Samsung PM9A3 All Flash Reference Platform

Jeffrey Wong, Edward Aguilar, Bridget Davis, Qirui Yang, and Devasena Inupakutika, *Datacenter Technology and Cloud Solutions, Samsung Semiconductor Inc.*

Abstract

In this document, we present an all-flash reference platform utilizing Weka high performance filesystem and Samsung SSD, that can be leveraged by IT strategists to meet the performance challenges faced in today's datacenters. Our performance results show that the Weka high performance filesystem paired with Samsung's latest datacenter NVMe SSD, the PM9A3, delivers exceptional performance for a wide variety of demanding workloads found in today's datacenters.

Weka is a software defined, distributed, parallel file system designed to deliver the highest-performance file services by leveraging NVMe flash. The distributed file system distributes the data across multiple servers and devices, providing high performance access for both throughput intensive and I/O intensive applications. The Weka distributed file system can fully leverage the combination of high performance and enterprise functionality capabilities of Samsung's flagship datacenter SSD offering, PM9A3.

We worked closely with Weka engineering to architect and configure the Weka cluster; tuning the system to achieve the highest performance possible with the given hardware. In order to benchmark the environment, we used the SPECstorage Solution 2020 benchmark tool. This benchmark is designed to simulate application-based workloads and real-world scenarios. The benchmark is capable of simulating workloads of the modern large scale data centers. In this study, the workloads applied include artificial intelligence (AI), genomics, electronic design automation (EDA), and video data acquisition (VDA). Our methodology is to establish the hardware and software testbed environment, run a series of SPECstorage Solution 2020 benchmarks, and tune the system parameters until optimal performance is reached for all workloads that were the focus of this study. The result of this work was a compact, high-density system that achieved the highest ever recorded performance for 4 of the 5 SPECstorage Solution 2020 audited benchmarks. Results are published on SPEC website[1].

1. Introduction

Software Defined Storage continues to gain momentum in the modern datacenter storage market. As datacenter storage infrastructures and technologies evolve, the need for speed and flexibility in the storage system becomes evident. NVMe is at the forefront of these advancements. Weka is a software defined storage system that is designed to meet the needs of today's I/O intensive workloads. It is designed to take advantage of the parallelism provided by NVMe flash. To showcase Samsung's datacenter class NVMe, test their capabilities to power storage systems, and promote NVMe adoption, we chose to architect a compact all flash scale-out reference platform that includes Samsung's newest PM9A3 NVMe devices and Weka's scale out, parallel file system. We put the system to the test by applying realistic workloads that are seen in today's datacenter.

In addition to showcasing Samsung NVMe, we aim to provide IT professionals with a reference architecture, meaning if the same equipment is purchased and configured as described in this work, the system will achieve the same performance. To achieve this, we chose benchmarks provided by the Standard Performance Evaluation Corporation (SPEC) to validate performance. SPEC benchmarks are fully audited by a team of industry experts. All hardware, software, and configuration details must be disclosed in order to be published. This ensures the repeatability of the performance results.

To showcase the solution, we wanted to apply a broad range of common datacenter workloads that are storage intensive. SPECstorage Solution 2020 contains a comprehensive set of workloads that are both storage centric and representative of real-world workloads found in the modern datacenter.

1.1. Weka

WekaFS [2] is a fully-distributed, parallel file system that was written entirely from scratch to deliver the highest-performance file services by leveraging NVMe flash. The software also includes integrated tiering that seamlessly expands the namespace to and from hard disk drive (HDD) object storage, without the need for special data migration software or complex scripts; all data resides in a single namespace for easy access and management.

1.2. SPECstorage Solution 2020 Benchmark

The SPEC benchmark simulates several different workload types that are common among data centers today. The benchmark starts a number of processes on the client hosts that perform file operations resembling the workload under test. Each workload has a unique ratio of types of operations, file layout, and distribution of IO size. The business metric is a standard way to describe the size of workload for any of the workloads

[1]https://www.spec.org/storage2020/results/

as shown in Table 1. As the value of business metric increases, the number of load generating processes will linearly increase.

Workload	Description	Business Metric
VDA	Video data acquisition	STREAMS
EDA_BLENDED	Electronic design automation	JOBS
AI_IMAGE	AI Image processing	JOBS
GENOMICS	Genomic processing	JOBS

Table 1: Workloads description

Workload	Storage target capacity requirements	Client memory requirements
VDA	24 GiB per STREAM	100 MiB per STREAM
EDA_BLENDED	11 GiB per JOB	520 MiB per JOB
AI_IMAGE	100 GiB per JOB	1.7 GiB per JOB
GENOMICS	3.5 GiB per JOB	416 MiB per JOB

Table 2: Workloads memory requirements

Table 2 shows the amount of data that is created for each business metric. The amount of client memory required per business metric is also described. It is to be noted that as business metric increases, it is uniformly distributed across the number of clients used in the test. There is one client that is designated the Prime Client which coordinates the load among all the other clients. This Prime Client is also the one which the user will start the benchmark program from. A configuration file is passed into the benchmark on runtime. This configuration file describes the workload to be run, the business metric of the tests, and clients to be used for the test.

