SPC BENCHMARK 1™

FULL DISCLOSURE REPORT

FUSIONSTACK® CO., LTD
FUSIONSTOR® SF6000

SPC-1 V3.4.0

SUBMISSION IDENTIFIER: A31008

SUBMITTED FOR REVIEW: SEPTEMBER 5, 2017

Second Edition – March 2018

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 FusionStack 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 People's Republic of China. FusionStack 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 FusionStack representative for information on products and services available in your area.

Trademarks

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

FusionStack, the FusionStack logo and FusionStor are trademarks or registered trademarks of FusionStack in the People's Republic of China 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

InfoSizing
The Right Metric For Sizing IT

Storage Performance Council
Certified Auditor

Wang JinKai
FisionStack Co., Ltd.
3rd Floor, Building No.20
North Olympic Science & Technology Park,
Beijing, China

September 2, 2017

I verified the SPC Benchmark 1™ (SPC-1™ Revision3.4.0) test execution and performance results of the following Tested Storage Product:

FUSIONSTACK® FUSIONSTOR® SF6000

The results were:

SPC-1 IOPS™	**801,083.63**
SPC-1 Price-Performance™	**¥ 0.80/SPC-1 IOPS™**
SPC-1 IOPS™ Response Time	0.296 ms
SPC-1 Overall Response Time	0.208 ms
SPC-1 ASU Capacity	11,520 GB
SPC-1 ASU Price	¥ 55.54/GB
SPC-1 Total System Price	¥ 639,742.00

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 3.0.2 Build g823a. 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 **A31008**.

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:

- None.

Respectfully Yours,

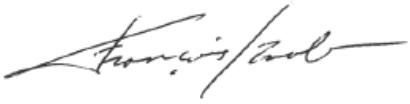

François Raab, Certified SPC Auditor

20 Kreg Lane • Manitou Springs, CO 80829 • 719-473-7555 • www.sizing.com

LETTER OF GOOD FAITH

FUSION 华云
STACK 网际

September 1, 2017

To: Francois Raab, SPC Auditor
InfoSizing
20 Kreg Ln.
Manitou Springs, CO 80829

Subject: SPC-1 Letter of Good Faith for the FusionStor® SF6000

FusionStack® is the SPC-1 test sponsor for the above listed product. To the best of our knowledge and belief, the required SPC-1 results and materials we have submitted for that product are complete, accurate, and in full compliance with version 3.4 of the SPC-1 benchmark specification.

In addition, we have reported any items in the Benchmark Configuration and execution of the benchmark that affected the reported results even if the items are not explicitly required to be disclosed by the SPC-1 benchmark specification.

Sincerely,

王金开

2017, 9, 1

WANG JinKai
CTO & Co-founder
FusionStack®

September 1, 2017

北京市海淀区宝盛南路奥北科技园 20 号楼 3 层 010-62920010 网址：www.fusionstack.cn

SPC BENCHMARK 1™

EXECUTIVE SUMMARY

FUSIONSTACK® CO., LTD
FUSIONSTOR® SF6000

SPC-1 IOPS™	**801,083**
SPC-1 Price-Performance™	**$ 121.73/SPC-1 KIOPS™**
SPC-1 IOPS™ Response Time	0.296 ms
SPC-1 Overall Response Time	0.208 ms
SPC-1 ASU Capacity	11,520 GB
SPC-1 ASU Price	$ 8.47/GB
SPC-1 Total System Price	$ 97,509.95
Data Protection Level	Protected 2 (mirroring and full redundancy)
Physical Storage Capacity	32,600 GB
Pricing Currency / Target Country	USD / People's Republic of China

SPC-1 V3.4.0

SUBMISSION IDENTIFIER: A31008

SUBMITTED FOR REVIEW: SEPTEMBER 5, 2017

Benchmark Configuration Diagram

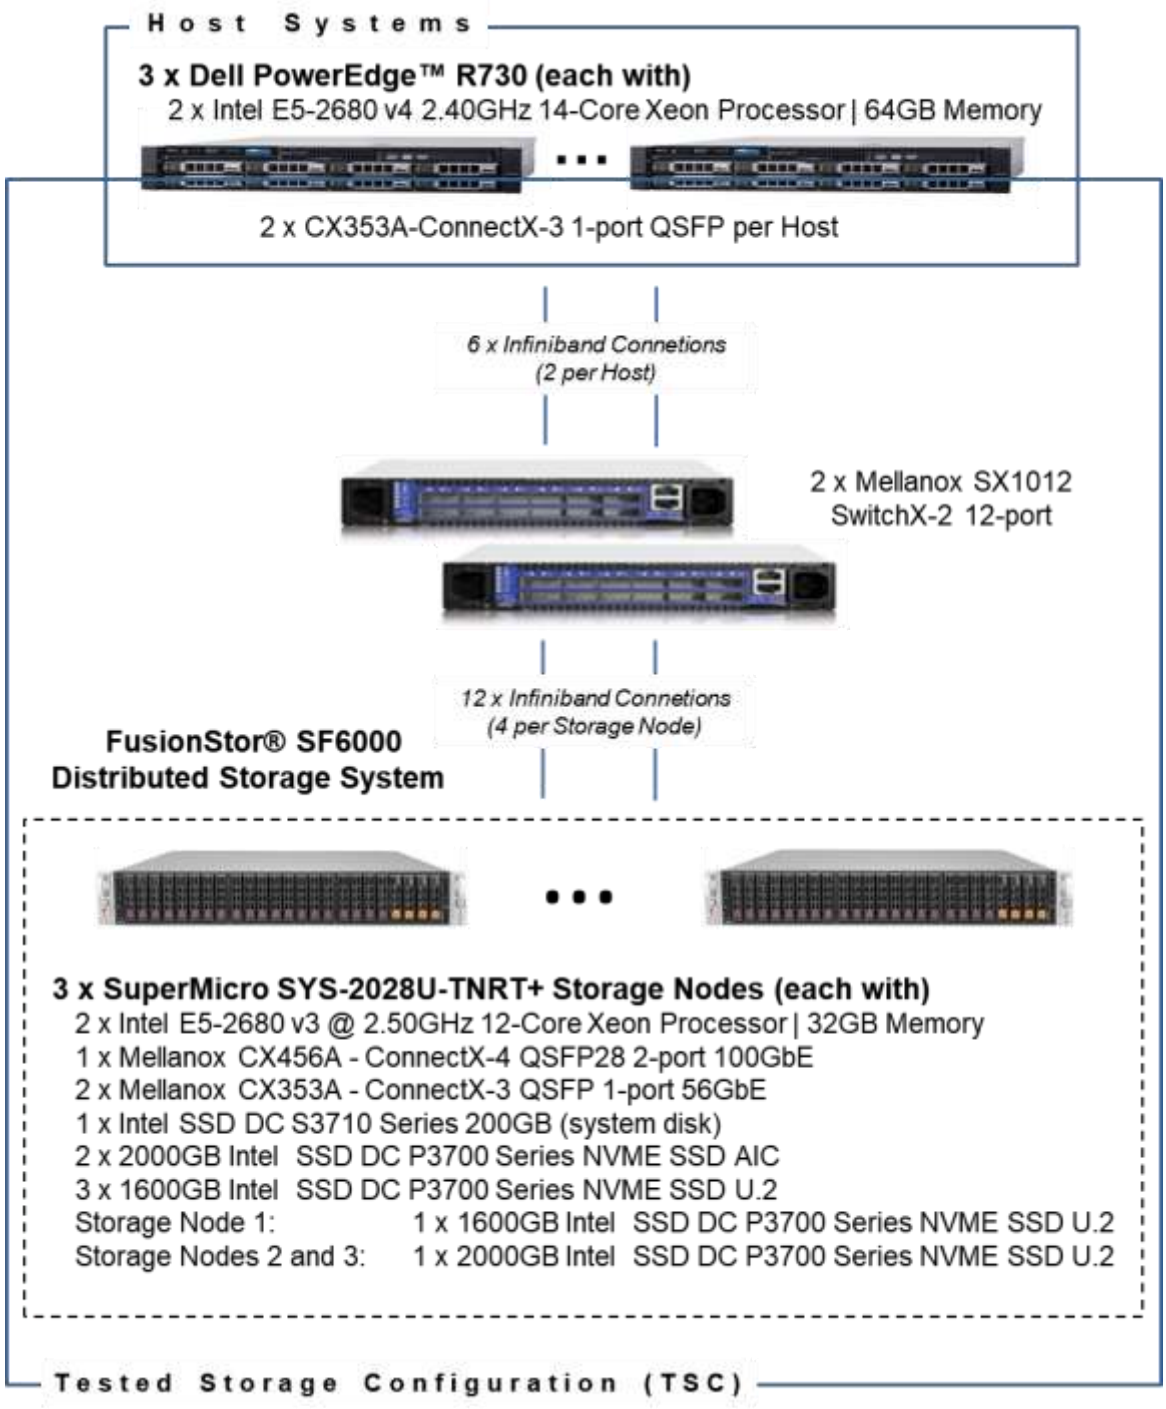

Tested Storage Product Description

FusionStor is a software solution offering distributed, high availability, high performance shared storage resource pool with linear scalability and easy management.

The latest release for FusionStor SF6000 is version 4.0. FusionStack owns full copyrights to the software. It is distributed and running on commodity x86 servers and is optimized for most types of storage medias.

FusionStor supports Ethernet and Infiniband system interconnects and most storage system interface standards; such as iSCSI, NBD, iSER, NVMe-oF, Cinder, etc.

For more details, visit:

http://www.fusionstack.cn/fusionstor_en?_l=en

Priced Storage Configuration Components

6 x CX353A-ConnectX-3 1-port QSFP
3 x SuperMicro SYS-2028U-TNRT+ Storage Nodes (each with) 2 x Intel E5-2680 v3 @ 2.50GHz 12-Core Xeon Processor | 32GB Memory 1 x Mellanox CX456A - ConnectX-4 QSFP28 2-port 100GbE 2 x Mellanox CX353A - ConnectX-3 QSFP 1-port 56GbE 1 x Intel SSD DC S3710 Series 200GB (system disk) 2 x 2000GB Intel SSD DC P3700 Series NVME SSD AIC 3 x 1600GB Intel SSD DC P3700 Series NVME SSD U.2 1 x 1600GB Intel SSD DC P3700 Series NVME SSD U.2 (in Storage Node 1) 1 x 2000GB Intel SSD DC P3700 Series NVME SSD U.2 (in Storage Nodes 2 & 3)

Storage Configuration Pricing

	Description	Source	Qty	Unit Price	Ext. Price	Disc.	Disc. Price
Hardware & Software							
SYS-2028U-TNRT+	SuperMicro SuperServer 2028U-TNRT+	2	3	2,713.08	8,139.24		8,139.24
CM8064401439612	Intel® Xeon E5-2680 v3	2	6	1,554.69	9,328.14		9,328.14
M393A2K40BB0-CPB	DDR4 RECC 16GB 2400MT/s	2	12	152.42	1,829.04		1,829.04
SSDSC2BA200G401	Intel® DC S3710 200GB	3	3	205.77	617.31		617.31
SSDPEDMD020T401	Intel® DC P3700 2.0TB AIC	3	6	2,865.50	17,193.00		17,193.00
SSDPE2MD020T401	Intel® DC P3700 2.0TB U.2	4	2	2,895.98	5,791.96		5,791.96
SSDPE2MD016T401	Intel® DC P3700 1.6TB U.2	4	10	2,277.16	22,771.60		22,771.60
MSX1012B-2BFS	SwitchX®-2 based 40GbE, 1U Open Ethernet Switch with MLNX-OS, 12 QSFP+ ports, 2 Power Supplies (AC), short depth, PPC460, P2C airflow, 12 ports licenses.	5	2	6,096.80	12,193.60		12,193.60
MCX456A-ECAT	Mellanox CX456A- ConnectX-4 QSFP28 2-port	5	3	1,656.51	4,969.53		4,969.53
MCX353A-FCBT	Mellanox CX353A- ConnectX-3 QSFP 1-port	5	6	757.23	4,543.38		4,543.38
MC2207128-002	Mellanox passive copper cable	5	12	82.31	987.72		987.72
FS-S6000-E10-S	FusionStor License per TB	1	32	228.63	7,316.16	50%	3,658.08
					Hardware & Software Subtotal		**92,022.60**
Support & Maintenance							
TS-FS-S6000-E10-S	3Years Installation,Upgrade,Support Service 24x7 4 Hour Rsponse(per TB)	1	56	22.87	1,280.72	50%	640.36
OS4HR3	SuperServer 2028U-TNRT+ 3-year onsite 24x7x4 service	2	3	640.17	1,920.51		1,920.51
SUP-SX1012-3S-4H	Mellanox SUP-SX1012-3S-4H Technical Support and Warranty - Silver 3 Year with 4 Hours On-Site Support	5	2	1,463.24	2,926.48		2,926.48
					Support & Maintenance Subtotal		**5,487.35**
SPC-1 Total System Price							**97,509.95**
SPC-1 IOPS™							801,083.63
SPC-1 Price-Performance™ ($/SPC-1 KIOPS™)							**121.73**
SPC-1 ASU Capacity (GB)							11,520
SPC-1 ASU Price ($/GB)							**8.47**

Third-Party Reseller: FusionStack Co., Ltd (Source 1) provides a software solution for distributed storage systems. All hardware components are provided through the following third-party resellers:

- Source 2 Beijing fuyangweixin Technology Co., Ltd
- Source 3 Digital China Macao Service Co., Ltd
- Source 4 Shenzhen Baotongzhiyuan Technology Co., Ltd
- Source 5 Sainuo Xinzhi Software Technology Co., Ltd

The above reflects the pricing quoted by these third-party resellers. See Appendix B of the Full Disclosure Report for a copy of the third-party resellers' quotations.

