SPC BENCHMARK 1™

FULL DISCLOSURE REPORT

CHINA ELECTRONICS CLOUD TECHNOLOGY CO., LTD.

CEASTOR 18116E

SPC-1™ V3.10.0

SUBMISSION IDENTIFIER: A32026

SUBMITTED FOR REVIEW: APRIL 25, 2023

First Edition – April 2023

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 China Electronics Cloud Technology Co., Ltd. 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. China Electronics Cloud Technology Co., Ltd. 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 China Electronics Cloud Technology Co., Ltd. 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.

CeaStor and the CECTC logo are trademarks or registered trademarks of China Electronics Cloud Technology Co., Ltd. in 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.spcresults.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

Yao Mi
N3013,3F,N R&D building, A.I. Technology Park,
Economic and Technological Development Zone,
Wuhan,Hubei,China

April 23, 2023

I verified the SPC Benchmark 1™ (SPC-1 ™ v3.10.0) test execution and performance results of the following Tested Storage Product:

CeaStor 18116E

The results were:

SPC-1 IOPS™	**10,000,690**
SPC-1 Price-Performance	**$165.82/SPC-1 KIOPS™**
SPC-1 Total System Price	1,658,233.00
SPC-1 IOPS Response Time	0.548 ms
SPC-1 Overall Response Time	0.352 ms
SPC-1 ASU Capacity	127,131 GB
SPC-1 ASU Price	$13.05/GB

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 v3.0.2. 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 China Electronics Cloud Technology Co., Ltd., stating the accuracy and completeness of the documentation and testing data provided in support of the audit of this result.

Page **1** of **2**

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

A Full Disclosure Report for this result was prepared by InfoSizing, reviewed and approved by China Electronics Cloud Technology Co., Ltd., and can be found at www.spcresults.org under the Submission Identifier A32026.

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

- The physical capacity of the data repository (349,440 GB).
- The total capacity of the Application Storage Unit (127,131 GB).
- 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 each persistence test.
- The compliance of the submitted pricing model.
- 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 in accordance with the SPC Policies:

None.

Respectfully Yours,

Doug Johnson, Certified SPC Auditor

63 Lourdes Dr. | Leominster, MA 01453 | 978-343-6562 | www.sizing.com

LETTER OF GOOD FAITH

April 22, 2023

To; Doug Johnson, SPC Auditor
PerfLabs, Inc. DBA InfoSizing
63 Lourdes Drive
Leominster, MA 01453-6709
USA

Subject: SPC-1 Letter of Good Faith for the CeaStor18116E

China Electronics Cloud Technology Co., Ltd is the SPC-1 test sponsor for the above listed product. To the best of our knowledge and belief, the required SPC-1 result and materials we have submitted for that product are complete, accurate, and in full compliance with version 3.10 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 result even if the items are not explicitly required to be disclosed by the SPC-1 benchmark specification.

Sincerely,

Signed: 

Date: 2024-4-22

Yao Mi
Storage Product Department
China Electronics Cloud Technology Co., Ltd

SPC Benchmark 1™

Executive Summary

CeaStor 18116E

SPC-1 IOPS™	**10,000,690**
SPC-1 IOPS Response Time	**0.548 ms**
SPC-1 Overall Response Time	0.352 ms

SPC-1 Price Performance	**$165.82/SPC-1 KIOPS™**
SPC-1 Total System Price	**$1,658,233.00**
SPC-1 Overall Discount	61.52%
Currency / Target Country	USD / China
Availability Date	May 31, 2023

Extensions

SPC-1 Data Reduction	NA
SPC-1 Encryption	NA
SPC-1 NDU	NA
SPC-1 Synchronous Replication	NA
SPC-1 Snapshot	NA

Storage Metrics

SPC-1 Data Protection Level	Protected 2
SPC-1 Physical Storage Capacity	349,440 GB
SPC-1 ASU Capacity	127,131 GB
SPC-1 ASU Price	$13.05/GB

Priced Storage Configuration Summary

40	Mellanox MCX623106AN-CDAT
1	CeaStor 18116E
30	Storage Nodes
15,360	GB Total Cache
60	Total Front-End Ports
420	Total Storage Devices (240x 256 GB Optane, 180x 1.6 TB NVMe SSD)
4	100Gb Huawei CE8851-32CQ8DQ-P Switches
68	Total RUs

RAMPD Average Response Time (ms) vs. IOPS

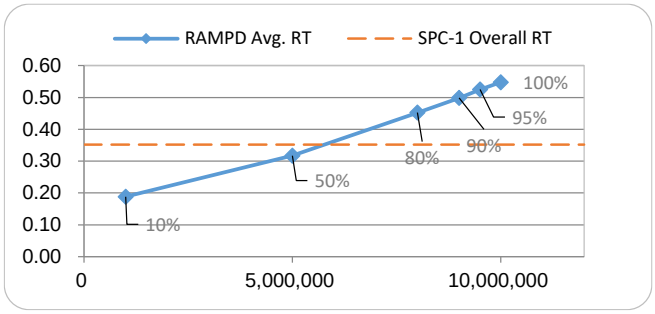

SUSTAIN Response Time (ms)

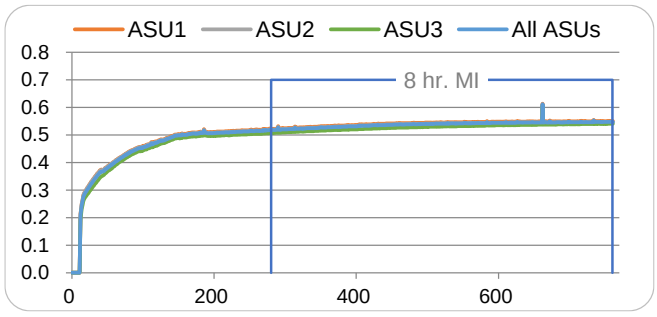

RAMPD_100 Response Time (ms)

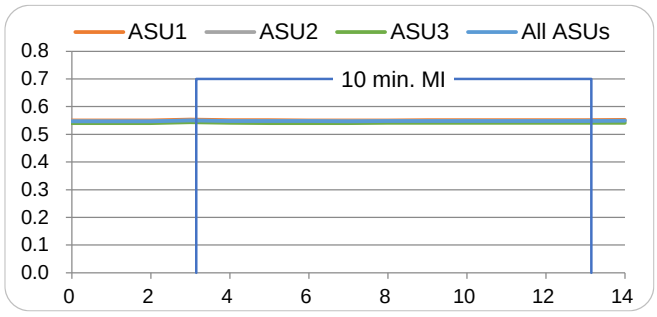

SPC Benchmark 1™ Specification Revision	v3.10.0
SPC Benchmark 1™ Workload Generator Revision	v3.0.2

Submitted for Review	April 25, 2023
Submission Details	www.storageperformance.org/r/A32026

PRICING DETAILS

Part No.	Description	Source	Qty	Unit Price	Ext. Price	Disc.	Disc. Price
Hardware & Software							
100409617284018	CeaStor 18116E Storage Node (NVMe All-Flash， 2U Rackmount) Case, Support 16 * 2.5-inch NVMe SSDs (hot swappable).	1	30	10,532.00	315,960.00	65%	110,586.00
100400010825874	Intel Xeon Platinum 8358P Processor Unit (2.60GHz, 32Core)	1	60	9,887.00	593,220.00	65%	207,627.00
100400010582171	32GB DDR4 2933MHz RDIMM.	1	480	945.00	453,600.00	65%	158,760.00
100400010194837	Intel Optane 200 Series PMEM, 256GB.	1	240	3,050.00	732,000.00	65%	256,200.00
100400010223011	SATA SSD, 480GB.	1	60	892.00	53,520.00	65%	18,732.00
100410026688821	NVMe SSD, INTEL DC P4610 1.6TB.	1	180	4,104.00	738,720.00	65%	258,552.00
100400010784751	Network Card, 100GbE Dual-port.	1	100	4,221.00	422,100.00	65%	147,735.00
100400010594085	100Gb CE8851-32CQ8DQ-P Switch.	1	4	75,401.00	301,604.00	50%	150,802.00
100400010854602	Mellanox Optical Module SFP 100G MM	1	200	165.00	33,000.00	50%	16,500.00
100400010998567	5M LC-LC MPO Optical Fiber Cable	1	100	60.00	6,000.00	50%	3,000.00
100407949637484	CeaStor Software, support Block/File/Object Functions, Licensed by Node.	1	30	18,000.00	540,000.00	50%	270,000.00
100400003390207	Storage RACK 42U AC Cabinet	1	2	1,200.00	2,400.00	50%	1,200.00
						Hardware & Software Subtotal	**1,599,694.00**
Support & Maintenance							
100307950095961	CeaStor 18000 Premier 24x7x4H Onsite Service, 36Month(s).	1	30	2,583.00	77,490.00	50%	38,745.00
100109389160418	Switche Premier 24x7x4H Onsite Service, 36Month(s).	1	4	9,897.00	39,588.00	50%	19,794.00
						Support & Maintenance Subtotal	**58,539.00**
SPC-1 Total System Price							**1,658,233.00**
SPC-1 IOPS™							10,000,690
SPC-1 Price-Performance™ ($/SPC-1 KIOPS™)							**165.82**
SPC-1 ASU Capacity (GB)							127,131
SPC-1 ASU Price ($/GB)							**13.05**

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

Warranty: The warranty provides 24x7x4H arrival of service within designated city and distance. The service includes 24x7 contact to the CeaStor call center with 4-hour on-site hardware replacement or troubleshooting, and online software support with access to all new software updates or troubleshooting.

Differences Between Tested and Priced Storage Configurations

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

PUBLICATION DETAILS

This section provides contact information for the test sponsor and auditor, a revision history of this document, and a description of any exceptions or waivers associated with this publication.

Contact Information

Role	Name	Details
Test Sponsor Primary Contact	China Electronics Cloud Technology Co., Ltd. Yao Mi	https://www.cecloud.com/ miyao@cestc.cn
SPC Auditor	InfoSizing Doug Johnson	www.sizing.com doug@sizing.com

Revision Information

Date	FDR Revision	Details
April 25, 2023	First Edition	Initial Publication

Anomalies, Exceptions, Waivers

There were no anomalies, exceptions or waivers associated with the audit of the CeaStor 18116E.