Table 3 presents the list of notations that are used in the paper. Each workload consists of the subcomponents that are defined by the combination of various parameters corresponding to categories such as file operation distribution, execution, global, content and access patterns, and read-, write-transfer size distributions. Genomics workload has an additional set of parameters for file size distribution category. Further details about the SPECstorage workloads and the start, end bytes for multiple read, write transfer sizes and file size distribution slots can be referred in [1].

1.2.1 AI Image

This workload simulates the AI Tensorflow image processing environments. The original traces were collected using Nvidia DGX systems running COCO, Resenet50, and CityScape datasets. There are four components which make up this workload, two are the pre-processing phase and the other two are the AI training phase.

- AI_SF: small image file ingest
- AI_TF: tensorflow record creation
- AI_TR: training consumption of tensor flow records
- AI_CP: checkpoint functionality

Notations	Description
AI	Artificial intelligence
VDA	Video data acquisition
EDA	Electronic design automation
WL	Workload
BM	Business metric
ROR	Requested op rate
AOR	Achieved Op rate in Ops/s
Cl Proc	Number of processes per client
CDS	Client Data Set size in MiB
SDS	Total Starting Data Set size in MiB
MFS	Max File Space in MiB

Table 3: List of notations for reading easiness

JOBS is the business metric for the AI_Image workload (Table 1). AI_SF and AI_TR consists of 37% and 95% read operations, respectively. AI_TF and AI_CP are 100% write workloads. AI_CP occurs infrequently during the typical run of the AI Image workload. Only AI_SF has a separate set of global parameters that consist of 4 processes per BM (JOBS for AI Image workload), a rate of 100 operations per process (henceforth, referred to as oprate), 1 MiB average file size, 3 directories per process, and 200 files per directory. The other salient set of parameter-values for the four subcomponent workloads can be accessed in [1].

1.2.2 Genomics

This workload simulates the genomics pipeline in both commercial and research facilities that perform genetic analysis. Traces were taken during these activities and the original genome data has been sanitized. The I/O behavior was captured in traces and synthesized by the benchmark. JOBS is the business metric for the Genomics workload (Table 1).

The salient set of parameter-value pairs for the Genomics workload are as follows: 70% read, 8% write, 1% unlink, 12% stat, 2% rand read, 1% rand write, 1% create, 4% access, and 1% chmod operations under file operation distribution category. As execution parameters, the workload starts 4 processes per JOB, 250 oprate per process, average file size of 1613 KiB, 2 directories per proc, and 25 files per directory. The other set of parameter-values for the workload can be accessed in [1].

1.2.3 Video Data Acquisition

Video Data Acquisition (VDA) workload simulates the environment where video cameras are streaming and storing data to backend storage. The goal is to provide as many simultaneous streams as possible while maintaining a certain bit rate and meeting the fidelity constraints. The workload is made up of two components, the data stream (VDA1) and the companion application (VDA2). Each stream is about 36 Mib/s bit rate which is high-definition video. STREAMS is the business metric for the VDA workload (Table 1).

The salient set of parameter-value pairs for the VDA workload subcomponents are as follows: The data stream VDA1 workload has 100% write operations with 1 process per STREAM, 9 oprate per process, 1GiB of average file size, 1 directory per process, and 1 file per directory as execution parameters. Further details about the other set of parameter-values for the workload and its subcomponents can be accessed in [1].

1.2.4 Electronic Design Automation

Electronic Design Automation (EDA blended) workload simulates the designing of semiconductor chips. Many software tools are used to design a chip and fabricate it. A large number of small files and low number of large files comprise this dataset. There is a mix of random and sequential IO and two components make up this workload. The EDA_FRONTEND workload represents the EDA frontend processing applications and EDA_BACKEND represents the EDA backend applications that generate final output files. JOBS is the business metric for the EDA_BLENDED workload. Both frontend and backend EDA workloads have different set of configuration parameters as described in [1].

2. Test Environment

The testbed used for the study is outlined in Table 4 and described in the following subsections. Figure 1 schematically shows the test bed set up.

Component	Model	Type	Details (Model/ Name/ Size/ Version/ Quantity)
8 Compute servers as clients	Dell R7525 [7]	CPU	Dual AMD EPYC 7702
		RAM	512 GB
		NIC 1	1GbE Management network
		NIC 2	1 ConnectX-6 VP providing 2, 200Gb/s ports
		SW	8 Weka clients - Weka [2] v3.12.0 client
2 Switches			Mellanox SN3700
6 Storage Servers	Dell R7515 [3]	CPU	AMD EPYC 7702P
		Storage	15 SSDs - Samsung PM9A3 [4]
		RAM	512 GB
		SW	6 Weka cluster nodes - WekaFS [2] v3.12.0

Table 4: Complete testbed details

2.1. Clients

On each of the client hosts, the Weka filesystem is mounted using *wekafs* file system type. Each client has 8 cores and 30720 MB of memory dedicated to the *wekafs* mount. The default mount mode for weka client caching is write-back. This is the default setting used by the majority of weka customers and thus is the basis of our cluster configuration for testing in this work. Readers are directed to WekaFS architectural white paper [2] for additional details about default cache mode for weka clients. Both network ports are specified in the mount command giving 200 Gb/s connection each. Because they are dual port, a single card on each client can theoretically transmit and receive 400 Gb/s. However, the card is inserted in a 16 lane PCIe Gen 4 slot in the server. The theoretical limit of the PCIe Gen4 x16 slot is 256 Gb/s. Weka enables both ports-so the theoretical max transfer speed is 256 Gb/s per client. Although the theoretical limit is 256 Gb/s, after optimizing the networking configuration in the clients, we are able to achieve a maximum throughput of approximately 192 Gb/s.