Discount Details: The discounts shown are based on enabled storage capacity and are generally available.

Warranty: The TSC includes components from major distributors or vendors. The distributors are responsible for first line support and the vendors are responsible for second line support. Those distributors and vendors have the capability to cover all customers in the pricing areas. FusionStack provides support for the software solution. The maintenance and support agreements from FusionStack and from the hardware distributors include a 3-year term with 24x7 coverage, and a 4-hour response time commitment.

Availability Date: Currently available.

Response Time and Throughput Graph

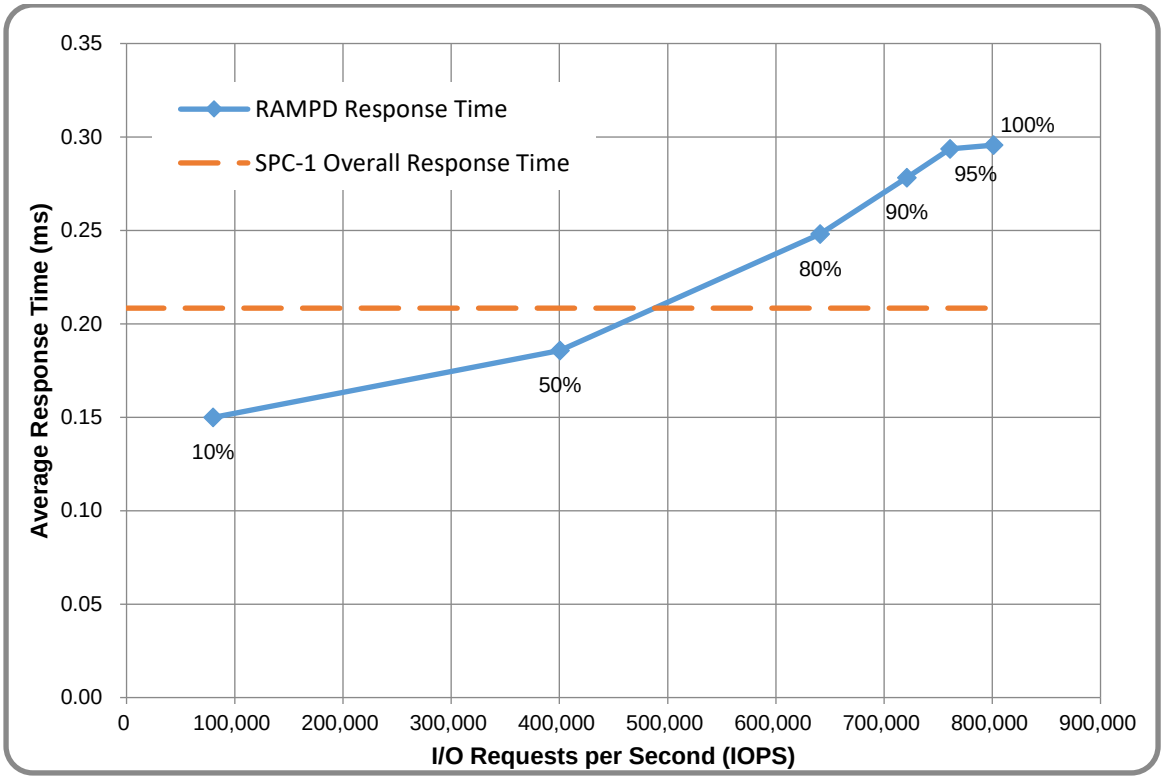

Contact Information	
Test Sponsor Primary Contact	FusionStack Co., Ltd – www.fusionstack.cn Zhang Kecheng – zhangkecheng@fusionstack.cn
SPC Auditor	InfoSizing – www.sizing.com Francois Raab – francois@sizing.com

Revision Information	
SPC Benchmark 1™ Revision	V3.4.0
SPC-1 Workload Generator Revision	V3.0.2 build g823a
Publication Revision History	• First Edition: September 5, 2017 • Second Edition: March 2, 2018 - Updated SPC-1 Price-Performance™ metric based on SPC-1 v3.6.0 definition. - Converted pricing to USD.

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).

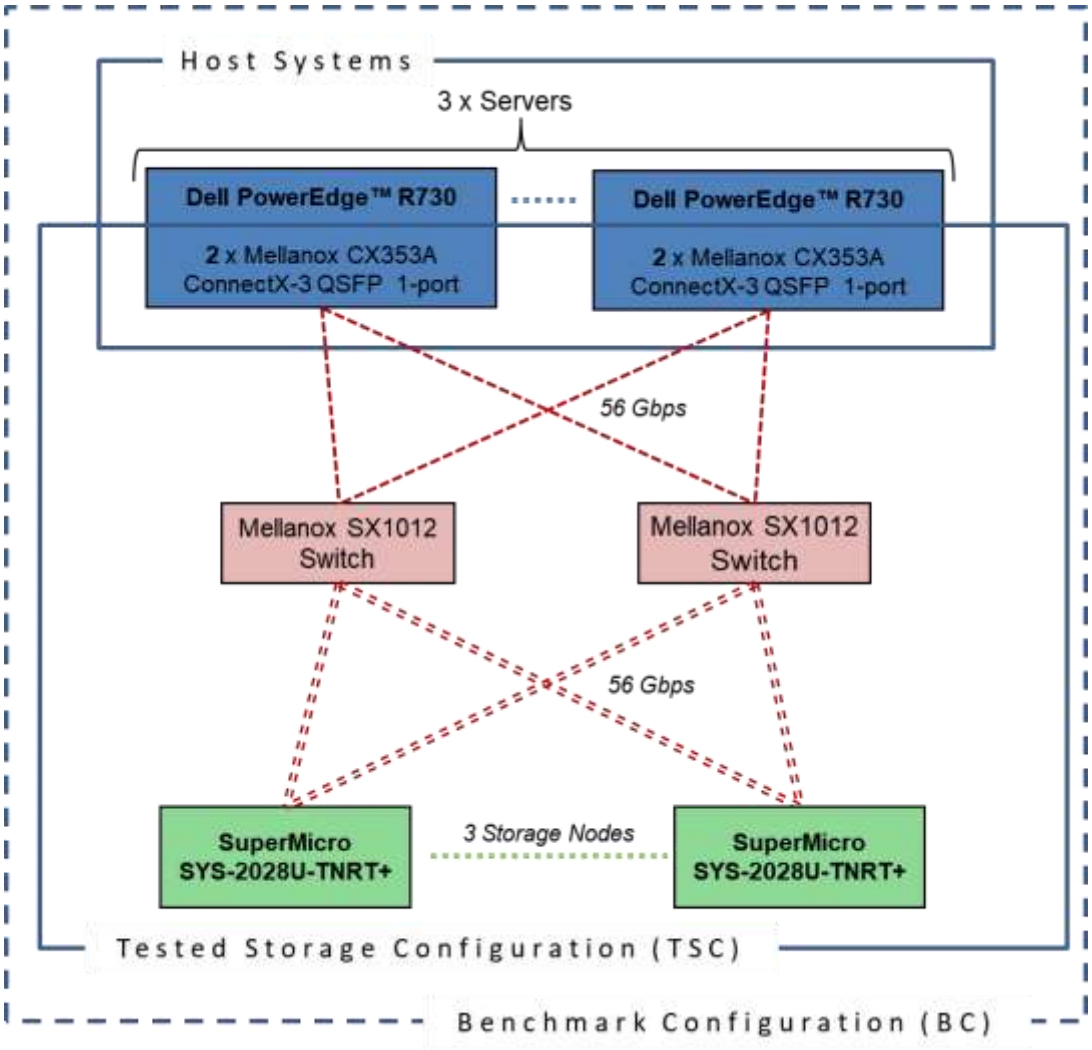

Storage Network Configuration

The Tested Storage Configuration (TSC) involved three FusionStor SF6000 Storage Nodes, driven by three Host Systems (Dell PowerEdge R730). Each R730 host was connected one-to-one to two Mellanox SX1012 Switches operating at 56Gbps.

Each of the three Storage Nodes were connected one-to-one to the two Mellanox SX1012 Switches using dual ports operating at 56Gbps.

Host System and Tested Storage Configuration Components

The FusionStor SF6000 software was deployed on the three FusionStor Storage Nodes.

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

Host Systems
3 x Dell PowerEdge R730 servers (each with) 2 x Intel E5 2680 V4 2.4 GHz 14-Core Xeon Processors 64GB Memory CentOS Linux release 7.2.1511
Tested Storage Configuration
6 x CX353A-ConnectX-3 1-port QSFP
3 x SuperMicro SYS-2028U-TNRT+ Storage Nodes (each with) 2 x Intel E5-2680 v3 @ 2.50GHz 12-Core Xeon Processor | 32GB Memory 1 x Mellanox CX456A - ConnectX-4 QSFP28 2-port 100GbE 2 x Mellanox CX353A - ConnectX-3 QSFP 1-port 56GbE 1 x Intel SSD DC S3710 Series 200GB (system disk) 2 x 2000GB Intel SSD DC P3700 Series NVME SSD AIC 3 x 1600GB Intel SSD DC P3700 Series NVME SSD U.2 1 x 1600GB Intel SSD DC P3700 Series NVME SSD U.2 (in Storage Node 1) 1 x 2000GB Intel SSD DC P3700 Series NVME SSD U.2 (in Storage Nodes 2 & 3)

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 components in 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	18	288.0	288.0	5,184.0	45.00%
ASU-2	18	288.0	288.0	5,184.0	45.00%
ASU-3	18	64.0	64.0	1,152.0	10.00%
			SPC-1 ASU Capacity	**11,520.0**	

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
2,000GB NVMe SSD	8	2,000.0	16,000.0
1,600GB NVMe SSD	10	1,600.0	16,000.0
200GB SSD	3	200.0	600.0
		Total Physical Capacity	**32,600.0**
		Physical Capacity Utilization	**35.33%**

Data Protection

The data protection level used for all logical volumes was **Protected 2**, which was accomplished by replicating all data on two separate storage nodes. In addition, all components and access paths from the Host Systems to the FusionStor storage nodes were redundant.

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.

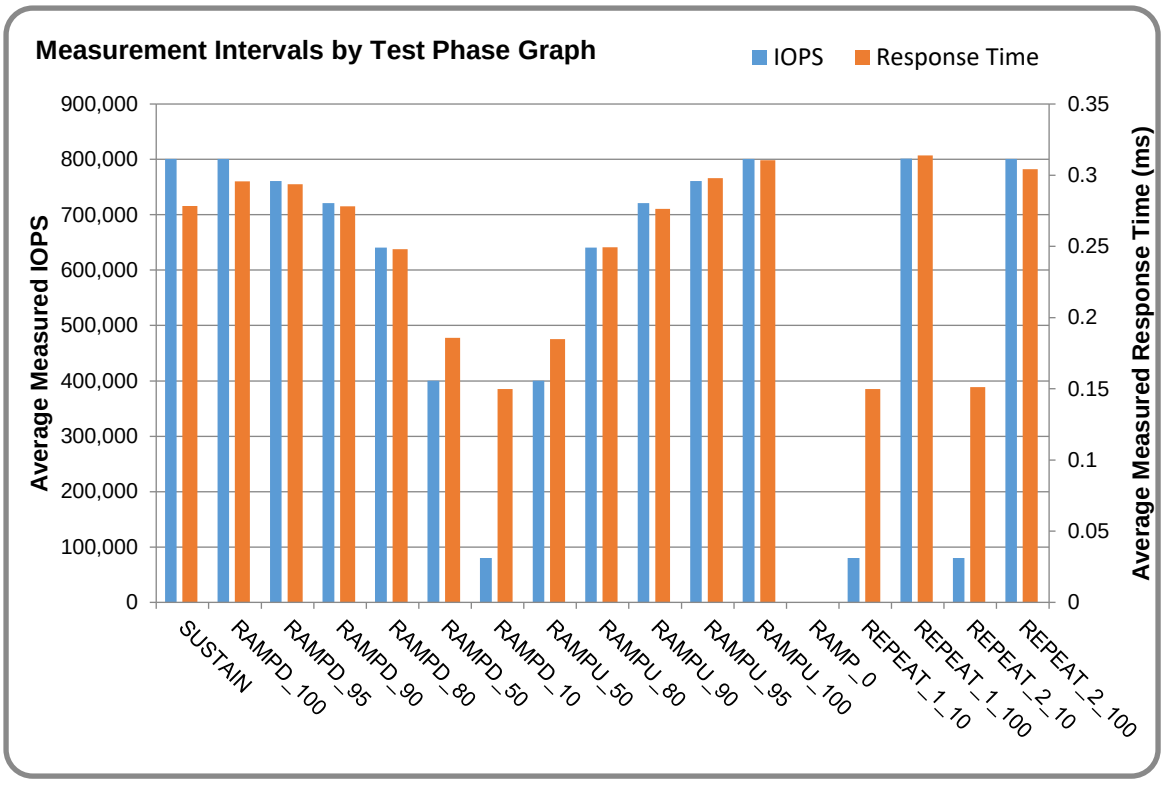

Exception and Waiver

During the course of the benchmark audit, no exceptions were encountered and no benchmark requirements were waived.