CONFIGURATION INFORMATION

Tested Storage Product Description

CeaStor 18116E is a next-generation storage product designed and developed by China Electronic Cloud for future storage architecture. It fully adopts a whole range of new storage hardware and new technologies, delivering industry-leading ultimate data access performance. This product features hyperscale, high availability, high reliability, elasticity, intelligence, simplicity, and efficiency. It can be widely used in cloud computing and cloud native environments to support key business applications and satisfy the strict requirements for data scale and access performance in emerging technologies and applications such as online transaction, supercomputing, AI, and autonomous driving.

Host System and Tested Storage Configuration Components

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

Host Systems
10x Huawei 2288H V6 2x Intel® Xeon® Platinum 8358P CPU @ 2.60 GHz 512 GB Main Memory CentOS Linus Release 8.4 10x Inspur NF5280M6 2x Intel® Xeon® Platinum 8358P CPU @ 2.60 GHz 512 GB Main Memory CentOS Linus Release 8.4
Tested Storage Configuration
40x Mellanox MCX623106AN-CDAT 100 Gb 2-port
30x CeaStor18116E storage nodes, each with: **2x Intel® Xeon® Platinum 8358P CPU @ 2.60 GHz** **512 GB cache (15,360 GB total)** **2x 100 Gbps Front End Ports** **240x Optane 256 GB, 180x NVMe SSD Storage Devices**
4x 100Gb Huawei CE8851-32CQ8DQ-P Switches

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

Configuration Diagrams

BC/TSC Configuration Diagram

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

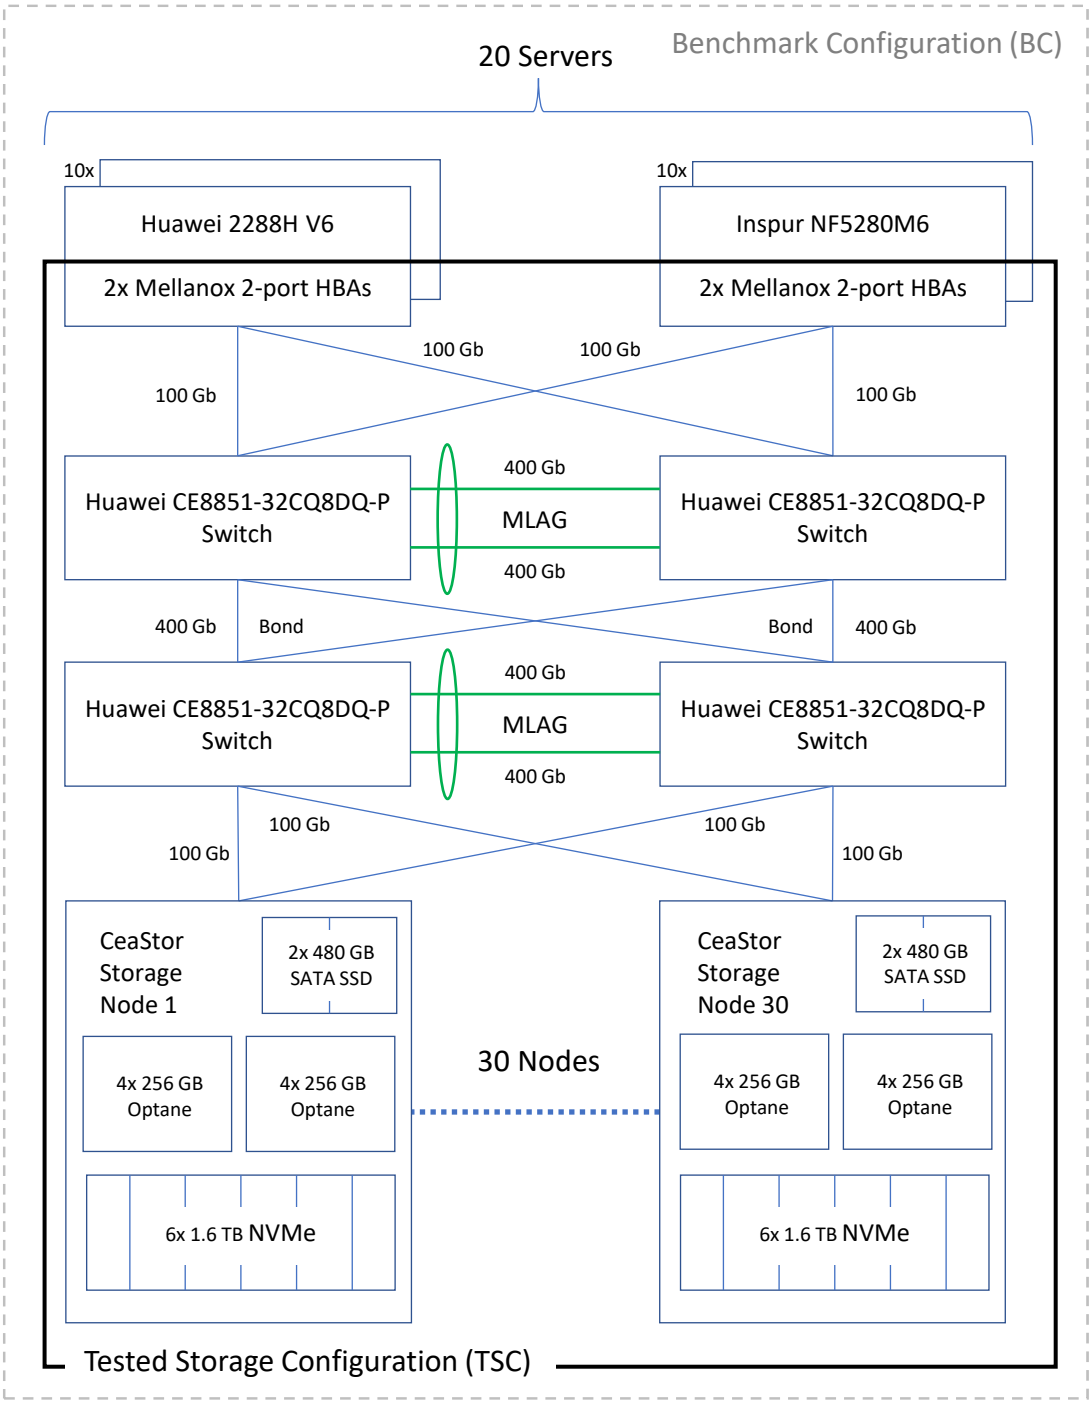

Storage Network Configuration

The Tested Storage Configuration (TSC) involved 30 Ceastor18116E Storage Nodes and 4 100Gb Huawei CE8851-32CQ8DQ-P switches, driven by 20 host systems.

Each host system had 2 Mellanox MCX623106AN-CDAT HBAs. Each HBA was connected to one of two switches. This is a total of 40x 100 Gb connections between the hosts and two of the switches.

Each CeaStor18116E storage node had 2 Mellanox MCX623106AN-CDAT HBAs. Each HBA was connected to one of the other two switches. This is a total of 60x 100 Gb connections between the storage nodes and the other two switches.

The 4 switches were interconnected with 400GE ports.

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 Application Storage Unit Mapping

The following table details the capacity of the Application Storage Units (ASUs) and how they are mapped to logical volumes (LVs). All capacities are reported in GB.

	LV per ASU	LV Capacity	Used per LV	Total per ASU	% ASU Capacity	Optimized*
ASU-1	9	6,356.5	6,356.5	57,208.9	45.0%	No
ASU-2	9	6,356.5	6,356.5	57,208.9	45.0%	No
ASU-3	2	6,356.5	6,356.5	12,713.1	10.0%	No
	SPC-1 ASU Capacity			127,131	*See Space Optimization Techniques	

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. All capacities are reported in GB.

Devices	Count	Physical Capacity	Total Capacity
Optane	240	256.0	61,440.0
NVMe SSD	180	1,600.0	288,000.0
	Total Physical Capacity		**349,440**
	Physical Capacity Utilization		**36.38%**

Data Protection

The data protection level used for all LVs was **Protected 2 (Replication)**, which was accomplished by providing fully redundant pathways from each host to the storage cluster where all data was replicated and distributed on two separate storage nodes.

Space Optimization Information

Description of Utilized Techniques

The TSC did not use any space optimization techniques.

Physical Free Space Metrics

The following table lists the Physical Free Space as measured at each of the required points during test execution. If space optimization techniques were not used, "NA" is reported.

Physical Free Space Measurement	Free Space (GB)
After Logical Volume Creation	NA
After ASU Pre-Fill	NA
After Repeatability Test Phase	NA

Space Optimization Metrics

The following table lists the required space optimization metrics. If space optimization techniques were not used, "NA" is reported.

Metric	Value
SPC-1 Space Optimization Ratio	NA
SPC-1 Space Effectiveness Ratio	NA

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).

Measurement Intervals by Test Phase Graph

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

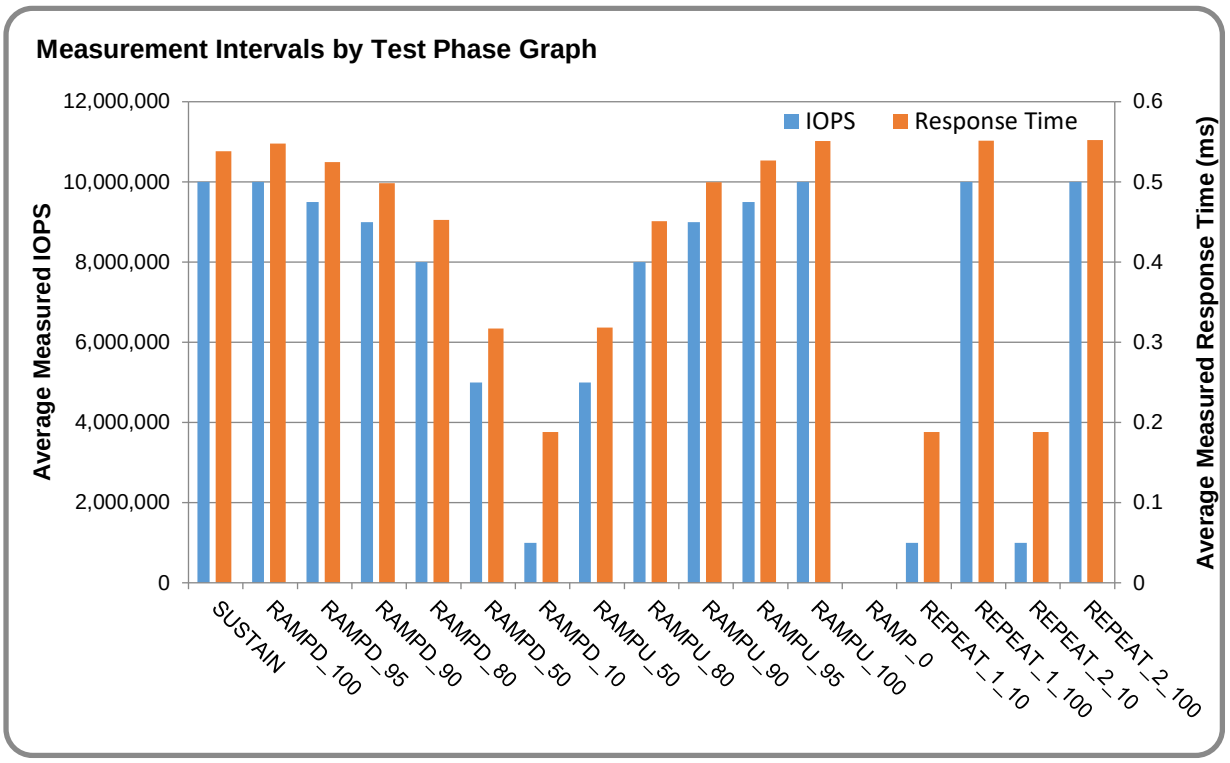

Response Time vs. Throughput Graph

The following graph presents the average Response Times versus the average IOPS for RAMPD_100 to RAMPD_10.