2.2. Targets

The backend cluster is setup to use all 6 hosts as storage with a WekaFS distributed data protection scheme of 4+2 erasure coding at the target host level [1], removing the hot spares. Each host is running 3 Weka containers. Each of the 3 containers has 5 PM9A3 assigned to them. One container on each host has 19 CPUs and the other two have 17 CPUs. Two CPUs on the container with 19 CPUs are dedicated to handling traffic from the clients. Of the remaining 17 CPU, 5 CPUs are dedicated to the storage device (one CPU per device), and the rest are compute cores. The raw storage with a total of 90 PM9A3s (3.84 TiB raw SSD capacity each) on the backend cluster is 314 TiB (calculated as $15 \times 6 \times 3.84 \times 〖10〗^{12}/2^{40}$ and 4+2 protection leads to 209 TiB [6]. WekaFS, in addition, reserves 10% of the capacity for internal use. Thus, one filesystem was created, utilizing all 188 TiB (0.9×209 TiB) of usable storage on the backend cluster. The client nodes have one dual ported 200 Gb/s NIC, and each target node has two, dual ported 200 Gb/s NICs in a 16 lane PCIe Gen4 slot (shown in Figure 1). The theoretical maximum throughput of the configuration is 800 Gb/s per target. However as discussed above the limit on the PCIe Gen4 x16 slot is 256 Gb/s per NIC. With two NICs, the theoretical

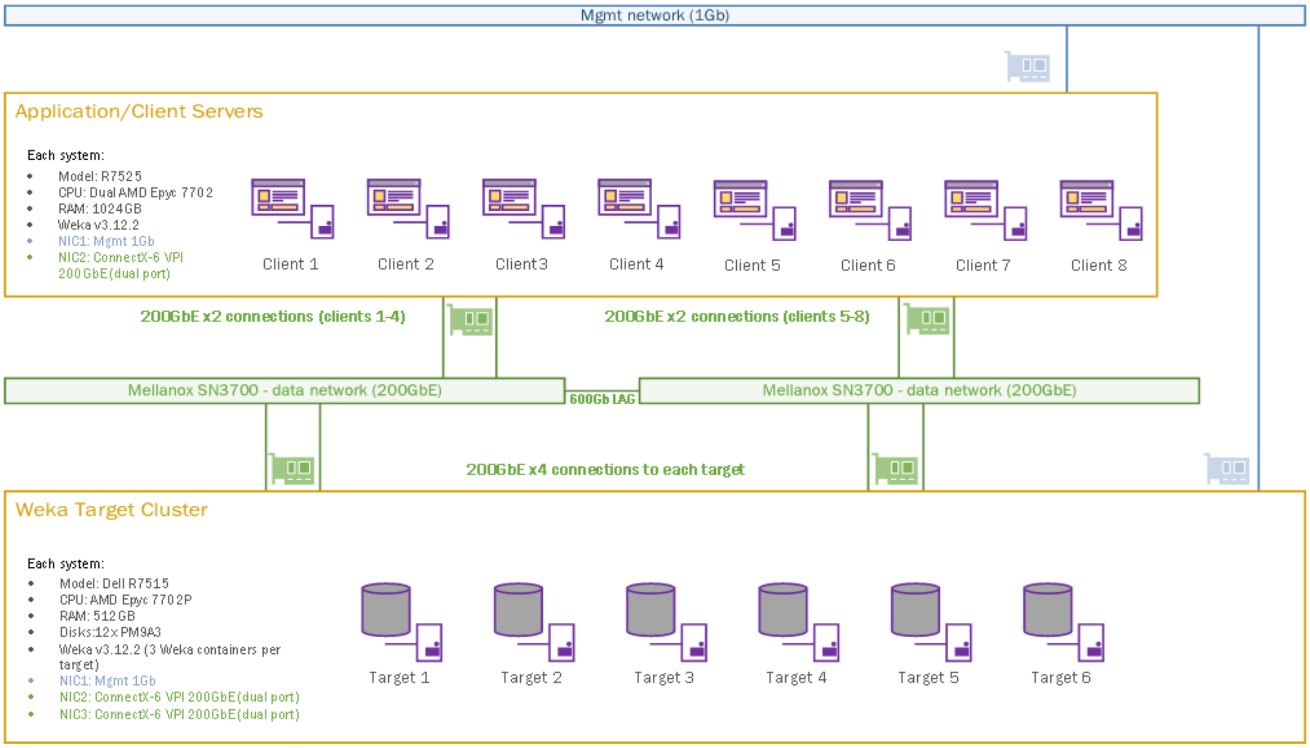

Figure 1: Weka Storage Test Environment

WL	BM	ROR	AOR	Avg Lat (ms)	Total MBps	Read MBps	Write MBps	Cl Proc	CDS MiB	SDS MiB	MFS MiB
AI_ Image	1400	609000	609005.30	3.635	136911.98	134946.72	1965.26	2975	15457750	123662000	134904000
VDA	8000	80000	80054.74	2.668	36846.36	3089.38	33756.98	2000	22528000	180224000	196608000
Genomics	960	960000	960029.83	0.546	81531.12	56203.12	25327.99	480	415851	3326812	3629250
EDA	3300	1485000	1485083.74	1.657	23965.24	12581.35	11383.90	2062	4558769	36470156	39785625