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 & Time	End Date &Time	Duration
Transition Period	19-Aug-17 00:31:14	19-Aug-17 01:21:12	0:49:57
Measurement Interval	19-Aug-17 01:21:12	19-Aug-17 09:21:13	8:00:01

SUSTAIN – Throughput Graph

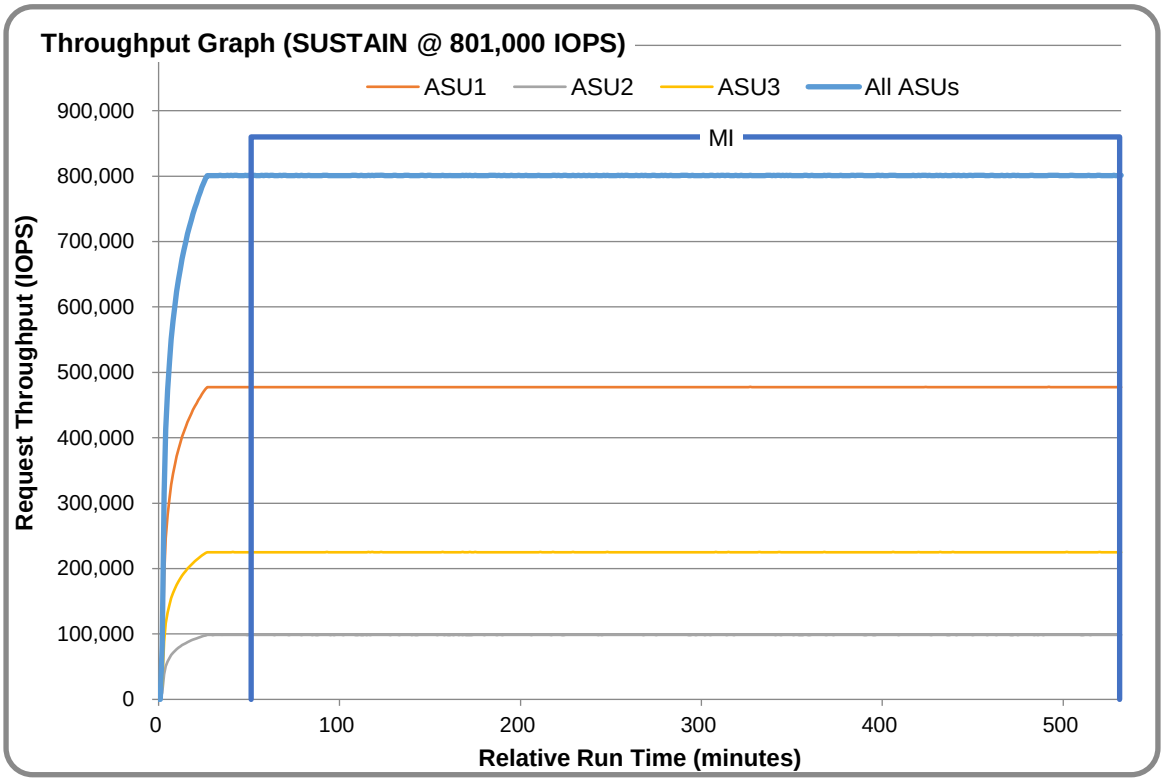

SUSTAIN – Response Time Graph

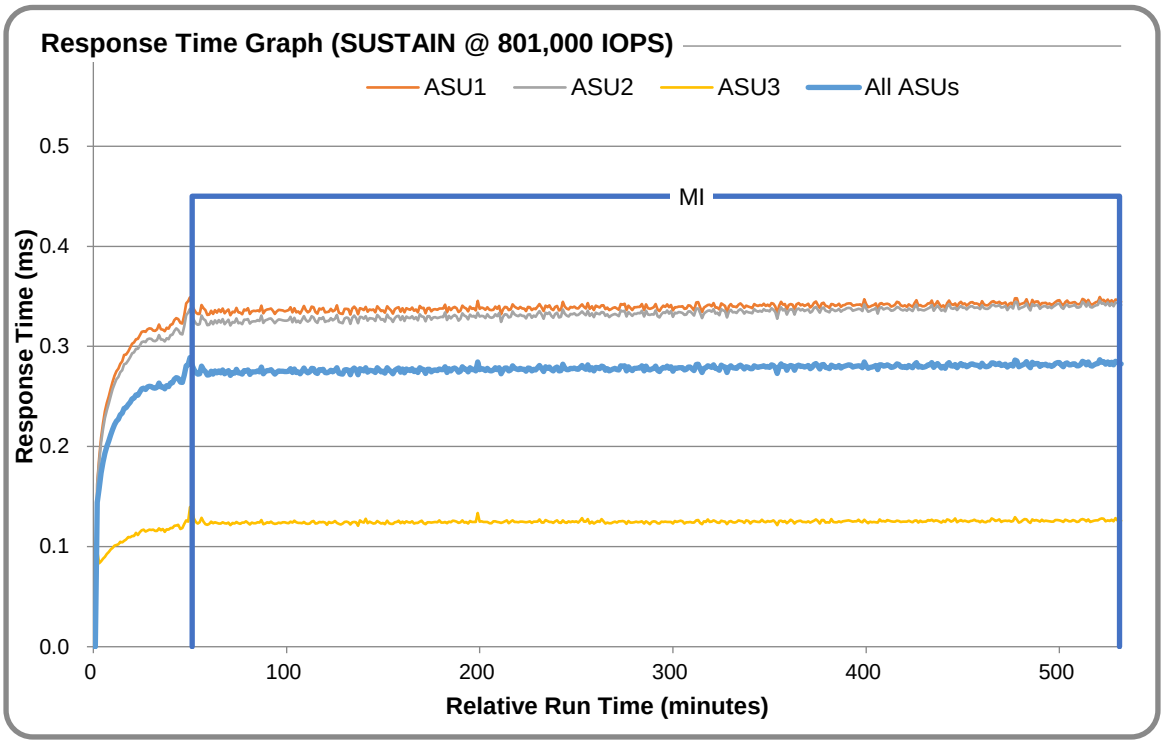

SUSTAIN – Data Rate Graph

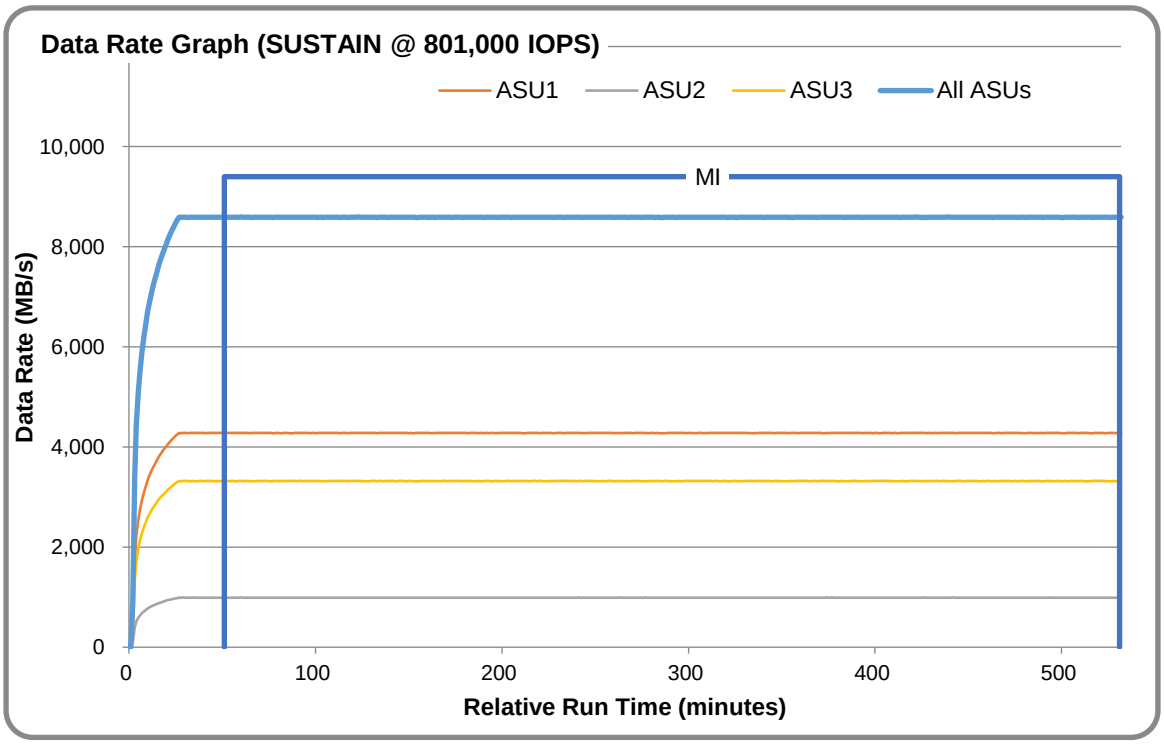

SUSTAIN – Response Time Frequency Graph

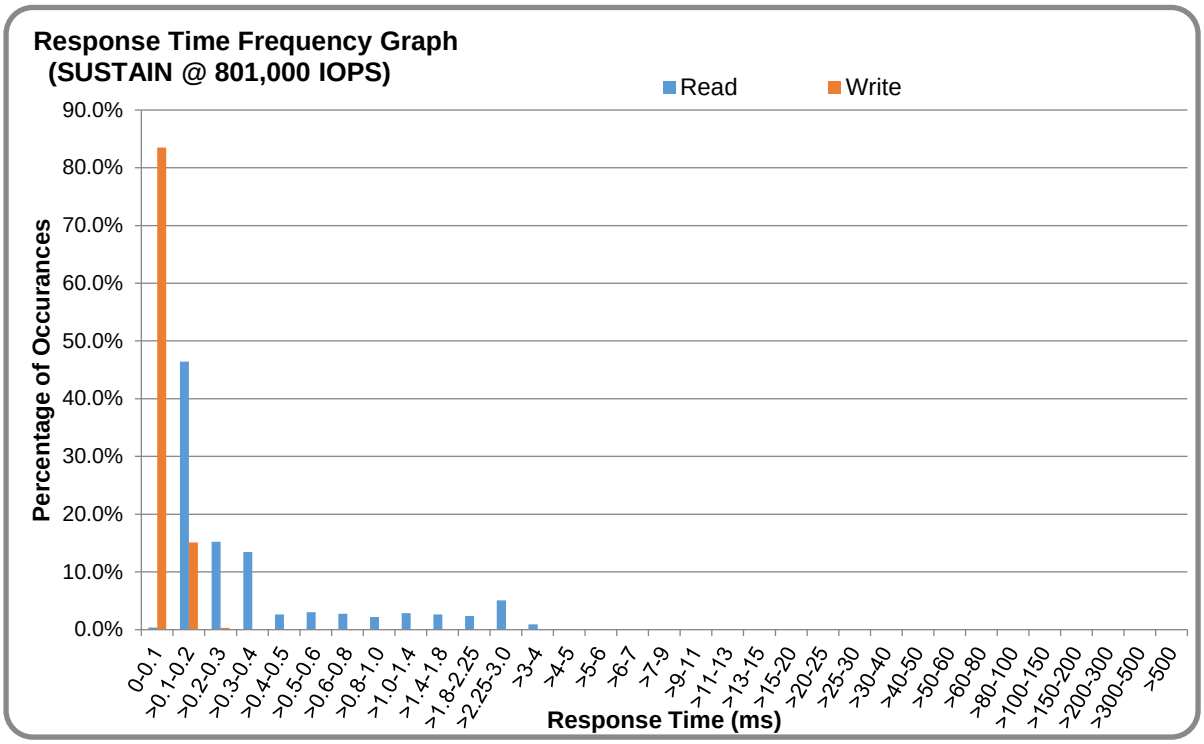

SUSTAIN – Intensity Multiplier

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	0.0007	0.0002	0.0006	0.0003	0.0012	0.0005	0.0007	0.0002
Difference	0.003%	0.002%	0.000%	0.000%	0.008%	0.010%	0.003%	0.002%