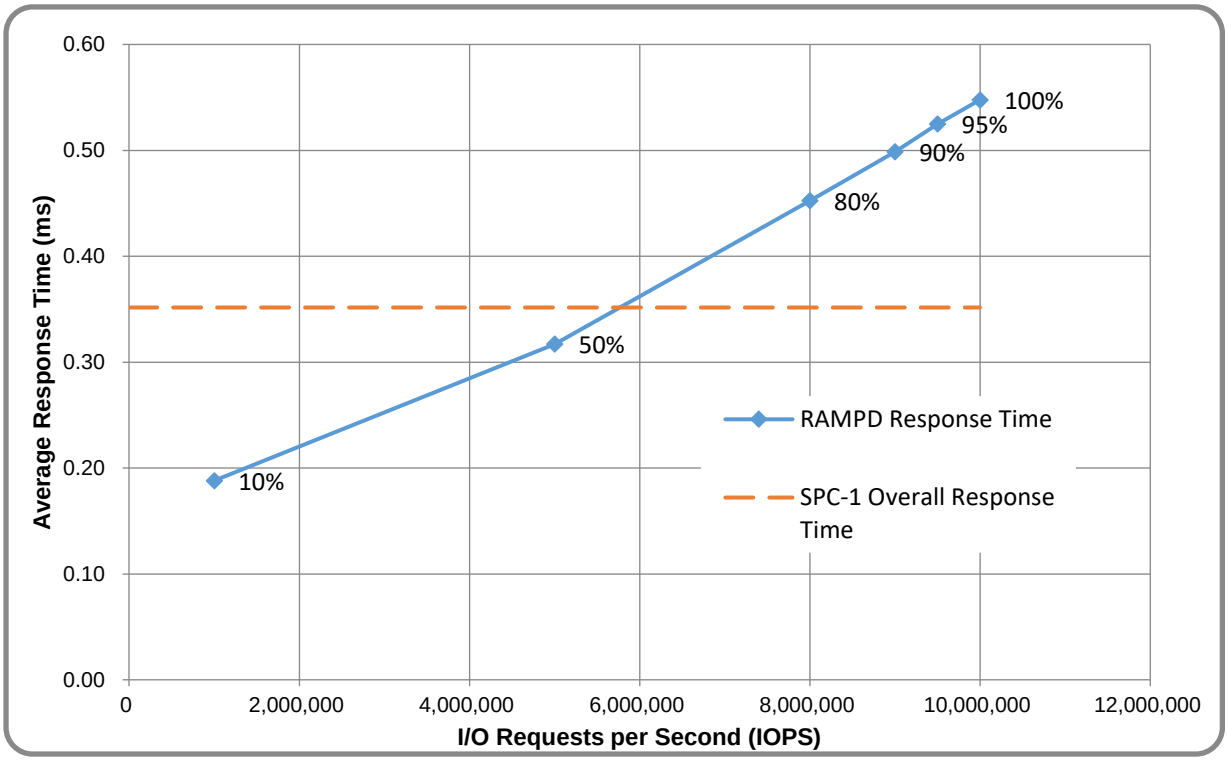

ASU Pre-Fill

The following table provides a summary of the Pre-Fill performed on the ASU prior to testing.

ASU Pre-Fill Summary			
Start Time	14-Apr-23 15:33:37	**Requested IOP Level**	50,000 MB/sec
End Time	14-Apr-23 16:17:22	**Observed IOP Level**	48,449 MB/sec
Duration	0:43:44	For additional details see the Supporting Files.	

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	14-Apr-23 16:32:10	14-Apr-23 21:02:06	4:29:56
Measurement Interval	14-Apr-23 21:02:06	15-Apr-23 05:02:07	8:00:01