Table 5: SPECstorage workloads results

maximum throughput of the target node is 512 Gb/s. Furthermore, as discussed earlier the actual throughput we anticipate is slightly lower than the theoretical limit. The PCIe Gen4 x4 PM9A3 has a maximum throughput of 6800 MB/s or 54 Gb/s. However, the limit on the Dell R7515 NVMe slot is 32 Gb/s. Each target has 15x PM9A3 devices, providing a maximum SSD throughput of 480 Gb/s. This is well under the 512 Gb/s network limit. The testbed that we have built, therefore, is storage performance limited.

2.3 Switches

The test bed, as seen in figure 1, includes two 32 port switches with each port supporting a full duplex 200 Gb/s of throughput. Of the 64 ports available in the configuration, 16 ports were used for client connections, 24 ports were dedicated for storage nodes and 4 ports were dedicated to inter-switch links (ISLs).

3. Test Results

In an ongoing effort with both Weka and SPEC, we have been tuning the environment and running a series of SPECstorage benchmarking tests after each change. This document will not cover all the tuning that was done in order to achieve the best results. Instead, in the following paragraphs, we will discuss the final rounds of testing and present the best business metric result achieved while running the benchmark. The appendix however consists of the final results of all the business metrics from the benchmark.

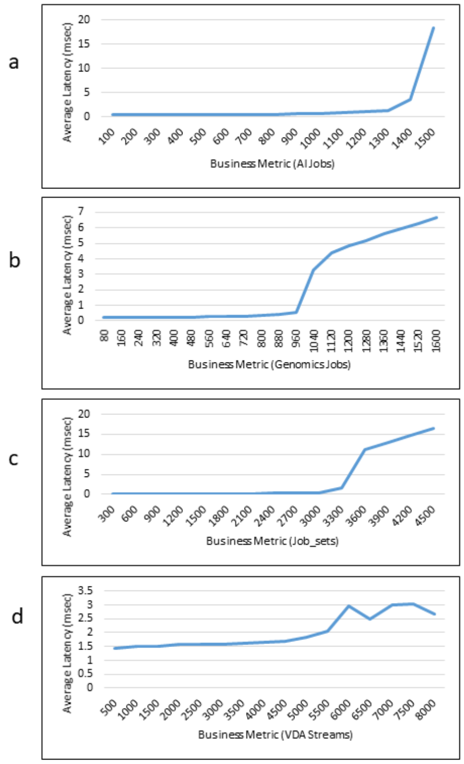

Figure 2: AI_Image, Genomics, EDA_Blended, and VDA SPEC Latency results

3.1. 8x6 Test Environment

The following test results were collected in an 8 by 6 environment, meaning the environment consists of 8 client hosts and 6 target hosts. It is important to note the number of clients used because the workload during the test is spread out across that number of physical machines. Table 5 shows the results from the highest business metric achieved from running the AI Image, VDA, Genomics, and EDA workloads with 8 clients for 300 seconds.

AI Image: The average file size was 9952 KiB. The initial file set is equal to the start data set for all the business metrics. With an incremental business metric of 100, we were able to run up to business metric 1400 before failing the test at 1500. Looking at the latency summary (figure 2 a) from the benchmark logs, at business metric 1400, the average latency jumps up from 4ms to 18ms, and the test fails as a result (figure 5 a). We can see from the benchmark logs that the failure at business metric 1500 is due to the operation rate falling below an acceptable threshold. If we observe the CPU utilization (figure 5 b) on the 8 client hosts, we can see that there is a noticeable increase as business metric increases. At the later portions of the timeline, we see CPU utilization on some of the clients nearly reach a maximum thus implying a CPU bottleneck. The AI Image workload reached a max business metric of 1400 with an average throughput of 137 GB/s (figure 3 a) against the storage system and achieved the oprate of 609K (figure 4 a). The clients hit a CPU bottleneck before the network, memory, or disk space is reached.