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 & Time	End Date &Time	Duration
Transition Period	19-Aug-17 09:22:12	19-Aug-17 09:25:12	0:03:00
Measurement Interval	19-Aug-17 09:25:12	19-Aug-17 09:35:13	0:10:00

RAMPD_100 – Throughput Graph

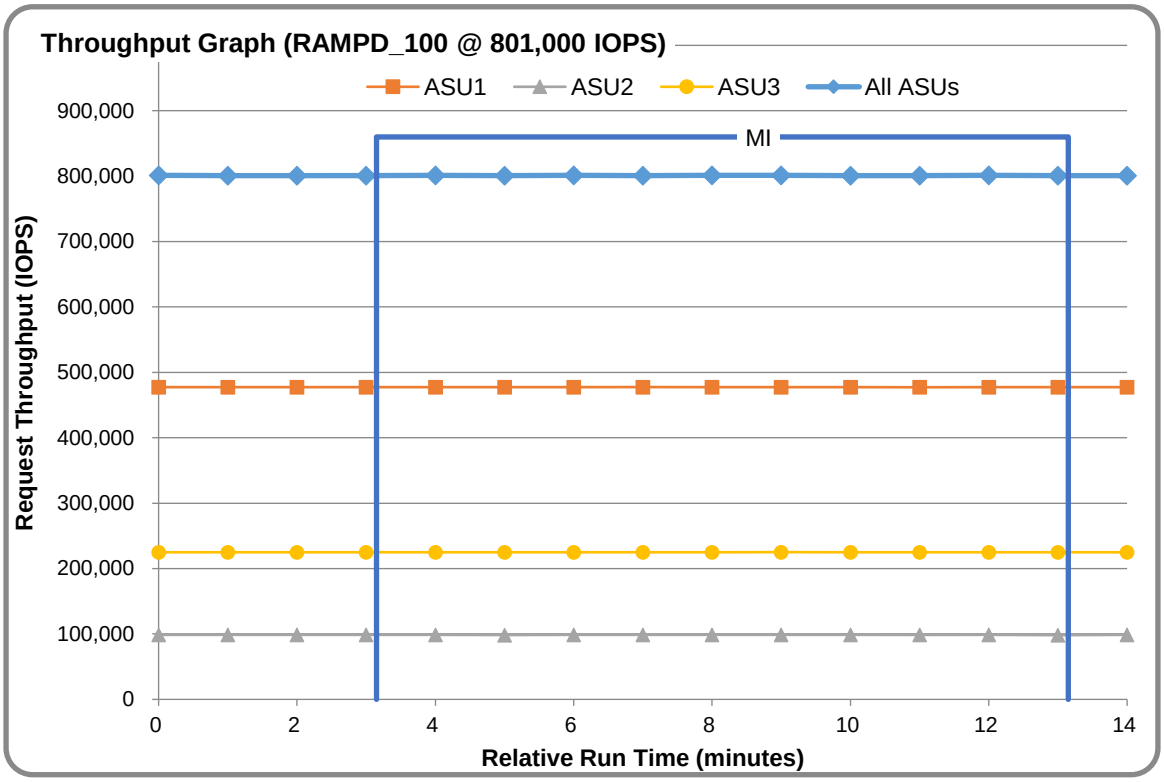

RAMPD_100 – Response Time Graph

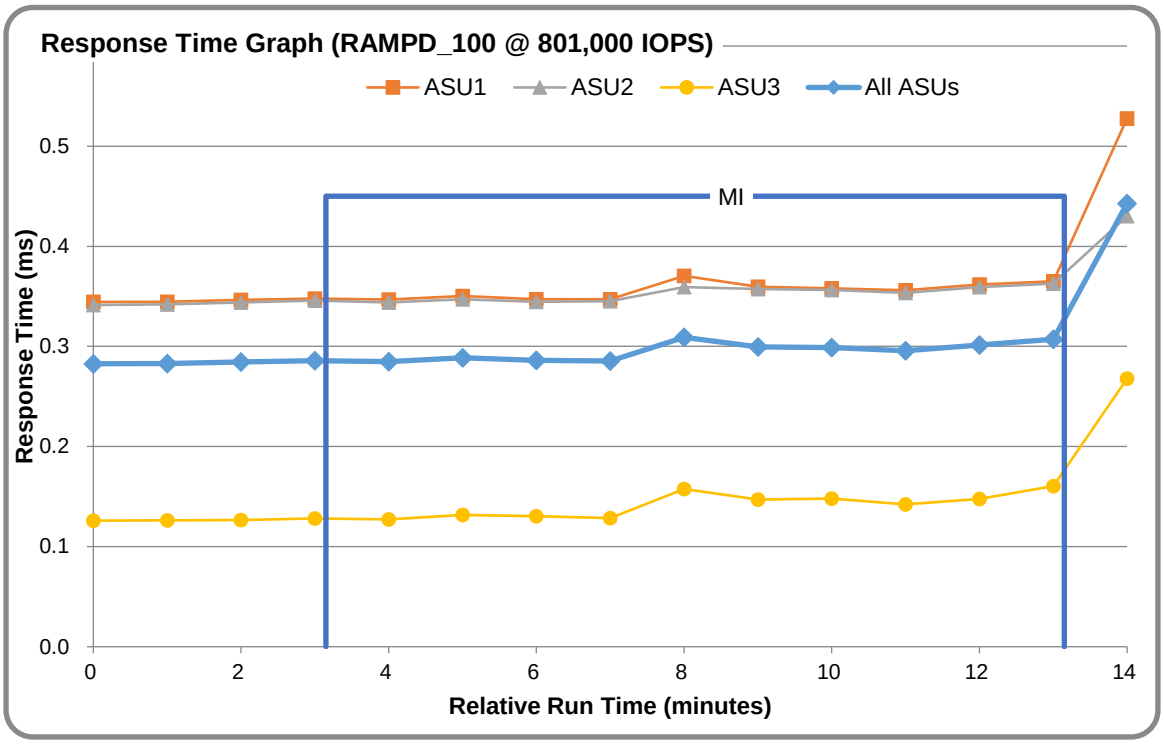

RAMPD_100 – Data Rate Graph

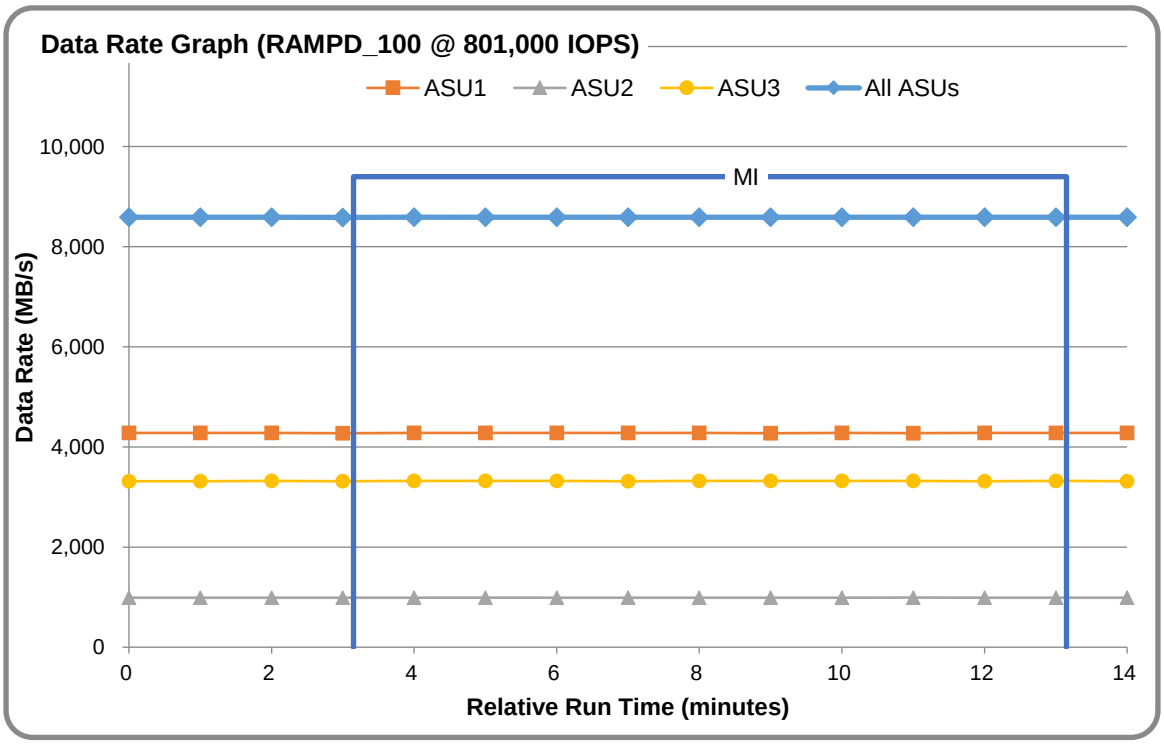

RAMPD_100 – Response Time Frequency Graph

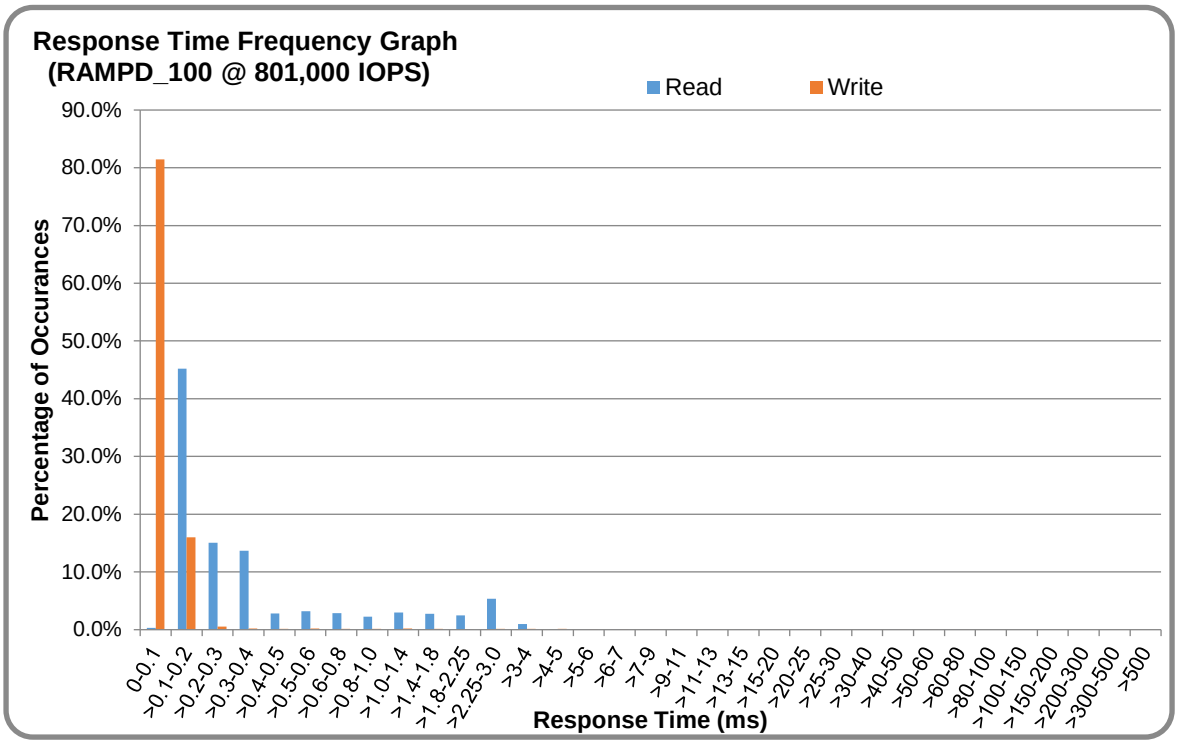

RAMPD_100 – Intensity Multiplier

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	0.0007	0.0004	0.0006	0.0004	0.0012	0.0008	0.0007	0.0002
Difference	0.030%	0.004%	0.009%	0.009%	0.035%	0.033%	0.028%	0.006%