SUSTAIN – Throughput Graph

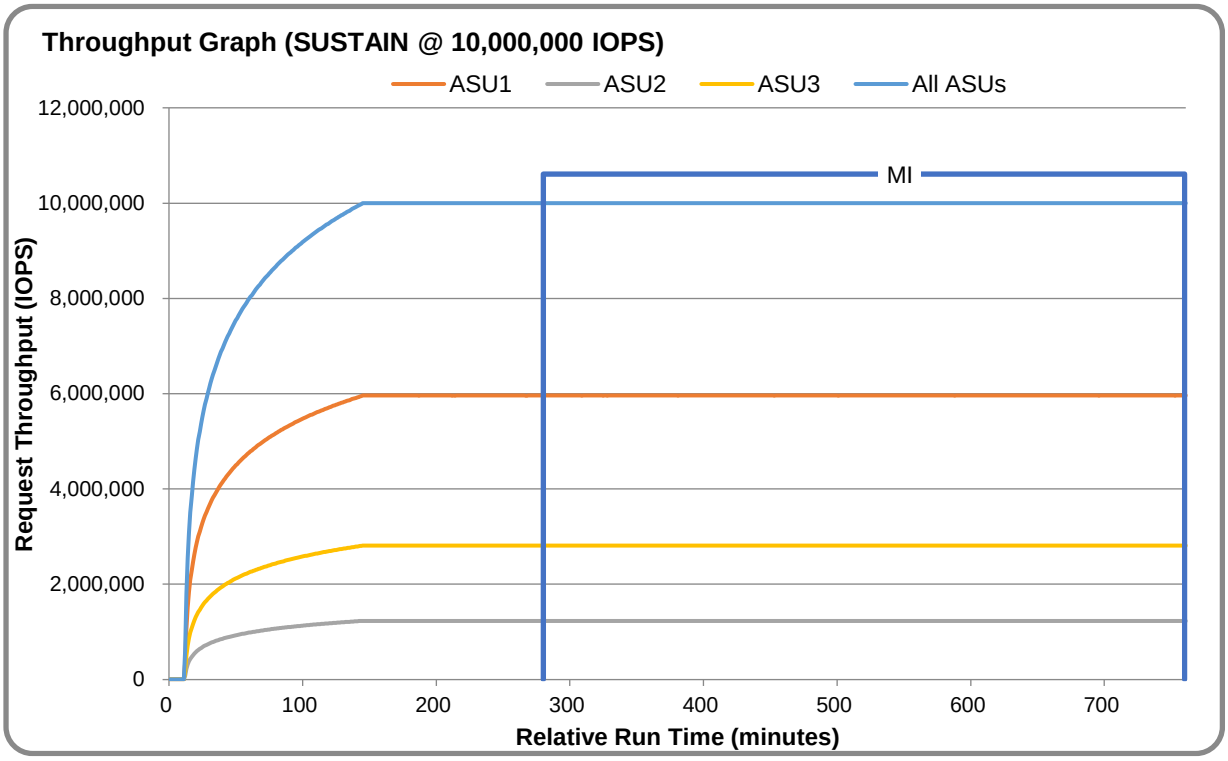

SUSTAIN – Response Time Graph

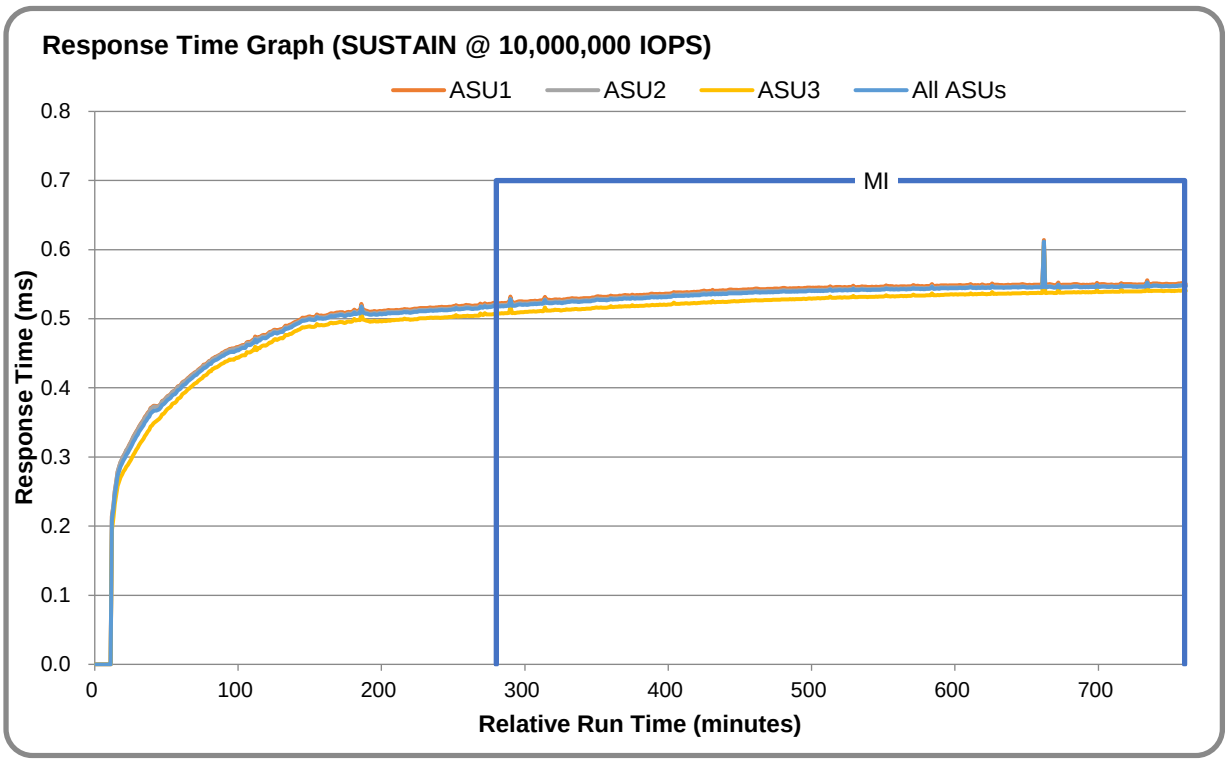

SUSTAIN – Data Rate Graph

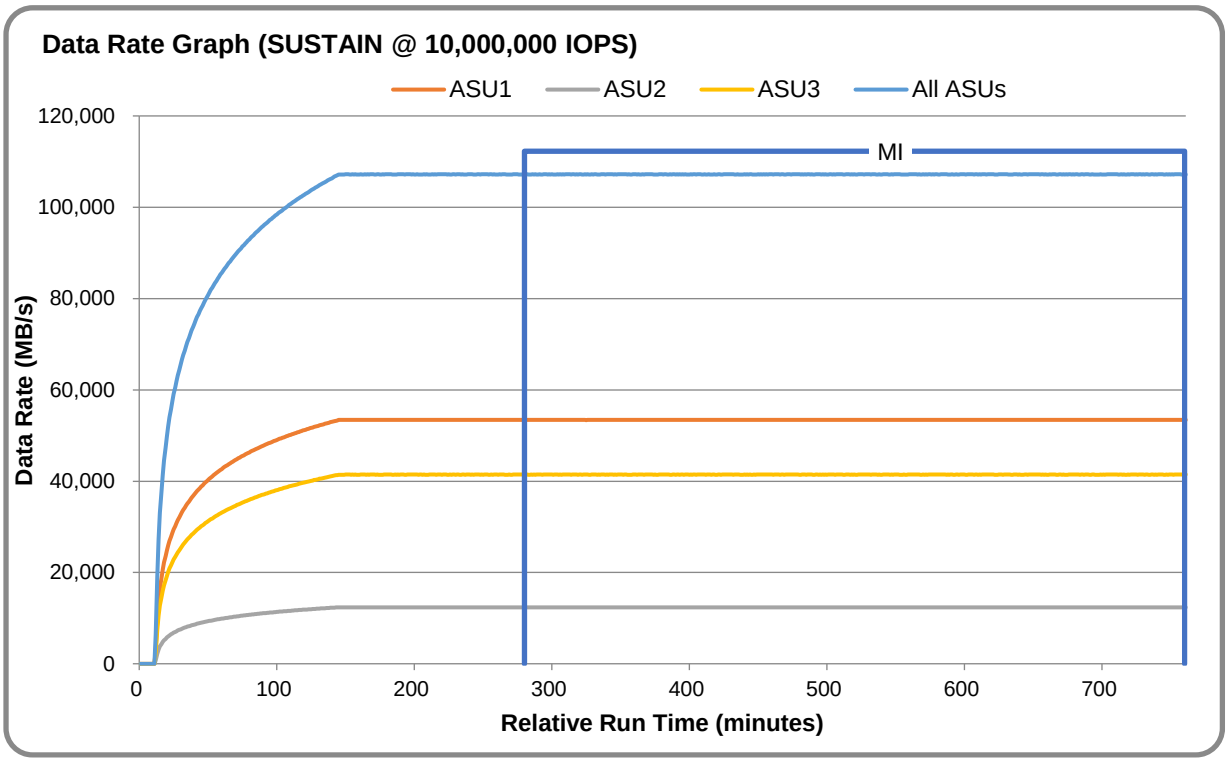

SUSTAIN – Response Time Frequency Graph

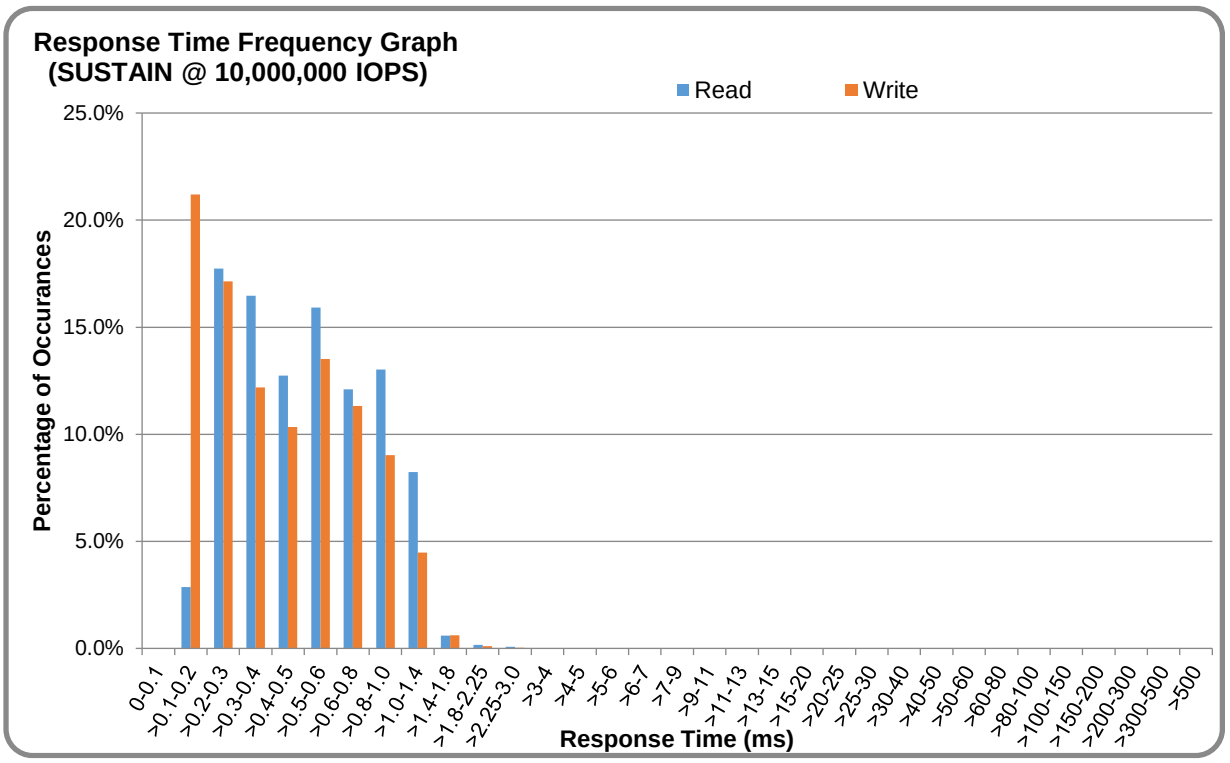

SUSTAIN – Intensity Multiplier

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 Defined 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	0.0002	0.0001	0.0001	0.0001	0.0003	0.0001	0.0002	0.0001
Difference	0.004%	0.002%	0.003%	0.001%	0.006%	0.005%	0.005%	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	15-Apr-23 05:03:06	15-Apr-23 05:06:07	0:03:01
Measurement Interval	15-Apr-23 05:06:07	15-Apr-23 05:16:07	0:10:00