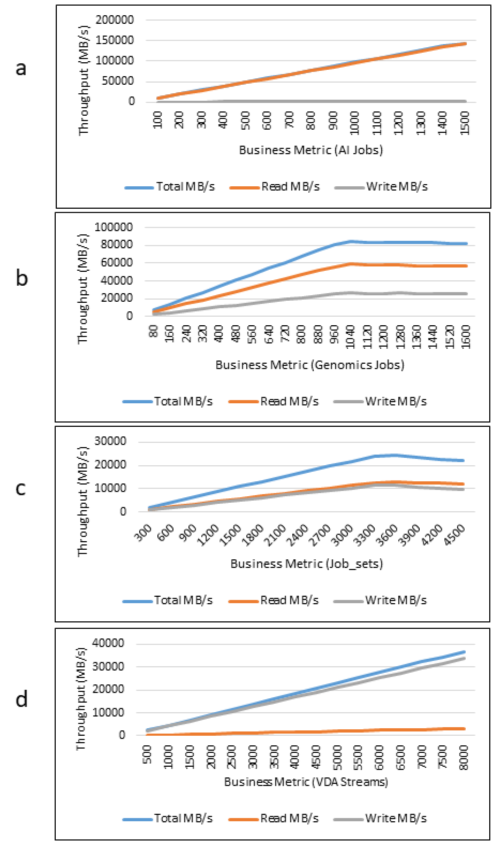

Figure 3: AI_Image, Genomics, EDA_Blended, and VDA SPEC Throughput results

Genomics: The average file size was 1613 KiB. The initial file set is equal to the start data set for all the business metrics. An incremental business metric of 80 was used and tests began to fail after business metric 1040. Looking at the latency curve in figure 2 b, we see that average latency is very low until business metric 960. At this point, the latency jumps up to above 6ms and benchmark tests fail. At business metric 1040, the benchmark log fails with the error message as in figure 6 a regarding operation rate threshold. When we observe the client side CPU utilization (figure 6 b), we see that CPU is bottlenecked at business metric of 1040 and afterwards. This is most likely the cause of the failing tests at that point. The Genomics workload reached a maximum throughput of 100 GB/s and achieved the maximum oprate of 1005K (figure 4 b) at a business metric of 1040 from which the tests started failing. A final business metric of 960 was achieved before the CPU became fully utilized.

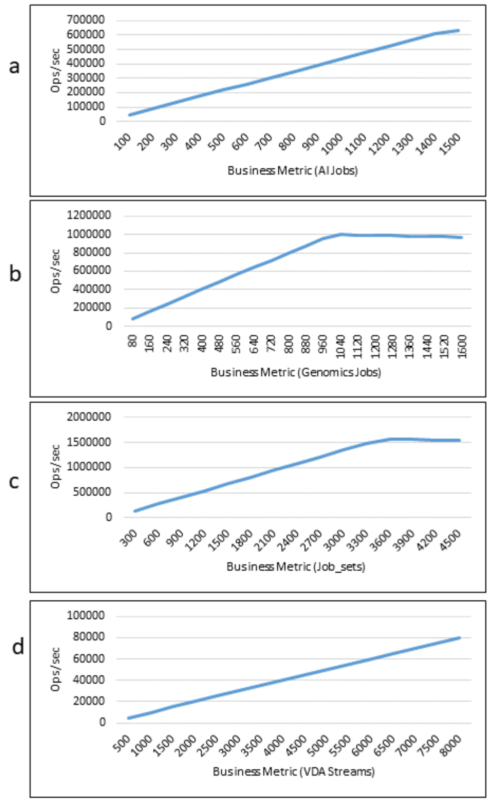

Figure 4: AI_Image, Genomics, EDA_Blended, and VDA SPEC IOPS results

EDA Blended: The average file size was 3424 KiB. The initial file set is equal to the start data set for all the business metrics. An incremental business metric of 300 was used and tests began to fail after a business metric of 3300. The average latency (figure 2 c) reported by the benchmark increases after 3600 business metric to over 10ms and continues up higher in all the failing tests. At business metric 3600, the benchmark test fails with the error about operation rate threshold as shown in figure 5 c. We see that CPU utilization (figure 5 d) on the clients begins to ramp up around the 3600 business metric and is nearly fully utilized on all the later tests. The EDA Blended workload did not see any improvement with the additional clients in the 8 client test. A final business metric of 3300 was achieved with an average throughput of 24 GB/s and the oprate of 1485K achieved on the storage before getting bottlenecked by client side CPU.

Video Data Acquisition: The average file size was 1048576 KiB. The initial file set is equal to the start data set for all the business metrics. An incremental business metric of 500, and a maximum business metric of 8000 was achieved. If we observe the latency results in figure 2 d across all the business metrics, we see only a small increase in the latency. The latency is well within the acceptable threshold and the environment should support higher business metrics if the storage nodes had higher capacity drives. The achieved throughput, and oprate are 37 GB/s (figure 3 d) and 80K (figure 4 d) respectively. The observed latency drops from a maximum of 3 ms to 2.7 ms (figure 2 d). The bottleneck preventing larger workloads is the storage capacity of the Weka cluster. The chart in figure 6 c shows the utilization % of the Weka cluster storage devices. You will see the capacity is increasing at a linear rate as business metric increases in subsequent tests until 100% utilization is reached. During the benchmarking, we do not observe any bottlenecks in CPU, memory, or network on either the clients or target nodes. In future work, we will add additional SSDs to the target nodes, in an attempt to reach higher business metrics.

4. Conclusion and Future Work

Samsung PM9A3 NVMe proved to be a highly performant SSD when integrated with Weka's software defined storage solution and put to the test with storage intensive workloads. This reference platform ranked number 1 in performance for 4 of the 5 SPECstorage Solution 2020 benchmarks. The final results for the workloads AI Image, EDA Blended, Genomics, and VDA are published on the SPEC website [5].

Limitations: The performance results for workloads EDA Blended, AI Image, and Genomics hit a client side CPU limitation. The backend storage servers (Weka) are capable of higher workloads if the clients were given more compute resources.

The VDA workload is limited to backend storage space. If more SSDs are provisioned to the Weka cluster, we believe we can test larger business metrics. Each Dell R7515 storage server supports attaching up to 24 NVMe devices; 54 more devices than the current configuration.

As a future work, we intend to conduct performance analysis with S3 and NFS protocols as services on top of the WekaFS file service and observe the impact of such configurations on the SPECstorage workloads.