RAMPD_100 – I/O Request Summary

I/O Requests Completed in the Measurement Interval	480,650,018
I/O Requests Completed with Response Time <= 30 ms	480,642,362
I/O Requests Completed with Response Time > 30 ms	7,656

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

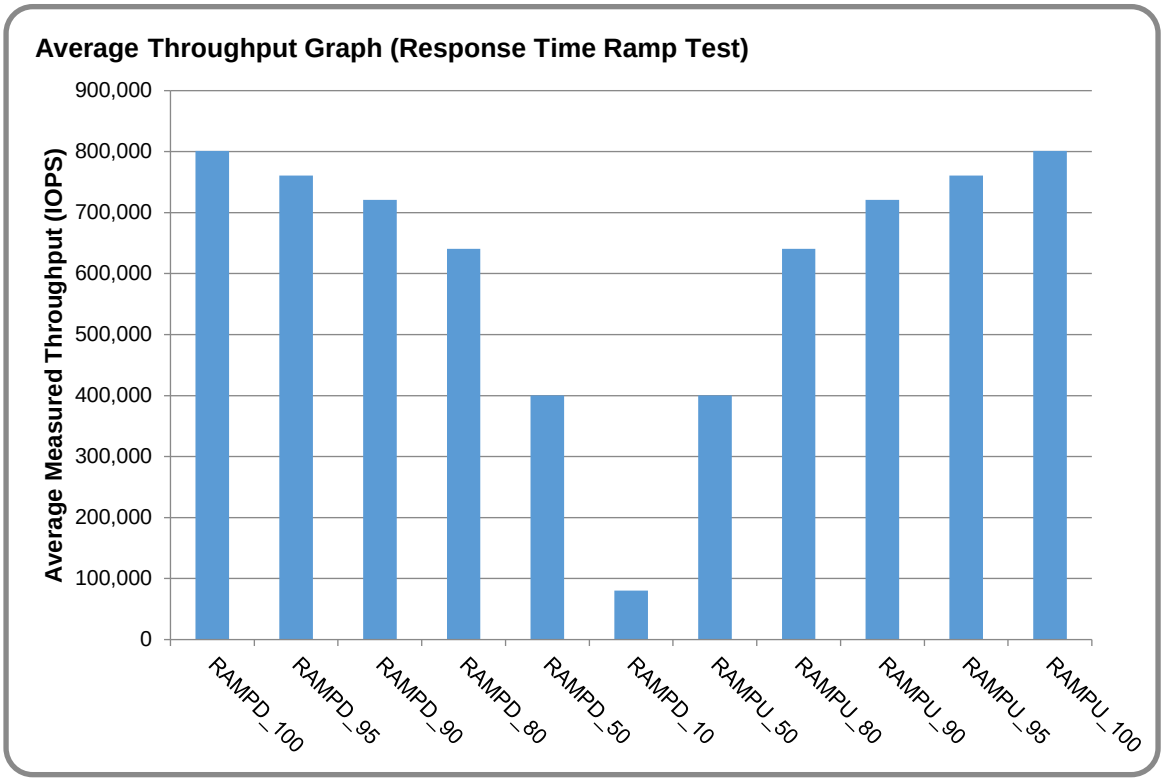

Response Time Ramp Test – Average Response Time Graph

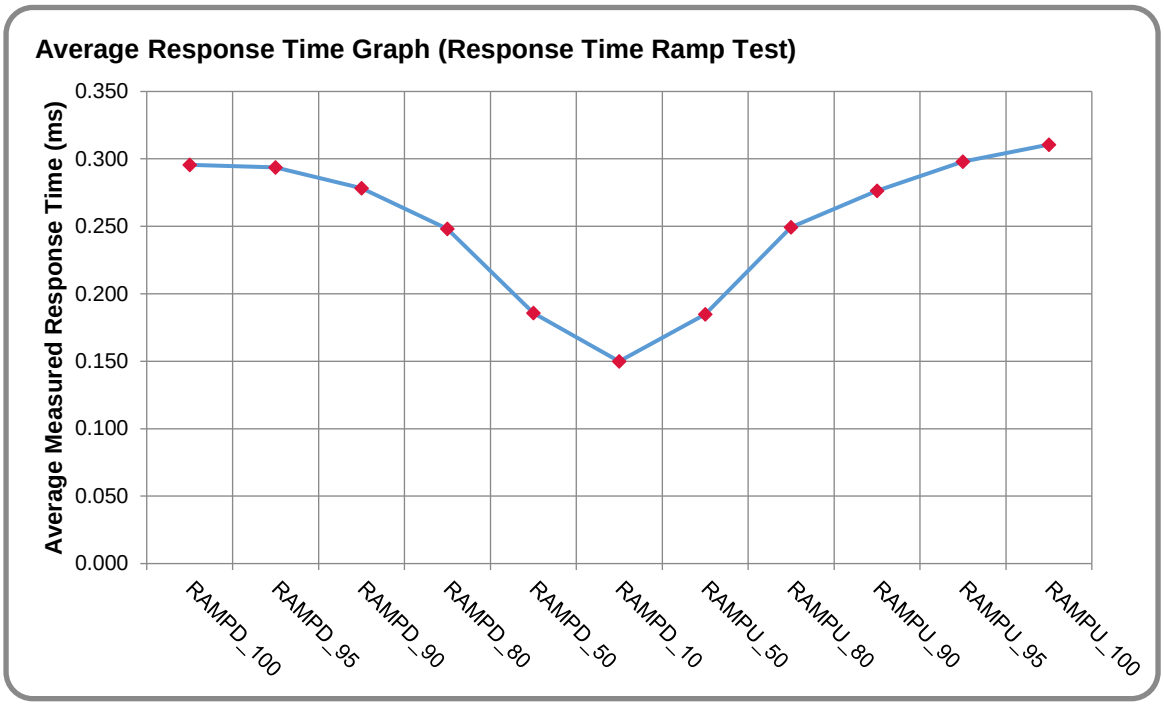

Response Time Ramp Test – RAMPD_10 Response Time Graph

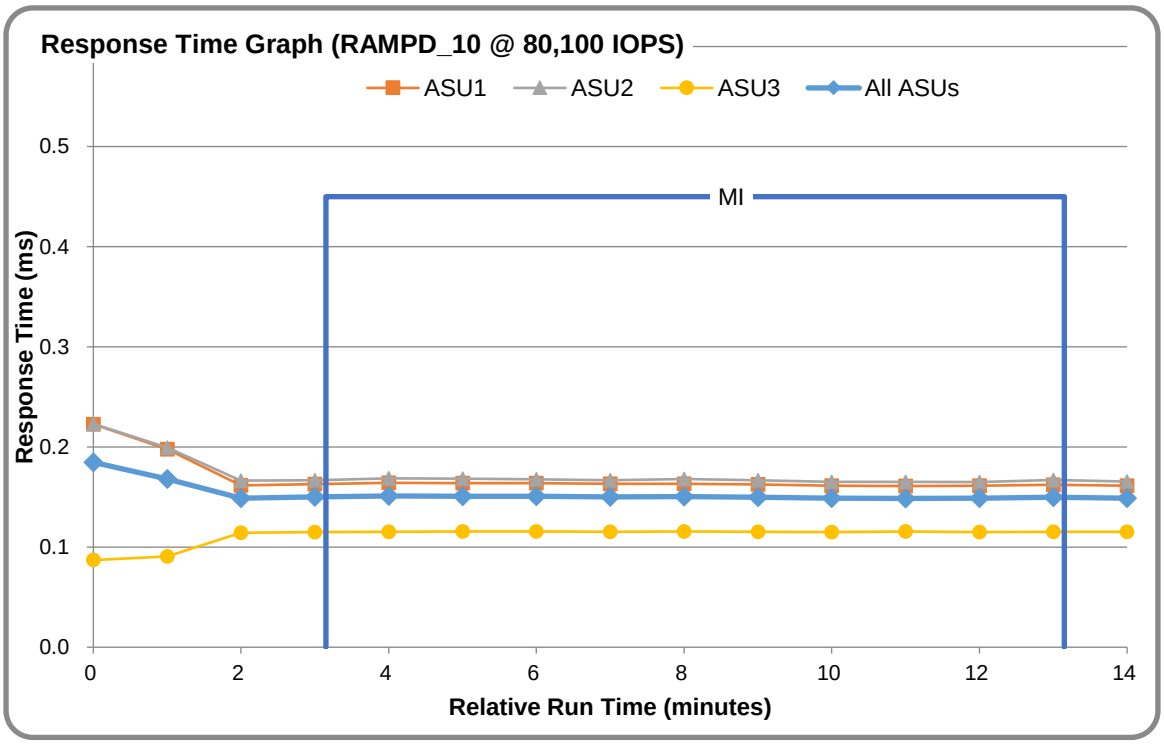

Repeatability Test

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	801,083.6	80,099.3
REPEAT_1	801,099.8	80,100.0
REPEAT_2	801,022.5	80,113.2