RAMPD_100 – Throughput Graph

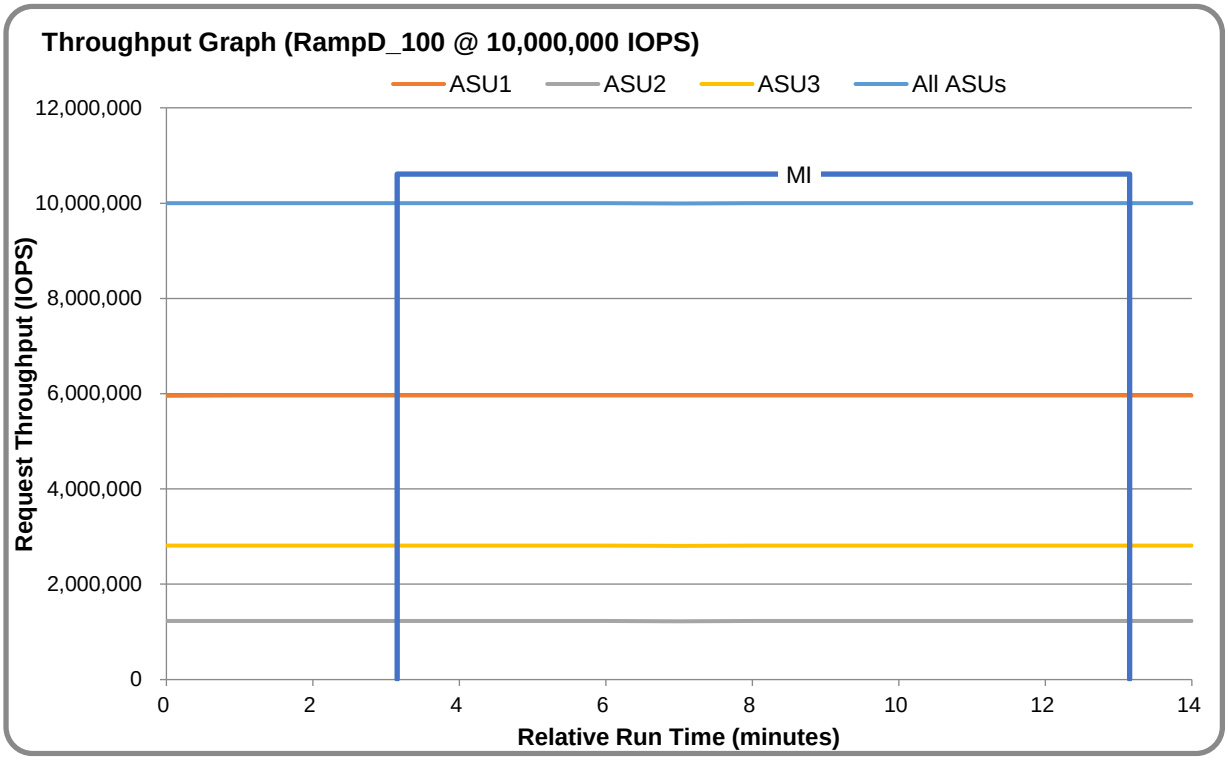

RAMPD_100 – Response Time Graph

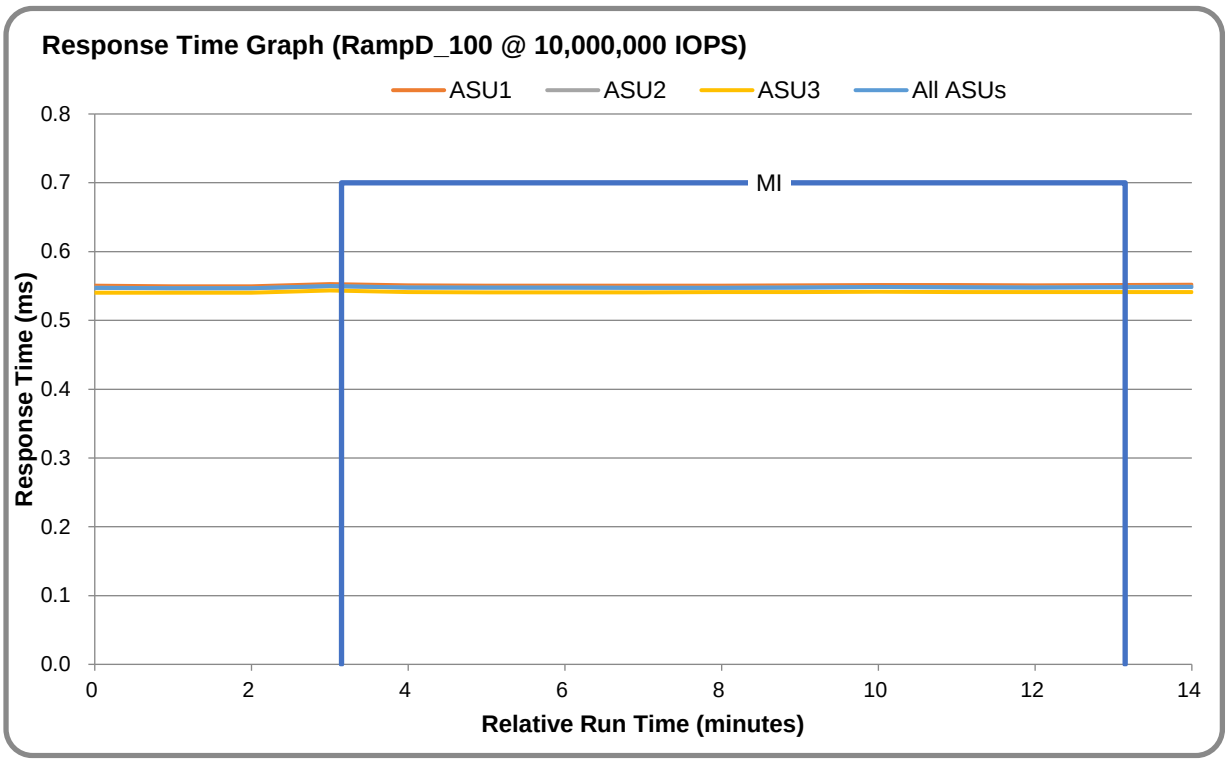

RAMPD_100 – Data Rate Graph

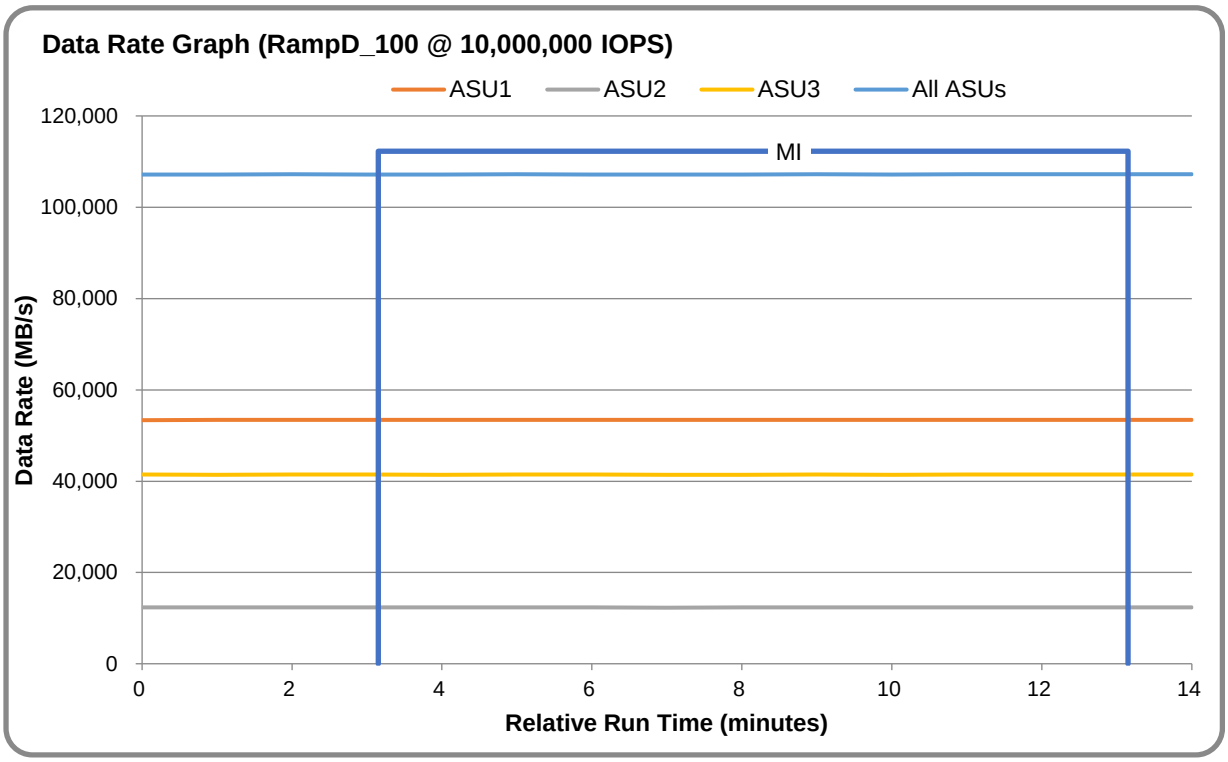

RAMPD_100 – Response Time Frequency Graph

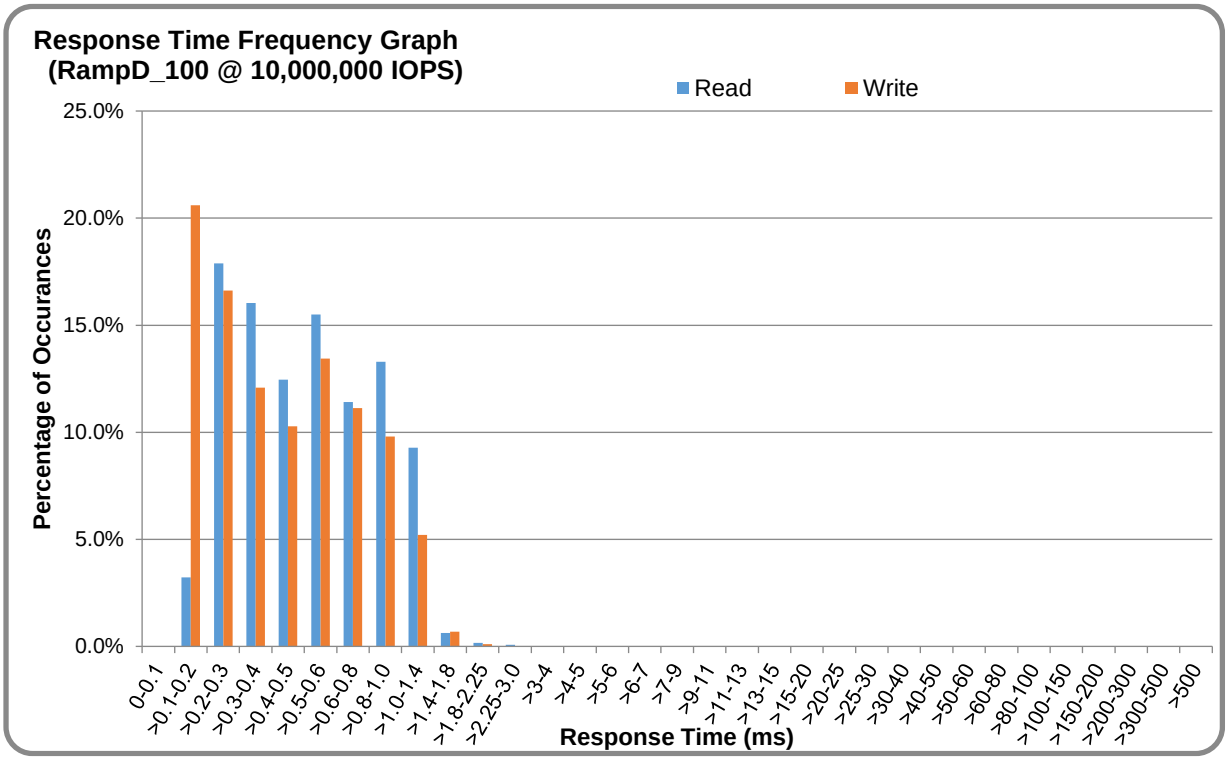

RAMPD_100 – Intensity Multiplier

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 Defined 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	0.0002	0.0001	0.0001	0.0001	0.0003	0.0002	0.0002	0.0001
Difference	0.011%	0.000%	0.012%	0.002%	0.005%	0.003%	0.000%	0.002%

RAMPD_100 – I/O Request Summary

I/O Requests Completed in the Measurement Interval	6,000,432,362
I/O Requests Completed with Response Time <= 30 ms	6,000,427,609
I/O Requests Completed with Response Time > 30 ms	4,753

