) { print $2 }'; done | awk -v ORS="" '{ print
(FNR==1?"":".") $1 }' `
echo $devlist0:$devlist1
$plink svctask mkarray -level raid10 -drive $devlist0:$devlist1 -name
md$arrcount thebiggroup
let e0="e0+(s0+5)/24"
let s0="(s0+5)%24"
let arrcount="arrcount+1"
done

```

step2_mkvdisk.cyg

```

#!/usr/bin/bash
#execute in cygwin command line
i=0
while [[ $i -le 79 ]]
do
let lode="1 + i%2"
iogrp=0

```

```
$plink svctask mkvdisk \  
    -size 400 -unit gb -mdiskgrp thebiggroup -iogrp io_grp$iogrp \  
    -name vd$i -node lode$lode -nofmtdisk  
let i="i+1"  
done
```

step3_mkhost.cyg

```
$plink svctask mkhost -force -name thea -fcwwpn \  
21000024FF2FCDBA:21000024FF2FCDBD:21000024FF3702CD:21000024FF2FCDBC:210  
00024FF3702CC:21000024FF2FCD73:21000024FF2FCDBB:21000024FF2FCD72
```

step4_mapfc_any.cyg

```
i=0  
while [[ $i -le 79 ]]  
do  
    $plink svctask mkvdiskhostmap -force -host thea vd$i  
    let i="i+1"  
done
```

APPENDIX E: CONFIGURATION INVENTORY

An inventory of the Tested Storage Configuration was collected during the execution the script *overnight.sh* (see Appendix F). It generated the following log files:

- ***lsarray.txt*** List of array MDisks
- ***lsdrive.txt*** list of configured storage devices
- ***lsmdisk.txt*** List of managed disks (MDisks)
- ***lsvdisk.txt*** View of volumes (VDisks)
- ***lspportfc.txt*** List of FibreChanel ports
- ***lspportsas.txt*** List of SAS ports

The above log files are included in the Supporting Files (see Appendix A).

APPENDIX F: WORKLOAD GENERATOR

The ASUs accessed by the SPC-1 workload generator, are defined based upon the LUNs dm-0, dm-1, ..., dm-79, using the script ***SPC1.asu***.

The phases of the benchmark are executed using the script ***overnight.sh***, up until the end of the PERSIST_1 test phase.

The PERSIST_2 test phase is executed using the script ***runpersist2.sh***.

The above scripts are included in the Supporting Files (see Appendix A) and listed below.

SPC1.asu

```
SIZE=400gib
OFFSET=0
--
asu=1
device=/dev/dm-0
device=/dev/dm-1
device=/dev/dm-2
device=/dev/dm-3
device=/dev/dm-4
device=/dev/dm-5
device=/dev/dm-6
device=/dev/dm-7
device=/dev/dm-8
device=/dev/dm-9
device=/dev/dm-10
device=/dev/dm-11
device=/dev/dm-12
device=/dev/dm-13
device=/dev/dm-14
device=/dev/dm-15
device=/dev/dm-16
device=/dev/dm-17
device=/dev/dm-18
device=/dev/dm-19
device=/dev/dm-20
device=/dev/dm-21
device=/dev/dm-22
device=/dev/dm-23
device=/dev/dm-24
device=/dev/dm-25
device=/dev/dm-26
device=/dev/dm-27
device=/dev/dm-28
device=/dev/dm-29
device=/dev/dm-30
device=/dev/dm-31
```

```
device=/dev/dm-32
device=/dev/dm-33
device=/dev/dm-34
device=/dev/dm-35
--
asu=2
device=/dev/dm-36
device=/dev/dm-37
device=/dev/dm-38
device=/dev/dm-39
device=/dev/dm-40
device=/dev/dm-41
device=/dev/dm-42
device=/dev/dm-43
device=/dev/dm-44
device=/dev/dm-45
device=/dev/dm-46
device=/dev/dm-47
device=/dev/dm-48
device=/dev/dm-49
device=/dev/dm-50
device=/dev/dm-51
device=/dev/dm-52
device=/dev/dm-53
device=/dev/dm-54
device=/dev/dm-55
device=/dev/dm-56
device=/dev/dm-57
device=/dev/dm-58
device=/dev/dm-59
device=/dev/dm-60
device=/dev/dm-61
device=/dev/dm-62
device=/dev/dm-63
device=/dev/dm-64
device=/dev/dm-65
device=/dev/dm-66
device=/dev/dm-67
device=/dev/dm-68
device=/dev/dm-69
device=/dev/dm-70
device=/dev/dm-71
asu=3
--
device=/dev/dm-72
device=/dev/dm-73
device=/dev/dm-74
device=/dev/dm-75
device=/dev/dm-76
device=/dev/dm-77
device=/dev/dm-78
device=/dev/dm-79
```

overnight.sh

```
export PATH=/home/spc1:$PATH
ulimit -n 32768
spc1 -run SPC1_INIT -output full_run_output -iops 1600 -storage SPC1.asu
      -slave_iops 400 -iops_per_session 400
spc1 -run SPC1_VERIFY -output full_run_output -iops 100 -storage SPC1.asu
spc1 -run SPC1_METRICS -output full_run_output -iops 85000 -storage
      SPC1.asu
spc1 -run SPC1_PERSIST_1 -output full_run_output -iops 25000 -storage
      SPC1.asu
ssh -p26 perftb5hi svcinfo lsdrive -bytes -delim : > lsdrive.txt
ssh -p26 perftb5hi svcinfo lsarray -bytes -delim : > lsarray.txt
ssh -p26 perftb5hi svcinfo lsmdisk -bytes -delim : > lsmdisk.txt
ssh -p26 perftb5hi svcinfo lsvdisk -bytes -delim : > lsvdisk.txt
ssh -p26 perftb5hi svcinfo lsportfc -delim : > lsportfc.txt
ssh -p26 perftb5hi svcinfo lsportsas -delim : > lsportsas.txt
#spc1 -run SPC1_PERSIST_2 -output full_run_output -iops 25000 -storage
      SPC1.asu
```

runpersist2.sh

```
export PATH=/home/spc1:$PATH
ulimit -n 32768
spc1 -run SPC1_PERSIST_2 -output full_run_output -iops 25000 -storage
      SPC1.asu
```