REPEAT_1_100 – Throughput Graph

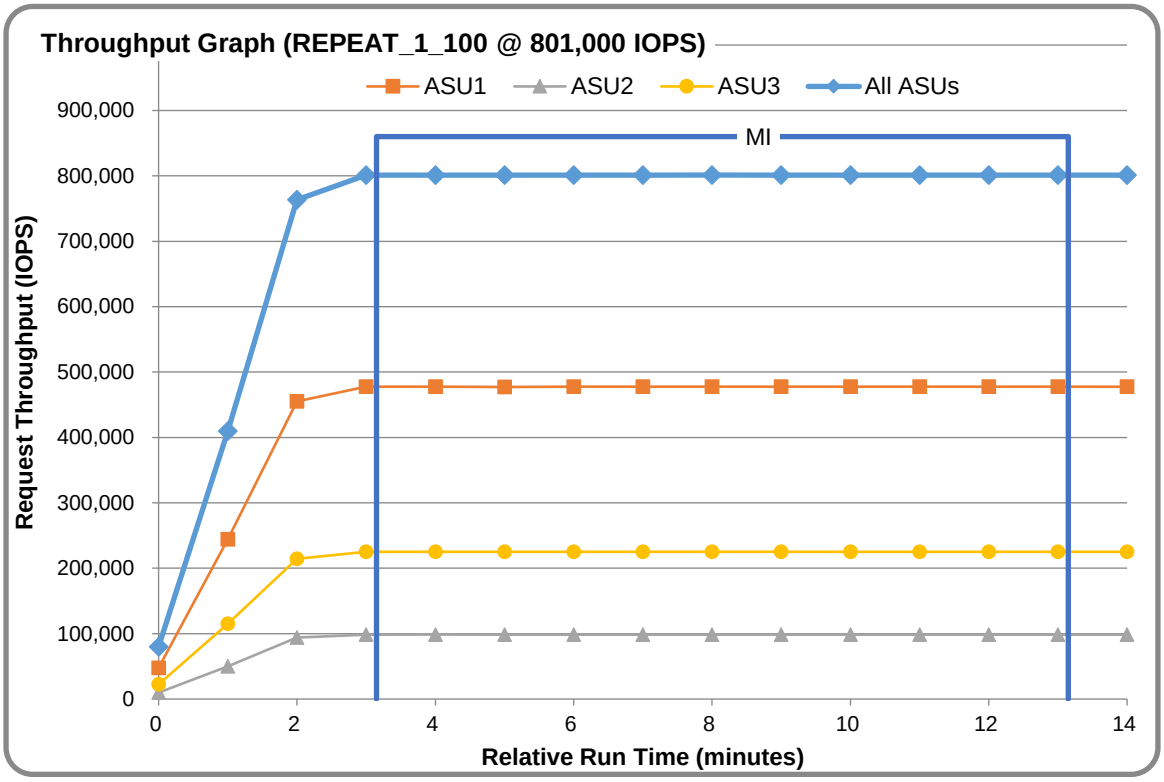

REPEAT_1_100 – Response Time Graph

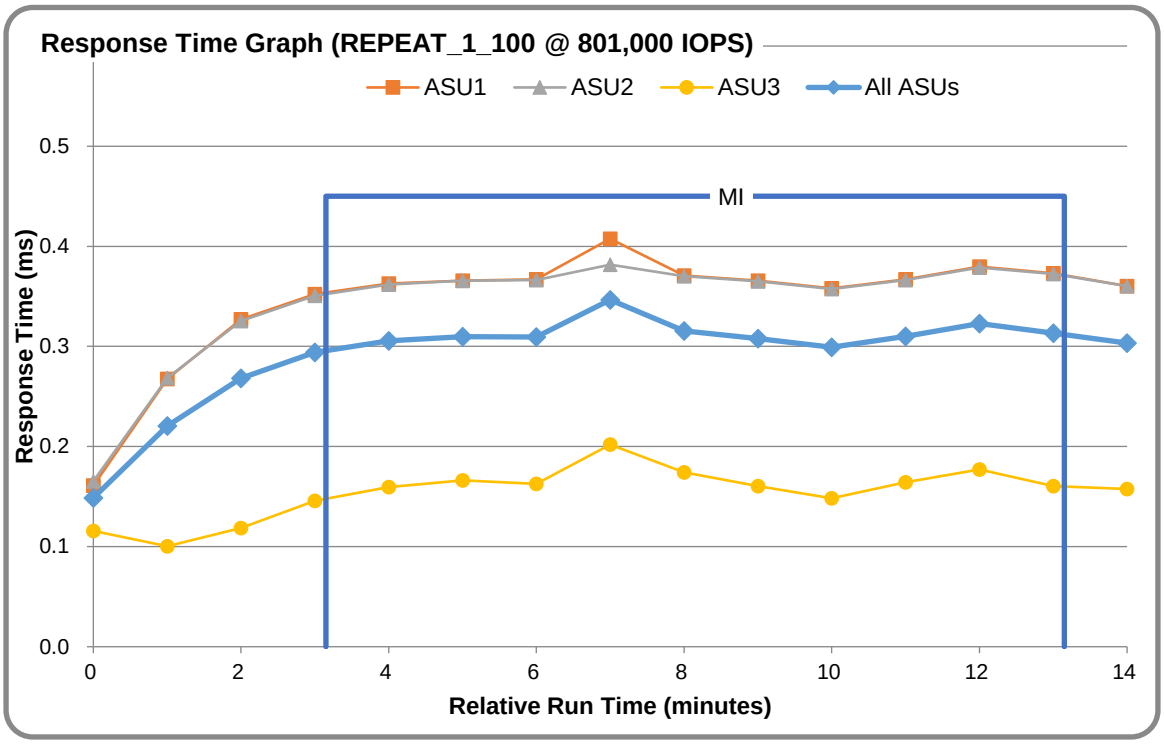

REPEAT_2_100 – Throughput Graph

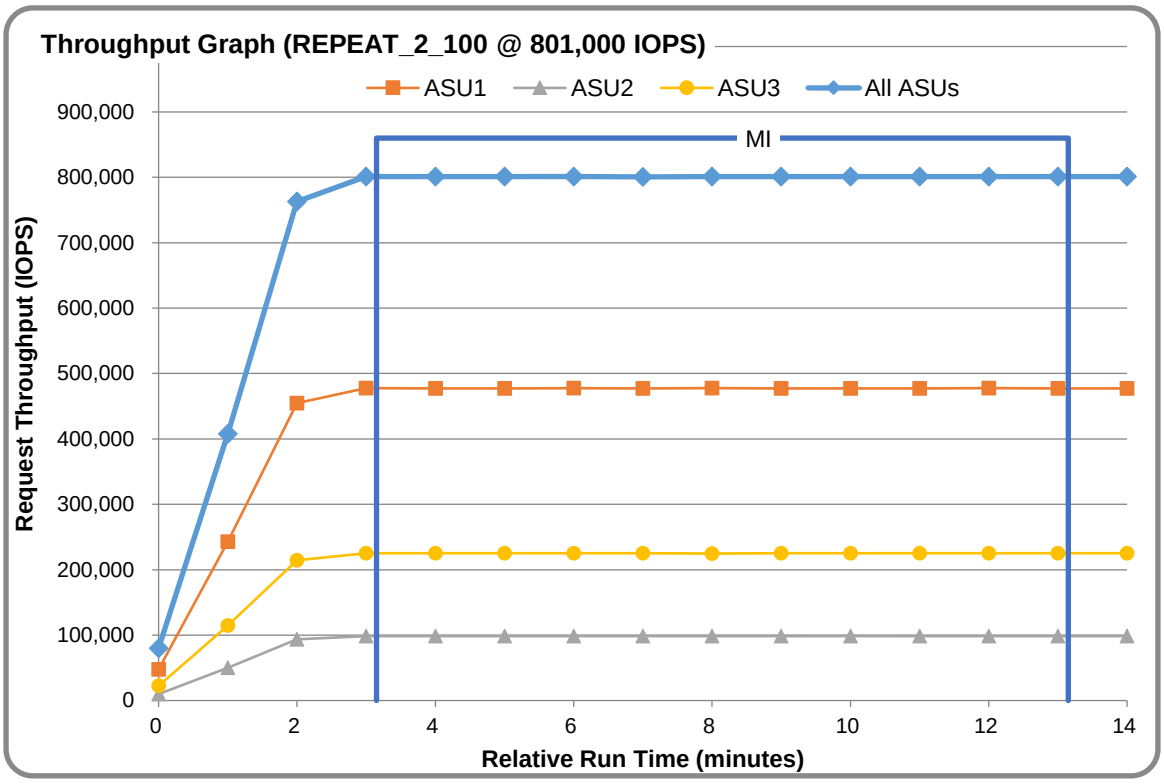

REPEAT_2_100 – Response Time Graph

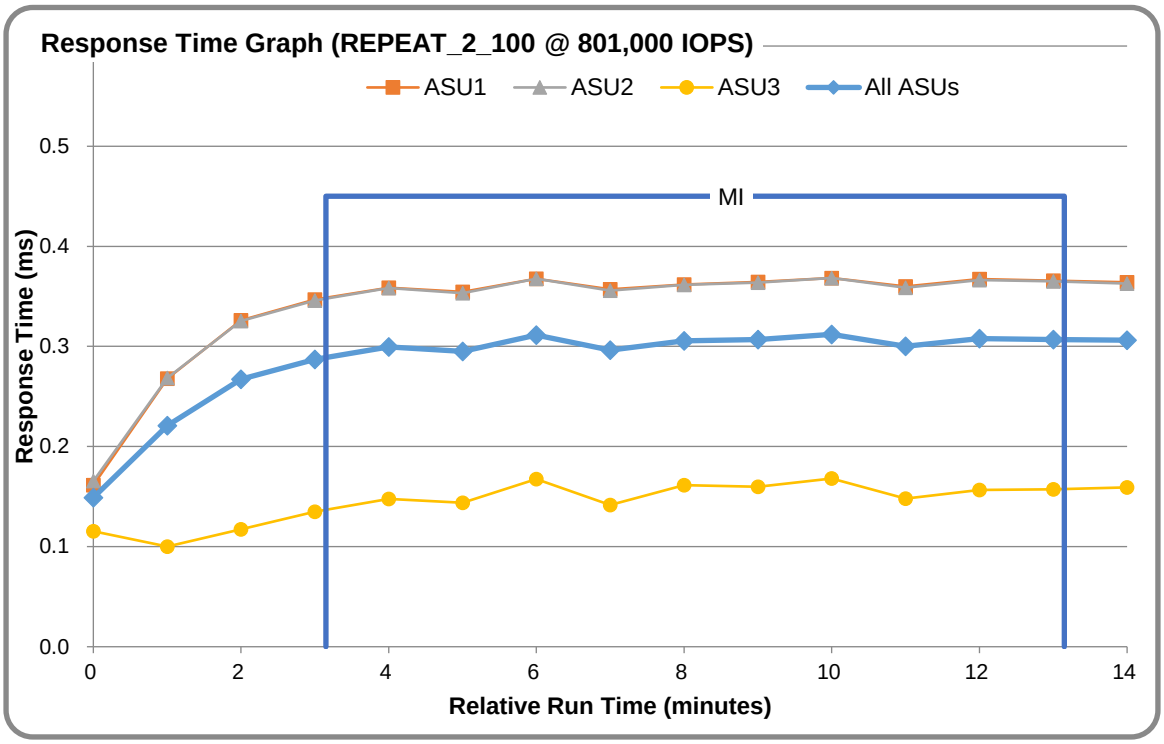

Repeatability Test – Intensity Multiplier

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 percent 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	0.2810	0.0700	0.2100	0.0180	0.0700	0.0350	0.2810
Variation	0.0006	0.0003	0.0007	0.0003	0.0010	0.0005	0.0005	0.0002
Difference	0.008%	0.011%	0.009%	0.024%	0.008%	0.033%	0.004%	0.003%

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	0.0700	0.0350	0.2810
Variation	0.0006	0.0002	0.0005	0.0003	0.0015	0.0004	0.0010	0.0002
Difference	0.006%	0.002%	0.002%	0.002%	0.036%	0.007%	0.008%	0.000%

Data Persistence Test

Data Persistence Test Result files

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	99,076,478
Total Number of Logical Blocks Verified	51,061,384
Total Number of Logical Blocks Overwritten	48,015,094
Total Number of Logical Blocks that Failed Verification	0
Time Duration for Writing Test Logical Blocks (sec.)	301
Size in Bytes of each Logical Block	8,192
Number of Failed I/O Requests During the Test	0

Committed Data Persistence Implementation

FusionStor splits each logical volume into many 1MB chunks and each chunk is replicated on two storage nodes, and the volume has one volume controller. If one of the replicated write or read failed, the volume controller would choose the next available device to recover the replication.

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 first VERIFY Test Phase	/SPC1_RESULTS
SPC1_VERIFY_1_Raw_Results.xlsx	Raw results for second VERIFY Test Phase	/SPC1_RESULTS
/C_Tuning	**Tuning parameters and options**	**root**
add_random.sh	Set add_random parameter	/C_Tuning
aio-max-nr.sh	Set maximum asynchronous I/O	/C_Tuning
max_sectors_kb.sh	Set maximum number of sectors	/C_Tuning
nomerges.sh	Set nomerges parameter	/C_Tuning
nr_request.sh	Increase value for nr_requests	/C_Tuning
queue_depth.sh	Change the queue depth	/C_Tuning
scheduler.sh	Change the I/O scheduler	/C_Tuning
/D_Creation	**Storage configuration creation**	**root**
create_volume_for_spc.sh	Create disk domains, storage pools, LUNs	/D_Creation
ln.sh	Collect PCI ids	/D_Creation
login.sh	Create the Logical Volumes	/D_Creation
/E_Inventory	**Configuration inventory**	**root**
collectStorageInfo.sh	Capture storage configuration inventory	/E_Inventory
collectStorageInfo_1.log	Storage inventory before INIT	/E_Inventory
collectStorageInfo_2.log	Storage inventory after Persistence test	/E_Inventory
/F_Generator	**Workload generator**	**root**
slave_asu.asu	Defining LUNs hosting the ASUs	/F_generator
3host.HST	Host configuration file	/F_generator
full_run.sh	Execute all test phases	/F_generator

APPENDIX B: THIRD PARTY QUOTATION

Source 2 - Beijing fuyangweixin Technology (Chinese Version)

From:富扬维鑫超微总代 董晓快
Tel: 010-5128 6188-8001
Beijing fuyangweixin Technology Co,.Ltd
北京富扬维鑫科技有限公司 邮政编码:100080
北京市海淀区志新东路8号五层
Fax: 010-8237 5099
Mobile: 18910138309
Email: dongxiaokuai@fywx.cc

超微服务器

	产品型号	主要配置	数量	单价	总价	总价
平台	2028U-TNRT+	2028U-TNRT+	3	2713.08	8139.24	
处理器	CM8064401439612	E5-2680V3	6	1554.69	9328.14	
内存	M393A2G40DB1-CRC	DDR4 RECC 16GB	12	152.42	1829.04	$ 21,216.93
技术服务	OS4HR3	三年7x24小时 4小时响应服务	3	640.17	1920.51	
备注：本报价为市场公开报价，自2017年7月23日起90天内有效						

组件		数量	描述
Motherboard / Chassis	CSE-219UAC-R1K02	1	2U Chassis
	MBD-X10DRU-i+	1	Super X10DRU-i+ Motherboard
LAN	AOC-2UR8N4-I2XT	1	Standard LP, 2x 10GbE RJ45, PCI-E x8, Intel X540
Air Shroud	MCP-310-82912-0N	1	4 heavy duty fans,1 Air Shroud
Backplane	BPN-SAS3-216A-N4		24-port 2U SAS3 12Gbps hybrid backplane, support up to 20x 2.5-inch SAS3/SATA3 HDD/SSD and 4x SAS3/SATA3/NVMe Storage Devices
Cable	CBL-SAST-0658	4	Internal Mini-SAS HD cable for PCIe SSD NVMe, 12Gb/s,
	CBL-SAST-0658	2	Internal RA side exit (right) miniSAS to miniSAS HD 55cm w/ SB, 30AWG,RoHS/REACH
	CBL-PWEX-0673	3	8 pin female to 2x big 4 pin female power, 65/70CM,
Parts	MCP-240-82909-0N	1	Ultra I/O bracket for SC829U, 219U (i2XT),RoHS/REACH
Drive Tray(s)	MCP-220-00127-0B	4	Black gen-3 2.5 NVMe drive tray, Orange tab with lock (for hotswap NVMe drive),HF,RoHS/REACH
Label	LBL-0108	1	CAUTION LABEL FOR REDUNDANT PWR SYSTEMS
Management	SFT-OOB-LIC	1	OOB Management Package (per node license)
Manual	MNL-1695-QRG	1	SYS-2028U-T(N)R(4/4T/T/TP)+, QRG,HF,RoHS/REACH
Riser Card	RSC-R2UW-4E8	1	RSC-R2UW-4E8-O-P
	RSC-R1UW-E8R	1	RSC-R1UW-E8R-O-P
Heatsink / Retention	SNK-P0058PSU	2	2U Passive CPU Heat Sink for X10 2U 24-DIMM Ultra Series Servers,RoHS/REACH
Power Supply	PWS-1K02A-1R	2	1U 1000W Redundant Power Supply 73.5mm width Titanium Level,RoHS/REACH
Mounting Rails	MCP-290-00053-0N	1	Inner/outer rail set,quick-release type, 26.5" to 36.4"
	MCP-290-00060-0N	1	Threaded rail adapter default for round hole rack