References

[1] SPECstorage Solution 2020 Benchmark, "SPECstorage solution benchmark user guide," [Online]. Available: https://www.spec.org/storage2020/docs/usersguide.pdf, 2021.

[2] Weka File System Architecture, "WekaFS Architecture White Paper," [Online]. Available: https://www.weka.io/resources/wekaio-architectural-whitepaper/, 2022

[3] Dell EMC PowerEdge Server Specifications, "Dell EMC PowerEdge R7515 Rack Server," [Online]. Available: https://i.dell.com/sites/csdocuments/Product_Docs/en/poweredge-r7515-spec-sheet.pdf.

[4] Samsung Gen-4 NVMe SSDs, "Samsung PM9A3," [Online]. Available: https://www.samsung.com/semiconductor/global.semi.static/[S210809]_PM9A3_SSD_Whitepaper.pdf, 2021.

[5] SPECstorage 2020 benchmark results, "SPECstorage benchmark results," [Online]. Available: https://www.spec.org/storage2020/results/, 2022.

[6] WekaFS SSD capacity management, “Weka SSD capacity management,” [Online]. Available: https://docs.weka.io/overview/ssd-capacity-management#the-formula-for-calculating-ssd-net-storage-capacity, 2022.

[7] Dell EMC PowerEdge Server R7525 Rack Server, “Dell R7525,” [Online]. Available: https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/dell-emc-poweredge-r7525-spec-sheet.pdf, 2022.

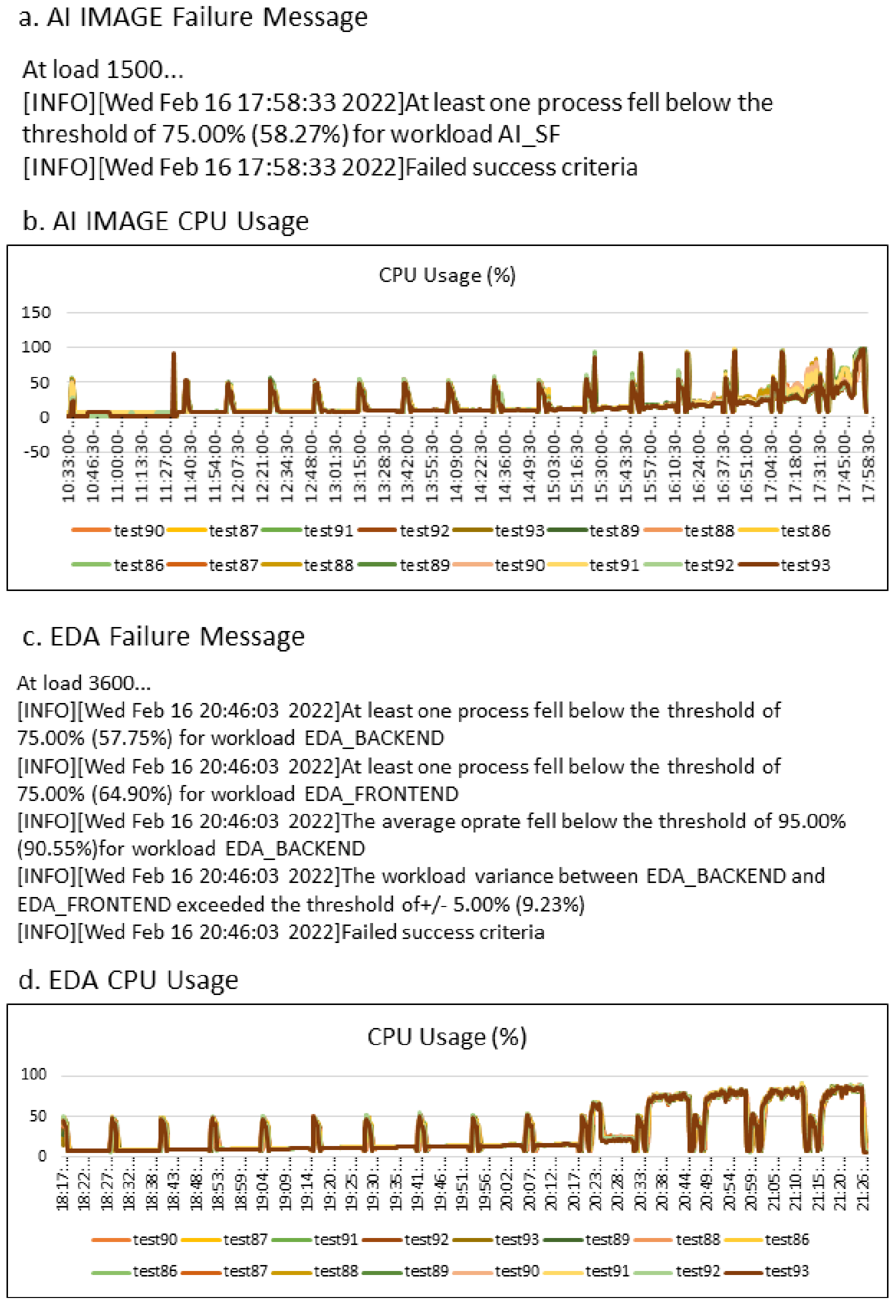

Figure 5: AI Image and EDA_Blended Workloads: Failure Message and CPU Usage

a. GENOMICS Failure Message

At load 1040...
[INFO][Thu Feb 17 02:38:00 2022]At least one process fell below the threshold of 75.00% (63.82%) for workload NGS
[INFO][Thu Feb 17 02:38:00 2022]Failed success criteria