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

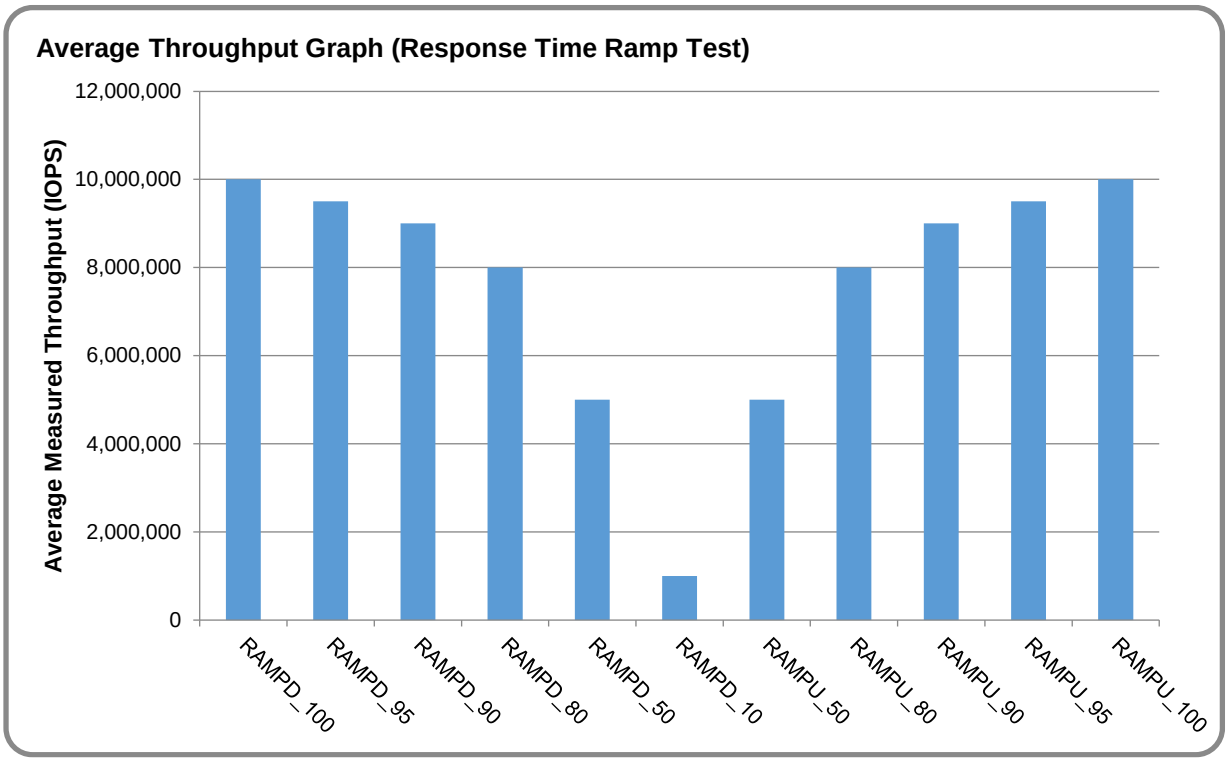

Response Time Ramp Test – Average Response Time Graph

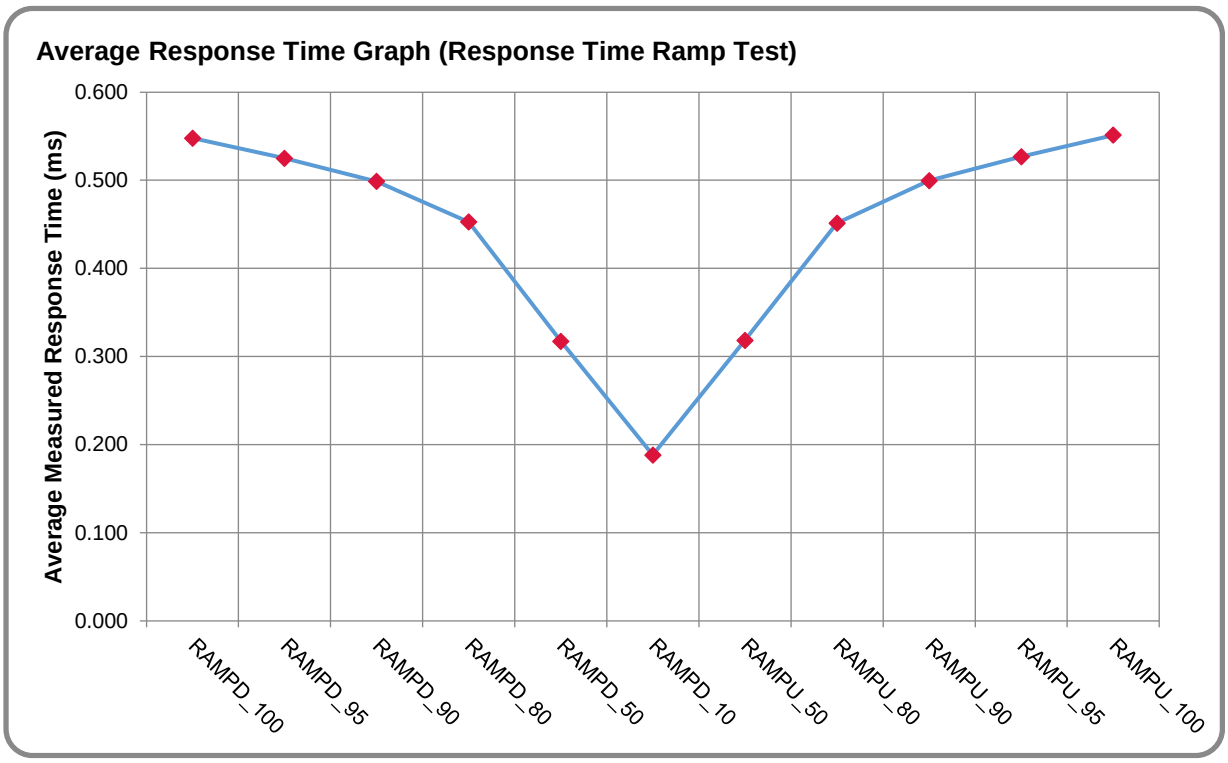

Response Time Ramp Test – RAMPD_10 Response Time Graph

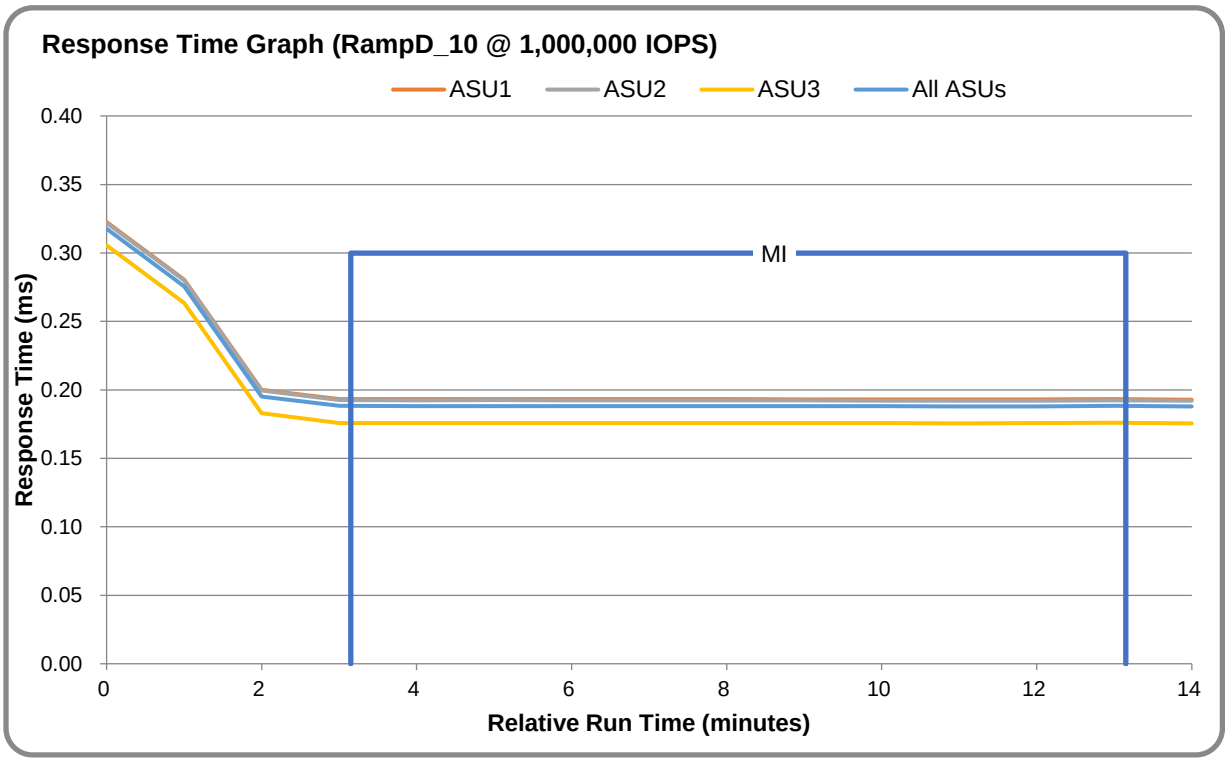

Repeatability Test

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 table below.

Test Phase	100% IOPS	10% IOPS
RAMPD	10,000,690.3	1,000,048.2
REPEAT_1	10,000,517.8	1,000,011.9
REPEAT_2	10,000,422.1	1,000,154.2