Source 2 - Beijing fuyangweixin Technology (English Version)

From:Supermicro General Distributor Dong Xiaokuai Tel: 010-5128 6188-8001

Beijing fuyangweixin Technology Co,.Ltd Fax: 010-8237 5099
ZIP Code:100080 Mobile: 18910138309
5th Floor No. 8, Zhixin East Rd, Haidian, Beijing, C Email: dongxiaokuai@fywx.cc

Supermicro Quotation

	Model	Configuration	Quantity	Unit Price	SUM	Total
Platform	2028U-TNRT+	2028U-TNRT+	3	2713.08	8139.24	$ 21,216.93
CPU	CM8064401439612	E5-2680V3	6	1554.69	9328.14	
Memory	M393A2G40DB1-CRC	DDR4 RECC 16GB	12	152.42	1829.04	
Support	OS4HR3	3 Years, 7x24, 4 Hr Response	3	640.17	1920.51	

NOTE : This is public quotation to China market, valid within 90 days from JUL 23rd 2017

Component		Quantity	Description
Motherboard / Chassis	CSE-219UAC-R1K02	1	2U Chassis
	MBD-X10DRU-i+	1	Super X10DRU-i+ Motherboard
LAN	AOC-2UR8N4-I2XT	1	Standard LP, 2x 10GbE RJ45, PCI-E x8, Intel X540
Air Shroud	MCP-310-82912-0N	1	4 heavy duty fans,1 Air Shroud
Backplane	BPN-SAS3-216A-N4		24-port 2U SAS3 12Gbps hybrid backplane, support up to 20x 2.5-inch SAS3/SATA3 HDD/SSD and 4x SAS3/SATA3/NVMe Storage Devices
Cable	CBL-SAST-0658	4	Internal Mini-SAS HD cable for PCIe SSD NVMe, 12Gb/s, 60cm,30AWG,RoHS/REACH
	CBL-SAST-0658	2	Internal RA side exit (right) miniSAS to miniSAS HD 55cm w/ SB, 30AWG,RoHS/REACH
	CBL-PWEX-0673	3	8 pin female to 2x big 4 pin female power, 65/70CM, 18AWG,RoHS/REACH
Parts	MCP-240-82909-0N	1	Ultra I/O bracket for SC829U, 219U (i2XT),RoHS/REACH
Drive Tray(s)	MCP-220-00127-0B	4	Black gen-3 2.5 NVMe drive tray, Orange tab with lock (for hotswap NVMe drive),HF,RoHS/REACH
Label	LBL-0108	1	CAUTION LABEL FOR REDUNDANT PWR SYSTEMS
Management	SFT-OOB-LIC	1	OOB Management Package (per node license)
Manual	MNL-1695-QRG	1	SYS-2028U-T(N)R(4/4T/T/TP)+, QRG,HF,RoHS/REACH
Riser Card	RSC-R2UW-4E8	1	RSC-R2UW-4E8-O-P
	RSC-R1UW-E8R	1	RSC-R1UW-E8R-O-P
Heatsink / Retention	SNK-P0058PSU	2	2U Passive CPU Heat Sink for X10 2U 24-DIMM Ultra Series Servers,RoHS/REACH
Power Supply	PWS-1K02A-1R	2	1U 1000W Redundant Power Supply 73.5mm width Titanium Level,RoHS/REACH
Mounting Rails	MCP-290-00053-0N	1	Inner/outer rail set,quick-release type, 26.5" to 36.4"
	MCP-290-00060-0N	1	Threaded rail adapter default for round hole rack

Source 3 - Digital China Macao Service (Chinese Versions)

报价单

神州数码 Digital China

联系人:

电话:

序列	型号	料号	单价	数量	总价
1	S3710系列200G	SSDSC2BA200G401	$ 205.77	3	$ 617.31
2	P3700系列2TB	SSDPEDMD020T401	$ 2,865.50	6	$ 17,193.00
3					
4					
5					
合计		$			17,810.31

TERMS :

* Manufacturer : Intel
* Lead Time :
* Payment : T/T in advance
* Validity : 90Days

州数码澳门离岸商业服务有限公

7/25/2017

Source 3 - Digital China Macao Service (English Versions)

QUOTATION

神州数码 Digital China

Contact:

Phone:

No.	Model	P/N	Unit Price	Qty	SUM
1	S3710系列200G	SSDSC2BA200G401	$ 205.77	3	$ 617.31
2	P3700系列2TB	SSDPEDMD020T401	$ 2,865.50	6	$ 17,193.00
3					
4					
5					
Total		$			17,810.31

TERMS :

* Manufacturer : Intel
* Lead Time :
* Payment : T/T in advance
* Validity : 90Days

Digital China Macao Service Co., Ltd

7/25/2017

Source 4 - Shenzhen Baotongzhiyuan Technology (Chinese Version)

深圳市宝通志远科技有限公司报价清单					
产品名称	产品型号	产品配置描述	单价	数量	小计
SSD P3700 2TB	SSDPE2MD020T401	Intel DC P3700 系列NVME U.2接口系列固态硬盘 容量2TB	$ 2,895.98	2	$ 5,791.96
SSD P3700 1.6TB	SSDPE2MD016T401	Intel DC P3700 系列NVME U.2接口系列固态硬盘 容量1.6TB	$ 2,277.16	10	$ 22,771.60
合 计					$ 28,563.56
注：本报价90天内有效(自2017年7月20号起)					

Source 4 - Shenzhen Baotongzhiyuan Technology (English Version)

Shenzhen Baotongzhiyuan Technology Co,.Ltd Quotation Page					
Product	Model	Description	Unit Price	Qty	SUM
SSD P3700 2TB	SSDPE2MD020T401	Intel DC P3700 Serial NVME U.2 Interface SSD, Capacity 2TB	$ 2,895.98	2	$ 5,791.96
SSD P3700 1.6TB	SSDPE2MD016T401	Intel DC P3700 Serial NVME U.2 Interface SSD, Capacity1.6TB	$ 2,277.16	10	$ 22,771.60
Total					$ 28,563.56
Note: This Quotation is Valid within 90 Days from JUL 20th, 2017 to Market					

Source 5 - Sainuo Xinzhi Software Technology (Chinese Version)

有效日期　　报价日后90日内(2017/7/25)

产品名称	产品描述	数量	单价	合计	到货周期
MSX1012B-2BFS	SwitchX®-2 based 40GbE, 1U Open Ethernet Switch with MLNX-OS, 12 QSFP+ ports, 2 Power Supplies (AC), short depth, PPC460, P2C airflow, Rail Kit must be purchased separately, RoHS6	2	$ 6,096.80	$ 12,193.60	目前现货，订货4-6周
MCX353A-FCBT	ConnectX®-3 VPI adapter card, single-port QSFP, FDR IB (56Gb/s) and 40/56GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6	12	$ 757.23	$ 9,086.76	目前现货，订货4-6周
MCX456A-ECAT	Mellanox CX456A- ConnectX-4 QSFP28 Dual-port	3	$ 1,656.51	$ 4,969.53	目前现货，订货4-6周
MC2207128-002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m	20	$ 82.31	$ 1,646.20	目前现货，订货4-6周
SUP-SX1012-3S-4H	Mellanox SUP-SX1012-3S-4H Technical Support and Warranty - Silver 3 Year with 4 Hours On-Site Support for SX1012 Series Switch	2	$ 1,463.24	$ 2,926.48	目前现货，订货4-6周
合计				$ 25,853.04	

报价单位：赛诺信致软件技术（北京）有限公司
联 系 人：王啸
联系电话：13811896638.
传　　真：010-68949110
E-mail:wangxiao@infohpc.com
付款方式：收到全款后发货

Source 5 - Sainuo Xinzhi Software Technology (English Version)

Valid Date 90 Days from JUL 25th, 2017

Product	Description	Qty	Unit Price	Sum	Delivery Plan
MSX1012B-2BFS	SwitchX®-2 based 40GbE, 1U Open Ethernet Switch with MLNX-OS, 12 QSFP+ ports, 2 Power Supplies (AC), short depth, PPC460, P2C airflow, Rail Kit must be purchased separately, RoHS6	2	$ 6,096.80	$ 12,193.60	Available, booking 4-6 week
MCX353A-FCBT	ConnectX®-3 VPI adapter card, single-port QSFP, FDR IB (56Gb/s) and 40/56GbE, PCIe3.0 x8 8GT/s, tall bracket, RoHS R6	12	$ 757.23	$ 9,086.76	Available, booking 4-6 week
MCX456A-ECAT	Mellanox CX456A- ConnectX-4 QSFP28 Dual-port	3	$ 1,656.51	$ 4,969.53	Available, booking 4-6 week
MC2207128-002	Mellanox® passive copper cable, VPI, up to 56Gb/s, QSFP, 3m	20	$ 82.31	$ 1,646.20	Available, booking 4-6 week
SUP-SX1012-3S-4H	Mellanox SUP-SX1012-3S-4H Technical Support and Warranty - Silver 3 Year with 4 Hours On-Site Support for SX1012 Series Switch	2	$ 1,463.24	$ 2,926.48	Available, booking 4-6 week
Total				$ 25,853.04	

Company: Sainuo Xinzhi software Technology(Beijing) Co,.Ltd
Contact: Wang Xiao
Phone: 13811896638.
FAX: 010-68949110
E-mail:wangxiao@infohpc.com
Payment: Deliver after receiving full payment

APPENDIX C: TUNING PARAMETERS AND OPTIONS

The following scripts, listed below, were used to set tuning parameters and options on each Host System:

- ***aio-max-nr.sh*** to set the maximum number of AIO operations to 1048576.
- ***nr_request.sh*** to set *nr_requests* to 1024.
- ***scheduler.sh*** to set the I/O scheduler to *noop*.
- ***queue_depth.sh*** to set *queue_depth* to 54.
- ***max_sectors_kb.sh*** to set *max_sectors_kb* to 128.
- ***nomerges.sh*** to set *nomerges* to 2.
- ***add_random.sh*** to set *add_random* to 1.