b. GENOMICS CPU Usage

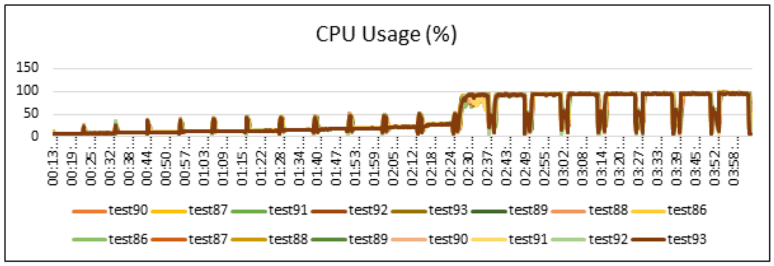

c. VDA Disk Usage

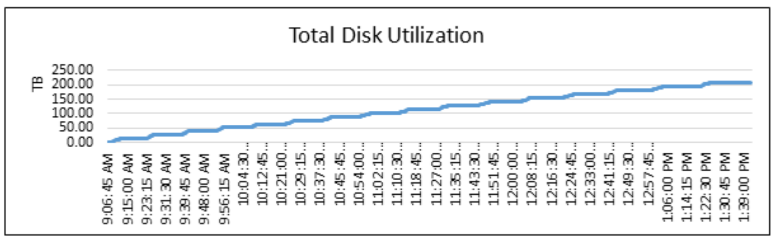

Figure 6: Genomics, and VDA: Failure Message, Disk and CPU Usage

Appendix

1. Workload Results

1.1. AI Image Spec Results

BM	ROR	AOR	Avg Lat (ms)	Total KBps	Read KBps	Write KBps	Cl Proc	CDS MiB	SDS MiB	MFS MiB
100	43500	43505.715	0.473	9782168.632	9642471.317	139697.315	212	1104125	8833000	9636000
200	87000	87011.437	0.48	19563861.53	19282042.96	281818.574	425	2208250	17666000	19272000
300	130500	130517.165	0.483	29338292.98	28916467.42	421825.558	637	3312375	26499000	28908000
400	174000	174022.868	0.486	39112621.71	38554204.66	558417.057	850	4416500	35332000	38544000
500	217500	217528.601	0.489	48898814.26	48195173.68	703640.58	1062	5520625	44165000	48180000
600	261000	261034.298	0.503	58673774.78	57833743.26	840031.517	1275	6624750	52998000	57816000
700	304500	304540.128	0.489	68453529.58	67472157.86	981371.722	1487	7728875	61831000	67452000
800	348000	348045.834	0.541	78243400.41	77117912.02	1125488.387	1700	8833000	70664000	77088000
900	391500	391551.542	0.632	88020984.12	86757684.99	1263299.13	1912	9937125	79497000	86724000
1000	435000	435057.161	0.727	97803375.75	96399024.92	1404350.824	2125	11041250	88330000	96360000
1100	478500	478563.037	0.864	107581854	106039488.4	1542365.653	2337	12145375	97163000	105996000
1200	522000	522068.519	1.068	117352330.4	115670611.3	1681719.13	2550	13249500	105996000	115632000
1300	565500	565574.367	1.398	127128326.9	125309398.6	1818928.324	2762	14353625	114829000	125268000
1400	609000	609005.301	3.635	136911977.2	134946719.9	1965257.338	2975	15457750	123662000	134904000
1500	652500	629039.066	18.349	144244658.2	142145111.6	2099546.632	3187	16561875	132495000	144540000

Table 1: AI Image SPEC results

2.2. Electronic Design Automation Spec Results

BM	ROR	AOR	Avg Lat (ms)	Total KBps	Read KBps	Write KBps	Cl Proc	CDS MiB	SDS MiB	MFS MiB
300	135000	135007.601	0.16	2178154.48	1144040.204	1034114.275	187	414433	3315468	3616875
600	270000	270015.132	0.161	4356604.182	2286706.392	2069897.789	375	828867	6630937	7233750

900	405000	405022.839	0.167	6535752.279	3431334.179	3104418.1	562	1243300	9946406	10850625
1200	540000	540030.689	0.177	8714324.103	4575211.009	4139113.095	750	1657734	13261875	14467500
1500	675000	675038.308	0.189	10894274.12	5718920.231	5175353.891	937	2072167	16577343	18084375
1800	810000	810045.866	0.205	13070278.21	6859865.729	6210412.48	1125	2486601	19892812	21701250
2100	945000	945053.559	0.223	15248380.98	8006149.499	7242231.481	1312	2901035	23208281	25318125
2400	1080000	1080059.967	0.247	17425711.08	9149159.525	8276551.55	1500	3315468	26523750	28935000
2700	1215000	1215069.284	0.277	19605730.32	10293041.01	9312689.303	1687	3729902	29839218	32551875
3000	1350000	1350075.187	0.324	21782436.48	11438576.01	10343860.47	1875	4144335	33154687	36168750
3300	1485000	1485083.744	1.657	23965243.89	12581348.53	11383895.36	2062	4558769	36470156	39785625
3600	1620000	1557147.802	11.143	24414445.88	12952178.26	11462267.62	2250	4973203	39785625	43402500
3900	1755000	1570177.614	13.075	23396375.43	12653345.58	10743029.85	2437	5387636	43101093	47019375
4200	1890000	1556294.406	14.803	22663905.96	12363174.41	10300731.55	2625	5802070	46416562	50636250
4500	2025000	1537560.95	16.439	21973854.37	12075076.85	9898777.528	2812	6216503	49732031	54253125