REPEAT_1_100 – Throughput Graph

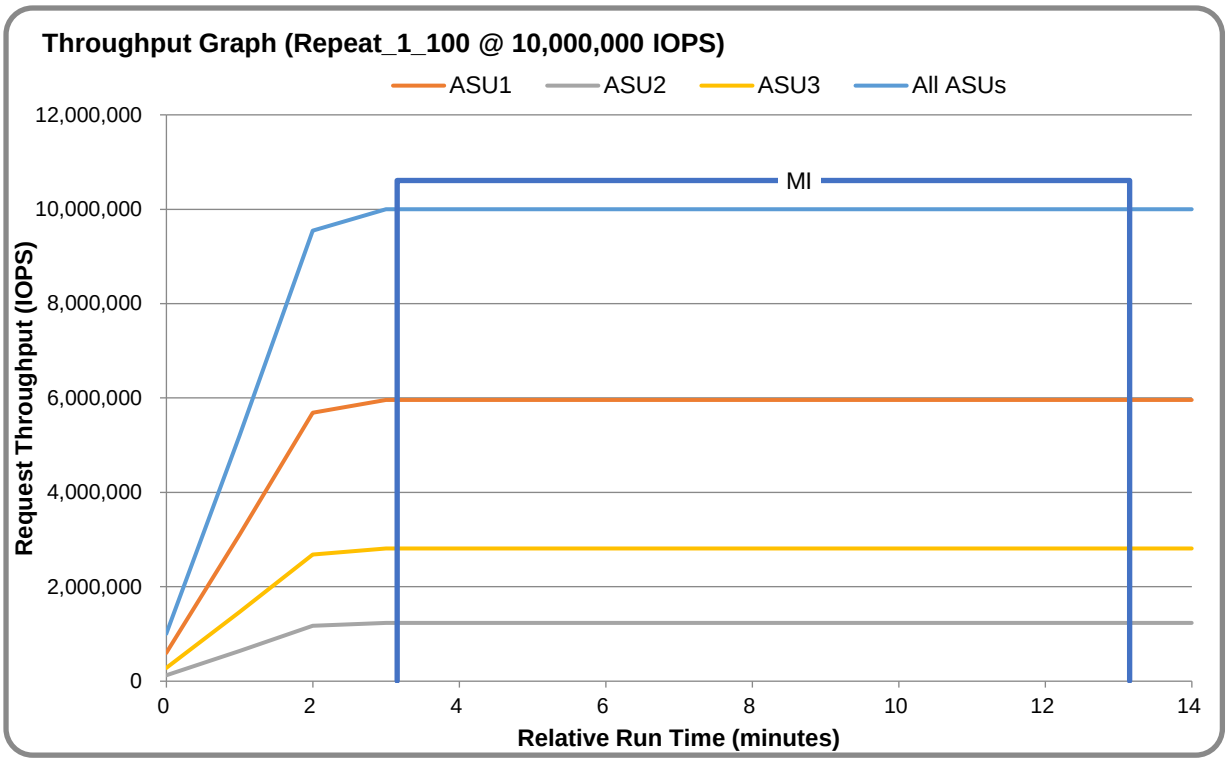

REPEAT_1_100 – Response Time Graph

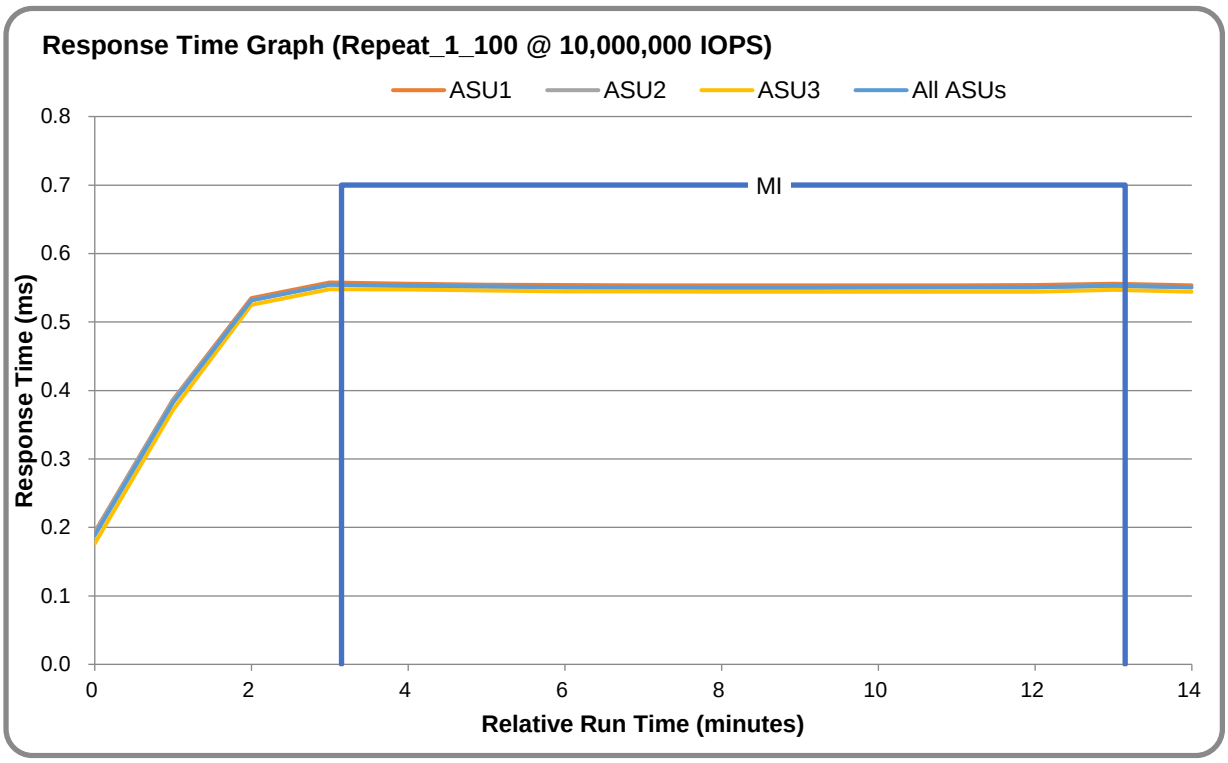

REPEAT_2_100 – Throughput Graph

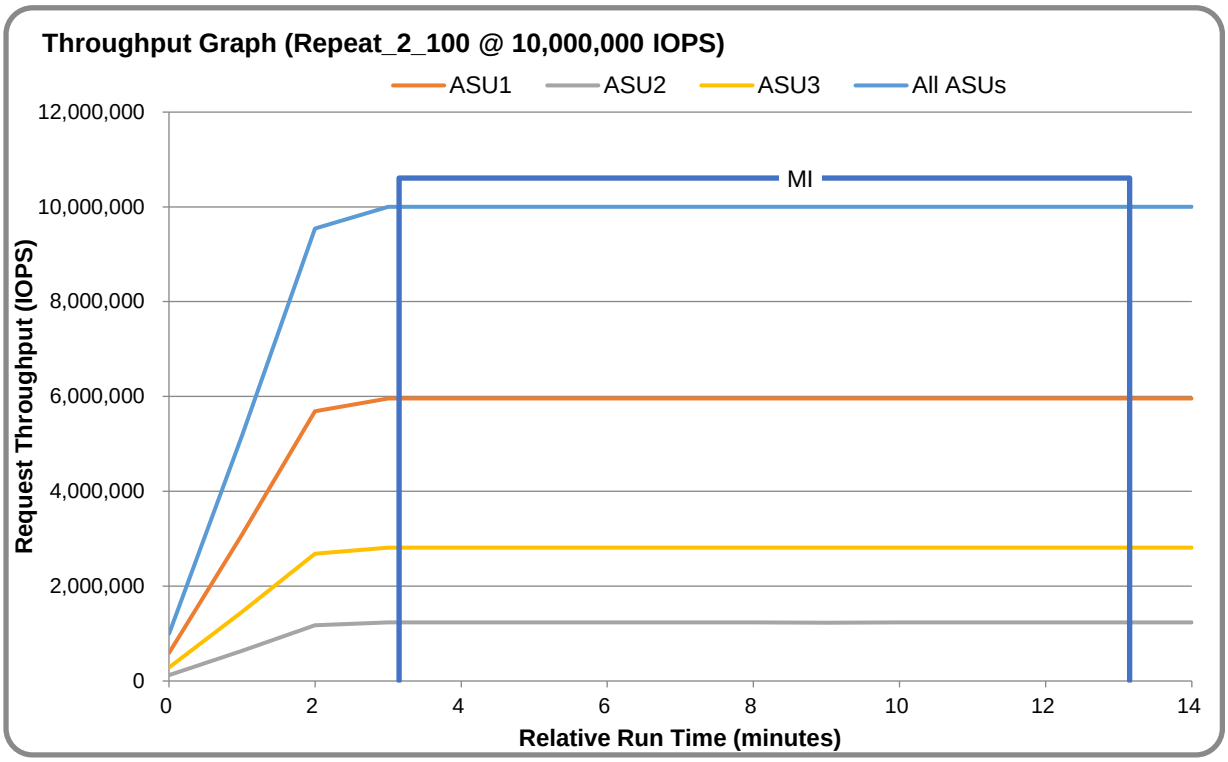

REPEAT_2_100 – Response Time Graph

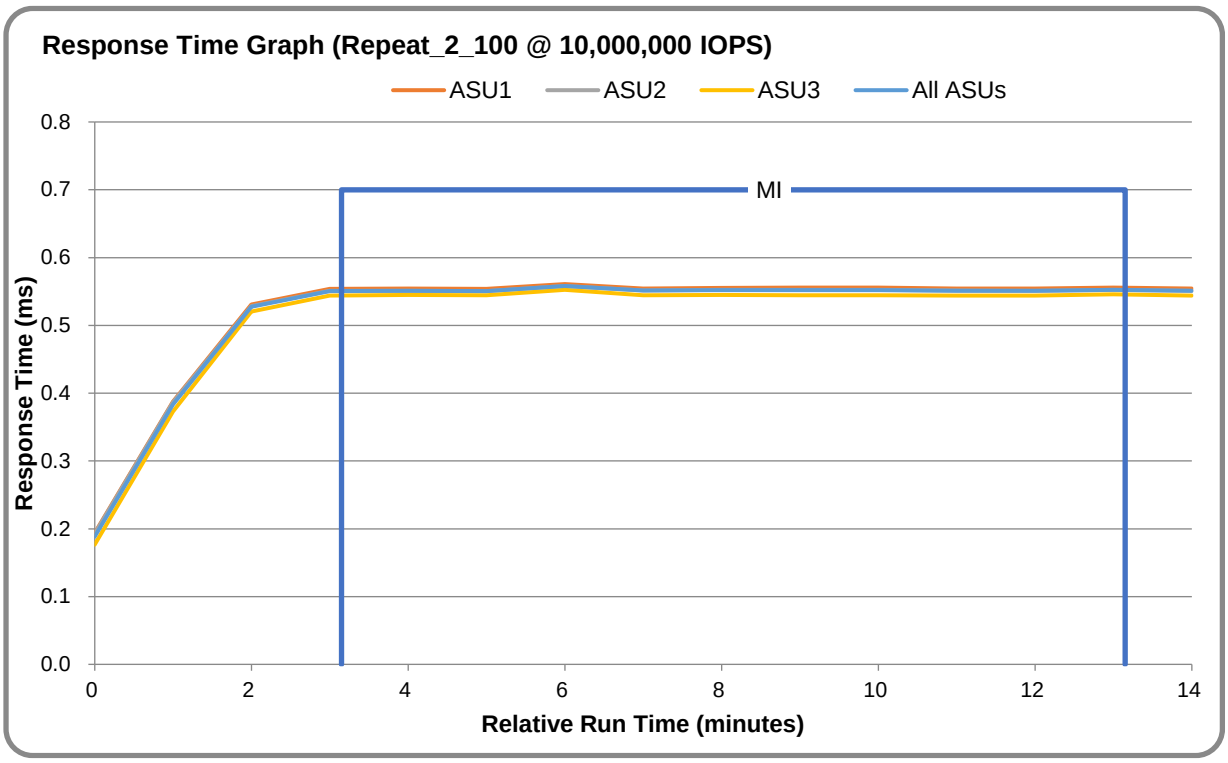

Repeatability Test – Intensity Multiplier

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 Defined 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	0.0002	0.0000	0.0002	0.0001	0.0002	0.0001	0.0003	0.0000
Difference	0.007%	0.001%	0.005%	0.003%	0.006%	0.005%	0.004%	0.001%

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	0.0003	0.0000	0.0001	0.0001	0.0002	0.0001	0.0002	0.0001
Difference	0.011%	0.002%	0.005%	0.000%	0.010%	0.002%	0.004%	0.000%

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	1,397,397,014
Total Number of Logical Blocks Verified	681,305,418
Total Number of Logical Blocks Overwritten	716,091,596
Total Number of Logical Blocks that Failed Verification	0
Time Duration for Writing Test Logical Blocks (sec.)	601
Size in bytes of each Logical Block	8,192
Number of Failed I/O Requests in the process of the Test	0

Committed Data Persistence Implementation

The Ceastor18116E adopts PMem which is a non-volatile memory to store meta data. When data is being written to the storage cluster, the I/O will not return success until all data (including replicated data and meta data) has been written into the PMem and NVME disks.