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

aio-max-nr.sh

```
echo 1048576 > /proc/sys/fs/aio-max-nr
```

nr_request.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 1024 > /sys/block/$dev/queue/nr_requests
done
```

scheduler.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 'noop' > /sys/block/$dev/queue/scheduler
done
```

queue_depth.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 54 > /sys/block/$dev/device/queue_depth
done
```

max_sectors_kb.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 128 > /sys/block/$dev/queue/max_sectors_kb
done
```

nomerges.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 2 > /sys/block/$dev/queue/nomerges
done
```

add_random.sh

```
devs=`ls -l /dev/disk/by-path/ | grep 'iscsi' | grep 'asu' | awk -F '/' '{print
    $NF}'`
for dev in $devs;
do
    echo 1 > /sys/block/$dev/queue/add_random
done
```

APPENDIX D: STORAGE CONFIGURATION CREATION

Step 1 - Create cluster, directory and logical volumes

The following actions are performed using the FusionStor SF6000 CLI from the master Host System:

1. The script ***ln.sh*** is executed to Collect PCI ids from the storage devices.
2. The script ***create_volume_for_spc.sh*** is executed to perform the following steps:

   Create a Control Cluster;

   Create a directory (pool);

   Create logical volumes; and

   Allocate logical volumes.

Step 2 - Mount the logical volumes on the Host System

The following actions are performed on the Host System:

1. The script ***login.sh*** is executed on the master Host System to perform the following steps:

   Mount the logical volumes;

   Create the software RAIDs; and

   Mount the logical volumes;

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

ln.sh

```
#!/bin/bash

/home/spdk-16.07/scripts/setup.sh
nvmeIds=`lspci | grep 'Non-Volatile memory' | awk '{print $1}'`
diskId=0

rm -rf /opt/fusionstack/data/disk/disk
mkdir -p /opt/fusionstack/data/disk/disk

for id in $nvmeIds;
do
{
        nvmeFile=/etc/pci:0000:$id
        touch $nvmeFile
        ln -s $nvmeFile /opt/fusionstack/data/disk/disk/$diskId.disk
        let diskId=$(($diskId + 1))
}
done

ls -l /opt/fusionstack/data/disk/disk
```

create_volume_for_spc.sh

```
#!/bin/bash

set -x

Fusionstack="spc-57 spc-58 spc-59"
node1=spc-57
node2=spc-58
node3=spc-59

echo '********************** Setp[1]: Prepare ***********************'
lich stop
for node in $Fusionstack;
do
        echo $node
        ssh $node 'iptables -F'
        ssh $node 'mkdir -p /home/log'
        ssh $node 'mv -f /opt/fusionstack/log /home/log/log_`date +%s`'
        ssh $node 'rm -rf /opt/fusionstack/data/*'
        ssh $node 'rm -rf /opt/fusionstack/etc/cluster.conf'
        ssh $node 'rm -rf /dev/shm/lich4/*'
        ssh $node 'rm -rf /opt/fusionstack/core/*'
        ssh $node 'rm -rf /dev/huagepages/*'

        ssh $node "bash /root/ln.sh"
done
sleep 10


echo '********************** Setp[2]: Create Cluster ***********************'
/opt/fusionstack/lich/bin/lich create $Fusionstack
sleep 20
lich stat


echo '********************** Setp[3]: Create Volume ***********************'
while [ 1 ]
do
        lichbd mkpool pool -p iscsi
        echo $?
        if [ $? -eq 0 ]; then
                break;
        elif [ $? -eq 17 ]; then
                break;
        fi

        sleep 5
done

for i in {0..5};  do lichbd create pool/volume$i -s 64G  -p iscsi;done
for i in {6..11}; do lichbd create pool/volume$i -s 288G -p iscsi;done
for i in {12..17};do lichbd create pool/volume$i -s 288G -p iscsi;done


for i in {18..23};do lichbd create pool/volume$i -s 288G -p iscsi;done
for i in {24..29};do lichbd create pool/volume$i -s 64G  -p iscsi;done
for i in {30..35};do lichbd create pool/volume$i -s 288G -p iscsi;done


for i in {36..41};do lichbd create pool/volume$i -s 288G -p iscsi;done
for i in {42..47};do lichbd create pool/volume$i -s 288G -p iscsi;done
```

```
for i in {48..53};do lichbd create pool/volume$i -s 64G  -p iscsi;done

sleep 20
echo 'create volume success'


echo '********************** Step[4]: Allocate Volume **********************'
for vol in $(lichfs --list /iscsi/pool |awk '{print $NF}');
do
       lich.inspect --allocate /iscsi/pool/$vol
done

while [ 1 ]
do
        incomplete=0
       for node in $Fusionstack
       do
                ssh $node1 'ps aux | grep "lich.inspect --allocate" | grep -v
   "grep"'
                if [ $? -eq 0 ]; then
                        let incomplete=$((incomplete+1))
                fi
       done

        if [ $incomplete -eq 0 ]; then
                break;
        fi
done
echo "Allocate Complete."
```

login.sh

```
#!/bin/bash

set -x

echo "Iser Initiator : ****************** `hostname` ******************"

server1=192.168.1.57
scp -r -o StrictHostKeyChecking=no root@$server1:/root/cpu.text /root/

ls /root/cpu.text
if [ $? -ne 0 ];then
   echo '/root/cpu.text not exist'
   exit
fi

devices=`find /dev  -maxdepth 1 -type b -name "md*"`
for device in $devices;
do
        mdadm -S $device
done

ib1=192.168.17.57
ib0=192.168.18.57
date && iscsiadm -m discovery -t st -p $ib0 -I iser
if [ $? -ne 0 ]; then
   echo "date && iscsiadm -m discovery -t st -p $ib0 -I iser"
   exit 1;
fi
```

```
for i in $(cat /root/cpu.text | grep 'spc-57' | grep "numa node \[0\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib0 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib0 -I iser -l"
      exit 1;
   fi
done

sleep 2
date && iscsiadm -m discovery -t st -p $ib1 -I iser
if [ $? -ne 0 ]; then
   echo "date && iscsiadm -m discovery -t st -p $ib1 -I iser"
   exit 1;
fi
for i in $(cat /root/cpu.text | grep 'spc-57' | grep "numa node \[1\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib1 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib1 -I iser -l"
      exit 1;
   fi
done


ib1=192.168.17.58
ib0=192.168.18.58
date && iscsiadm -m discovery -t st -p $ib0 -I iser
if [ $? -ne 0 ]; then
   echo "date && iscsiadm -m discovery -t st -p $ib0 -I iser"
   exit 1;
fi
for i in $(cat /root/cpu.text | grep 'spc-58' | grep "numa node \[0\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib0 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib1 -I iser -l"
      exit 1;
   fi
done

sleep 2
date && iscsiadm -m discovery -t st -p $ib1 -I iser
for i in $(cat /root/cpu.text | grep 'spc-58' | grep "numa node \[1\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib1 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib1 -I iser -l"
      exit 1;
   fi
```

```
done


ib1=192.168.17.59
ib0=192.168.18.59
date && iscsiadm -m discovery -t st -p $ib0 -I iser
if [ $? -ne 0 ]; then
   echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p
   $ib1 -I iser -l"
   exit 1;
fi
for i in $(cat /root/cpu.text | grep 'spc-59' | grep "numa node \[0\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib0 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib1 -I iser -l"
      exit 1;
   fi
done

sleep 2
date && iscsiadm -m discovery -t st -p $ib1 -I iser
if [ $? -ne 0 ]; then
   echo "date && iscsiadm -m discovery -t st -p $ib1 -I iser"
   exit 1;
fi
for i in $(cat /root/cpu.text | grep 'spc-59' | grep "numa node \[1\]" | awk
   '{print $2}');
do
   date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -p $ib1 -I
   iser -l;
   if [ $? -ne 0 ]; then
      echo "date && iscsiadm -m node -T iqn.2001-04-1234.com.meidisen:pool.$i -
   p $ib1 -I iser -l"
   fi
done

sleep 4

devs=`ls -l /dev/disk/by-path/ | grep iscsi | grep asu | awk -F '/' '{print $NF}'`
for dev in $devs;
do
   echo 54 > /sys/block/$dev/device/queue_depth
   echo 1024 > /sys/block/$dev/queue/nr_requests
   echo 'noop' > /sys/block/$dev/queue/scheduler
   echo 128 > /sys/block/$dev/queue/max_sectors_kb
   echo 1 > /sys/block/$dev/queue/add_random
   echo 2 > /sys/block/$dev/queue/nomerges
done

devices=`find /dev  -maxdepth 1 -type b -name "md*"`
for device in $devices;
do
        mdadm -S $device
done

echo 'y' | mdadm --create /dev/md0 -c 32 -l 0 -n $(ls /dev/disk/by-path/ | grep
   'asu1'  |  wc  -l)  $(ls  /dev/disk/by-path/  |  grep  'asu1'  |  sed  -n
   's/^/\/dev\/disk\/by-path\//p')
```

```
echo 'y' | mdadm --create /dev/md1 -c 32 -l 0 -n $(ls /dev/disk/by-path/ | grep
   'asu2'  |  wc  -l)  $(ls  /dev/disk/by-path/  |  grep  'asu2'  |  sed  -n
   's/^/\/dev\/disk\/by-path\//p')
echo 'y' | mdadm --create /dev/md2 -c 32 -l 0 -n $(ls /dev/disk/by-path/ | grep
   'asu3'  |  wc  -l)  $(ls  /dev/disk/by-path/  |  grep  'asu3'  |  sed  -n
   's/^/\/dev\/disk\/by-path\//p')

devices=`find /dev  -maxdepth 1 -type b -name "md*"`
for device in $devices;
do
        mdadm -D $device
done
```

APPENDIX E: CONFIGURATION INVENTORY

An inventory of the Tested Storage Configuration was collected before and after the test execution. The test execution script invokes ***collectStorageInfo.sh*** to collect the inventory profile of the storage configuration. The following log files were generated and are included in the Supporting Files (see Appendix A):

- ***collectStorageInfo_1.log*** List of configured storage and logical volumes before test execution.
- ***collectStorageInfo_2.log*** List of configured storage and logical volumes after TSC restart.

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

collectStorageInfo.sh

```
#!/bin/bash

set -x

rm -rf /root/$1

stor=192.168.1.57
#ssh $stor 'lich stop'
ssh $stor 'lich start' | tee -a /root/$1

sleep 20

while [ 1 ]
do
    count=0
    running=`ssh $stor 'ps aux | grep lichd | grep fusionstack'`
    status=`ssh $stor 'lich list -v' | awk -F ':' '{print $6}' | awk -F ',' '{print
    $1}'`
    echo $status | tee -a /root/$1
    for st in $status;
    do
        if [[ $st == 'admin' ]]; then
               let count=$((count+1))
        elif [[ $st == 'meta' ]]; then
               let count=$((count+1))
        else
               echo "`date` -- node not ready, status: $st" | tee -a /root/$1
        fi
    done

    sleep 5

    echo $count
    if [ $count -eq 3 ]; then
        echo "`date` -- Fusionstack is ready." | tee -a /root/$1
        break;
    fi
done

sleep 30
```

```
while [ 1 ]
do
   ssh -o StrictHostKeyChecking=no root@$stor 'bash /root/show_vol_cpu_info.sh >
   /root/cpu.text'
   if [ $? -eq 0 ]; then
      break
   fi

   echo   "retry:   ssh   -o   StrictHostKeyChecking=no   root@$stor   'bash
   /root/show_vol_cpu_info.sh > /root/cpu.text'"
   sleep 5
done

iserInitiator=`cat /root/SPC1_v340_20170515/Host.HST | grep 'HOST=' | awk -F '='
   '{print $2}'`
for node in $iserInitiator;
do
   ssh $node 'iptables -F' | tee -a /root/$1
   ssh $node 'bash /root/aio-max-nr.sh' | tee -a /root/$1
        ssh $node 'bash /root/login.sh' | tee -a /root/$1
   sleep 5
done
```

APPENDIX F: WORKLOAD GENERATOR

The host parameters for the SPC-1 workload generator were defined using the script ***3host.HST***.

The ASUs accessed by the SPC-1 workload generator are defined using the script ***slave_asu.asu***.

The initial test phases of the benchmark were invoked using the script ***full_run.sh***, which paused to allow for the Persistence Test shutdown. Once the TSC had been restarted, the PERSIST_2 test phase was invoked by resuming the execution of the script.

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

3host.HST

```
-- Common Information
LOGIN=root
PASSWORD=password
CONFIG=/root/SPC1_v340_20170515
EXEC=spc1
PORT=1001
WINDOWS=N
WEIGHT=1
STORAGE=slave_asu.asu
-- Host Entries
HOST=192.168.1.120
HOST=192.168.1.121
HOST=192.168.1.122
```

slave_asu.asu

```
ASU=1
OFFSET=0
SIZE=0
DEVICE=/dev/md0
ASU=2
OFFSET=0
SIZE=0
DEVICE=/dev/md1
ASU=3
OFFSET=0
SIZE=0
DEVICE=/dev/md2
```

full_run.sh

```
#!/bin/bash

if [[ $1 == "phase1" ]];then
mv                    -f                    /root/output/spc1_full_run_800k_iops
   /home/spc_full_test/spc1_full_run_800k_iops_`date +%s`
sh collectStorageInfo.sh collectStorageInfo_1.log
```

```
/root/SPC1_v340_20170515/spc1  -run  SPC1_INIT          -iops  6000     -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops -master 3Host.HST
/root/SPC1_v340_20170515/spc1  -run  SPC1_VERIFY       -iops  100        -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops
/root/SPC1_v340_20170515/spc1  -run  SPC1_METRICS       -iops  801000  -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops -master 3Host.HST

else
/root/SPC1_v340_20170515/spc1  -run  SPC1_VERIFY       -iops  100        -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops
/root/SPC1_v340_20170515/spc1  -run  SPC1_PERSIST_1  -iops  200250    -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops -master 3Host.HST

echo "Power cycle TSC, then Enter to continue"
read
sh collectSysShutdownLog.sh collectSysShutdown.log
sh collectStorageInfo.sh collectStorageInfo_2.log

/root/SPC1_v340_20170515/spc1  -run  SPC1_PERSIST_2  -iops  200250    -storage
   slave_asu.asu -output /root/output/spc1_full_run_800k_iops -master 3Host.HST
fi
```