Table2: EDA SPEC results

2.3. Genomics Spec Results

BM	ROR	AOR	Avg Lat (ms)	Total KBps	Read KBps	Write KBps	Cl Proc	CDS MiB	SDS MiB	MFS MiB
80	80000	80002.29	0.211	6793139.495	4682121.157	2111018.338	40	34654	277234	302437
160	160000	160004.052	0.212	13588610.83	9368620.315	4219990.517	80	69308	554468	604875
240	240000	240006.332	0.221	20379318.25	14051945.37	6327372.877	120	103962	831703	907312
320	320000	320008.21	0.23	27179202.6	18736044.91	8443157.69	160	138617	1108937	1209750

400	400000	400010.895	0.239	33976736.61	23418064.96	10558671.65	200	173271	1386171	1512187
480	480000	480012.9	0.255	40771062.2	28104074.32	12666987.88	240	207925	1663406	1814625
560	560000	560015.107	0.268	47557271.19	32783928.2	14773342.99	280	242580	1940640	2117062
640	640000	640018.026	0.288	54362774.18	37471630.09	16891144.1	320	277234	2217875	2419500
720	720000	720020.9	0.32	61158105.89	42155692.12	19002413.77	360	311888	2495109	2721937
800	800000	800022.604	0.353	67945790.38	46832986.3	21112804.08	400	346542	2772343	3024375
880	880000	880025.436	0.417	74744018.85	51524377.82	23219641.03	440	381197	3049578	3326812
960	960000	960029.827	0.546	81531120.47	56203120.77	25327999.71	480	415851	3326812	3629250
1040	1040000	1004657.76	3.252	85322746.66	58818581.72	26504164.93	520	450505	3604046	3931687
1120	1120000	988675.553	4.379	83966555.69	57879502.52	26087053.17	560	485160	3881281	4234125
1200	1200000	987021.828	4.869	83816766.84	57770936.09	26045830.74	600	519814	4158515	4536562
1280	1280000	990097.288	5.2	84073228.64	57950947.77	26122280.87	640	554468	4435750	4839000
1360	1360000	979173.429	5.608	83151425.51	57309548.49	25841877.02	680	589123	4712984	5141437
1440	1440000	982335.609	5.944	83409353.45	57489383.98	25919969.47	720	623777	4990218	5443875
1520	1520000	973569.091	6.305	82661889.82	56966791.78	25695098.04	760	658431	5267453	5746312
1600	1600000	972900.919	6.663	82600179.28	56926618.57	25673560.71	800	693085	5544687	6048750

Table 3: Genomics SPEC results

2.4. VDA Spec Results

BM	ROR	AOR	Avg Lat (ms)	Total KBps	Read KBps	Write KBps	Cl Proc	CDS MiB	SDS MiB	MFS MiB
500	5000	5003.372	1.423	2303210.461	192929.207	2110281.254	125	1408000	11264000	12288000
1000	10000	10006.838	1.504	4601634.062	385999.197	4215634.865	250	2816000	22528000	24576000
1500	15000	15010.264	1.519	6908380.357	578341.532	6330038.825	375	4224000	33792000	36864000
2000	20000	20013.648	1.565	9214477.372	771203.875	8443273.497	500	5632000	45056000	49152000
2500	25000	25017	1.583	11510404.54	966031.114	10544373.43	625	7040000	56320000	61440000
3000	30000	30020.441	1.581	13823528.64	1160172.173	12663356.47	750	8448000	67584000	73728000
3500	35000	35023.801	1.612	16120545.92	1349982.093	14770563.83	875	9856000	78848000	86016000

4000	40000	40027.339	1.643	18423030.43	1544387.277	16878643.15	1000	11264000	90112000	98304000
4500	45000	45030.748	1.686	20737312.41	1738812.483	18998499.93	1125	12672000	101376000	110592000
5000	50000	50033.981	1.82	23038524.13	1930115.96	21108408.17	1250	14080000	112640000	122880000
5500	55000	55037.399	2.04	25332750.13	2121190.941	23211559.19	1375	15488000	123904000	135168000
6000	60000	60041.117	2.948	27636711.52	2315410.487	25321301.03	1500	16896000	135168000	147456000
6500	65000	65044.461	2.507	29932066.53	2510407.395	27421659.14	1625	18304000	146432000	159744000
7000	70000	70047.81	3.001	32257148.38	2703644.81	29553503.57	1750	19712000	157696000	172032000
7500	75000	75051.303	3.027	34547198.41	2896686.243	31650512.16	1875	21120000	168960000	184320000
8000	80000	80054.741	2.668	36846363.16	3089381.841	33756981.31	2000	22528000	180224000	196608000

Table 4: VDA SPEC results