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**
See Appendix D Step 3	Adjust aio-max-nr	N/A
/D_Creation	**Storage configuration creation**	**root**
hosts.sh	Define environment	/D_Creation
cluster-deploy.sh	Create/format storage cluster	
client_deploy.sh	Create logical volumes	
lun_map.sh	Map volumes on host as NVMe disks	
pool_create.yaml	Define NVMe disk pool	
/E_Inventory	**Configuration inventory**	**root**
collect_after_restart.sh	Collect reboot logs after restart	/E_Inventory
collect_before_restart.sh	Collect reboot logs before restart	
collect_ceastor_info.sh	Collect CeaStor profile	
collect_dmidecode.sh	Collect node system info	
collect_lsblk.sh	Collect node disk info	
collect_network_card.sh	Collect node network info	
/F_Generator	**Workload generator**	**root**
cestc.HST	Host definition file	/F_generator
multi_lun.asu	ASU definition file	
persist_test_2.sh	Run Persist2 test	
spc_run.sh	Run Init – Persist1	

APPENDIX B: THIRD PARTY QUOTATION

All components are available directly through the Test Sponsor (China Electronics Cloud Technology Co., Ltd.).

APPENDIX C: TUNING PARAMETERS AND OPTIONS

See Appendix D Step 3.

APPENDIX D: STORAGE CONFIGURATION CREATION

Step 0: Edit configure file

Edit host.h for identify the IP address of all storage node and host node, as well as the size and number of logical volume.

Step 1: Create and format storage cluster:

Running cluster-deploy.sh:

```
[root@node55 script]# ./cluster-deploy.sh
start ceastor
format
Format Summary:
  Hosts                         SCM Devices Disk Devices
  -----                         ----------- ------------
  10.255.153.[2-17,19-26,28-33] 2           6
create pool
Pool created with 5.66%,94.34% storage tier ratio
-------------------------------------------------
  UUID                          : 98a2a8ff-c983-475d-9812-2d14abfe9d90
  Service Ranks                 : [0,10,29]
  Storage Ranks                 : [0-59]
  Total Size                    : 305 TB
  Storage tier 0 (Meta-Space) : 17 TB (288 GB / rank)
  Storage tier 1 (Data-Space) : 288 TB (4.8 TB / rank)

deploy finished
```

Step 2: Create logical volumes and map them as NVMe disks on host nodes

2.1 Running client-deploy.sh

Creating 1 container for creation logical volumes.

Creating 20 logical volumes each with 5920GiB(6207.5GB) and start agent on all host node.

```
start agent
create container BLK001
##########start init global#############
worker 116 ThreadPool worker start
worker 116 ThreadPool worker start
worker 116 ThreadPool worker start
start 65: contextcnt 4, threadcnt2
worker 116 ThreadPool worker start
start 78: ioctx#0 thread0
start 78: ioctx#0 thread1
start 78: ioctx#1 thread0
start 78: ioctx#1 thread1
start 78: ioctx#2 thread0
start 78: ioctx#2 thread1
start 78: ioctx#3 thread0
start 78: ioctx#3 thread1
start 78: ioctx#4 thread0
start 78: ioctx#4 thread1
start_watcher_tick
cbd_init 782 memery_list_aio_comp_size 0
cbd_init 783 memery_list_aio_comp_retry_count 0
##################container create test START####################
Pool uuid:98a2a8ff-c983-475d-9812-2d14abfe9d90
container uuid:eb6709f6-ccd9-5ebd-bee4-0320df87feee
container logic id:BLK001
######start query cont
######end query cont, start get prop entry
######end get entry
######start get entry roots
######verify BDM
######End verify BDM!
##################container create test END####################
stop 193 ThreadPool stop
worker 130 ThreadPool worker finish
worker 130 ThreadPool worker finish
worker 130 ThreadPool worker finish
worker 130 ThreadPool worker finish
stop 208 ThreadPool stopped
stop 98: stop ioctx0
stop 98: stop ioctx1
stop 98: stop ioctx2
stop 98: stop ioctx3
stop 98: stop ioctx4
create image BLK001-00000001, size=5920G
```

2.2 Connect Logical Volume to Host:

Run Lun_map.sh script to map all logical volume on host as NVMe disk.

Map logical volume BLK001-00000001 to BLK001-00000009 for ASU-1

Map logical volume BLK001-00000010 to BLK001-00000018 for ASU-2

Map logical volume BLK001-00000019 to BLK001-00000020 for ASU-3

```
<<<nqn.2014-08.org.nvmexpress:uuid:fd6257c5-7f35-440f-8ede-67475e8bee28>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_1.sock nvmf_create_subsystem nqn.2014-08.org.nvmexpre
ss:uuid:fd6257c5-7f35-440f-8ede-67475e8bee28 -a -s SPDK00000000000007 -d SPDK_Controller
listen on 10.255.153.34:4421
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_1.sock nvmf_subsystem_add_listener nqn.2014-08.org.nv
mexpress:uuid:fd6257c5-7f35-440f-8ede-67475e8bee28 -t rdma -a 10.255.153.34 -s 4421
nvme connect...
<<<5bb9718e-3791-46a8-be75-414fc380dd49 BLK001-00000008 5920G 6356551598080>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_2.sock bdev_cbd_create -b cbd8 5bb9718e-3791-46a8-be7
5-414fc380dd49 BLK001-00000008 6356551598080 512
cbd8
<<<nqn.2014-08.org.nvmexpress:uuid:0ef6cada-277d-4011-abe8-b90fd3bff5ee>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_2.sock nvmf_create_subsystem nqn.2014-08.org.nvmexpre
ss:uuid:0ef6cada-277d-4011-abe8-b90fd3bff5ee -a -s SPDK00000000000008 -d SPDK_Controller
listen on 10.255.153.134:4421
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_2.sock nvmf_subsystem_add_listener nqn.2014-08.org.nv
mexpress:uuid:0ef6cada-277d-4011-abe8-b90fd3bff5ee -t rdma -a 10.255.153.134 -s 4421
nvme connect...
<<<5bb9718e-3791-46a8-be75-414fc380dd49 BLK001-00000009 5920G 6356551598080>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_1.sock bdev_cbd_create -b cbd9 5bb9718e-3791-46a8-be7
5-414fc380dd49 BLK001-00000009 6356551598080 512
cbd9
<<<nqn.2014-08.org.nvmexpress:uuid:357a964b-bde5-44fc-a94f-ee531596e4b7>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_1.sock nvmf_create_subsystem nqn.2014-08.org.nvmexpre
ss:uuid:357a964b-bde5-44fc-a94f-ee531596e4b7 -a -s SPDK00000000000009 -d SPDK_Controller
listen on 10.255.153.34:4421
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_1.sock nvmf_subsystem_add_listener nqn.2014-08.org.nv
mexpress:uuid:357a964b-bde5-44fc-a94f-ee531596e4b7 -t rdma -a 10.255.153.34 -s 4421
nvme connect...
<<<5bb9718e-3791-46a8-be75-414fc380dd49 BLK001-00000010 5920G 6356551598080>>>
/home/zyh/protocol/spdk/scripts/rpc.py -s /var/tmp/tgt_2.sock bdev_cbd_create -b cbd10 5bb9718e-3791-46a8-be
75-414fc380dd49 BLK001-00000010 6356551598080 512
cbd10
```

After finished Mapping:

```
[root@node34 script]# lsblk
NAME        MAJ:MIN RM   SIZE RO TYPE MOUNTPOINT
sda           8:0    0 446.1G  0 disk
├─sda1        8:1    0   600M  0 part /boot/efi
├─sda2        8:2    0     1G  0 part /boot
└─sda3        8:3    0 444.5G  0 part
  ├─cl-root 253:0    0    70G  0 lvm  /
  ├─cl-swap 253:1    0     4G  0 lvm  [SWAP]
  └─cl-home 253:2    0 816.6G  0 lvm  /home
sdb           8:16   0 446.1G  0 disk
└─sdb1        8:17   0 446.1G  0 part
  └─cl-home 253:2    0 816.6G  0 lvm  /home
nvme0n1     259:3    0   5.8T  0 disk
nvme1n1     259:7    0   5.8T  0 disk
nvme2n1     259:11   0   5.8T  0 disk
nvme3n1     259:15   0   5.8T  0 disk
nvme4n1     259:19   0   5.8T  0 disk
nvme5n1     259:23   0   5.8T  0 disk
nvme6n1     259:27   0   5.8T  0 disk
nvme7n1     259:31   0   5.8T  0 disk
nvme8n1     259:35   0   5.8T  0 disk
nvme9n1     259:39   0   5.8T  0 disk
nvme10n1    259:41   0   5.8T  0 disk
nvme11n1    259:43   0   5.8T  0 disk
nvme12n1    259:45   0   5.8T  0 disk
nvme13n1    259:47   0   5.8T  0 disk
nvme14n1    259:49   0   5.8T  0 disk
nvme15n1    259:51   0   5.8T  0 disk
nvme16n1    259:53   0   5.8T  0 disk
nvme17n1    259:55   0   5.8T  0 disk
nvme18n1    259:57   0   5.8T  0 disk
nvme19n1    259:59   0   5.8T  0 disk
```

Step 3: Change the Scheduler on each Host System

Run command: echo 1048576 > /proc/sys/fs/aio-max-nr on all host nodes

APPENDIX E: CONFIGURATION INVENTORY

The scripts used to collect the configuration inventory and the log files that were generated are available in the Supporting Files (see Appendix_A)

APPENDIX F: WORKLOAD GENERATOR

The ASUs accessed by the SPC-1 workload generator are defined in multi_lun.asu. The test phases up through PERSIST1 are executed by spc_run.sh. PERSIST2 is executed by persist_test_2.sh.

The files are included in the Supporting Files (see Appendix_A).