intel + VAST

BREAKING SCALABILITY AND CONCURRENCY BARRIERS WITH VAST DATA

Benchmarking Guide to Maximize Client Performance (with NFS & S3)

JANUARY 2022

REVISION HISTORY

Revision Number	Description	Date
1.0	Initial Release	January 2022

KEY CONTACTS

Name	Title	Email
Chunhong Mao	Solutions Architect	chunhong.mao@intel.com
Frank Ober	Solutions Architect, Principal	frank.ober@intel.com
Allison Goodman	Director of IOG Solutions Architecture	allison.goodman@intel.com
Andy Pernsteiner	Field CTO, Manager	andy@vastdata.com
Sven Breuner	Field CTO	sven.breuner@vastdata.com
Rob Mallory	Field CTO	rob@vastdata.com

Executive Summary

VAST Data's Universal Storage is a next-generation, scale-out file and object storage solution that breaks decades of storage tradeoffs between performance, capacity, and cost. Through game-changing storage innovations that lower the cost of flash, VAST Data makes – for the very first time – flash affordable for all applications, from the highest performance datasets to the largest data archives.

To accomplish this monumental achievement, VAST Data pioneered a new scale-out storage architecture called DASE (Disaggregated Shared Everything). DASE is powered by three new technologies including low-cost hyperscale (QLC) flash, persistent storage class memory (such as high-performance Intel® Optane™ Technology),, and low-latency NVMe-over-Fabrics networking. Scalable to exabytes, this new architecture delivers linear performance scale by eliminating east-west communication across storage CPUs, while enabling high-speed random I/O access to billions of small files.

To demonstrate the scaling benefits of the DASE architecture, along with random I/O access benefits of an all-flash storage system, Intel and VAST Data jointly conducted several key benchmarking tests to simulate various degrees of concurrency and parallelism across 10's and 100's of simultaneous clients.

Highlights and key findings from the testing include:

- Near-perfect linear client performance scaling with 100GbE clients (Figure 15) and 10GbE clients (Figure 16).
- Achieved 117GB/s of read bandwidth (97% of theoretical max) with only 14 (100GbE) clients (Table 14).
- Sequential and random I/O access exhibit performance parity (+/- 2%) using same block size (4K) (Table 12).
- S3 read performance delivered 110GB/s (94% of NFS performance using same file/object size) (Table 13).

The remainder of this paper provides detailed system and network configurations used, multiple testing methodologies and tools, and step-by-step instructions to obtain the results outlined.

intel

Table of Contents

1) Hardware Environment

1.1 Block Diagram

The following diagram depicts the top level of logical connectivity between the 100 client cluster and the VAST Data 3x3 cluster.

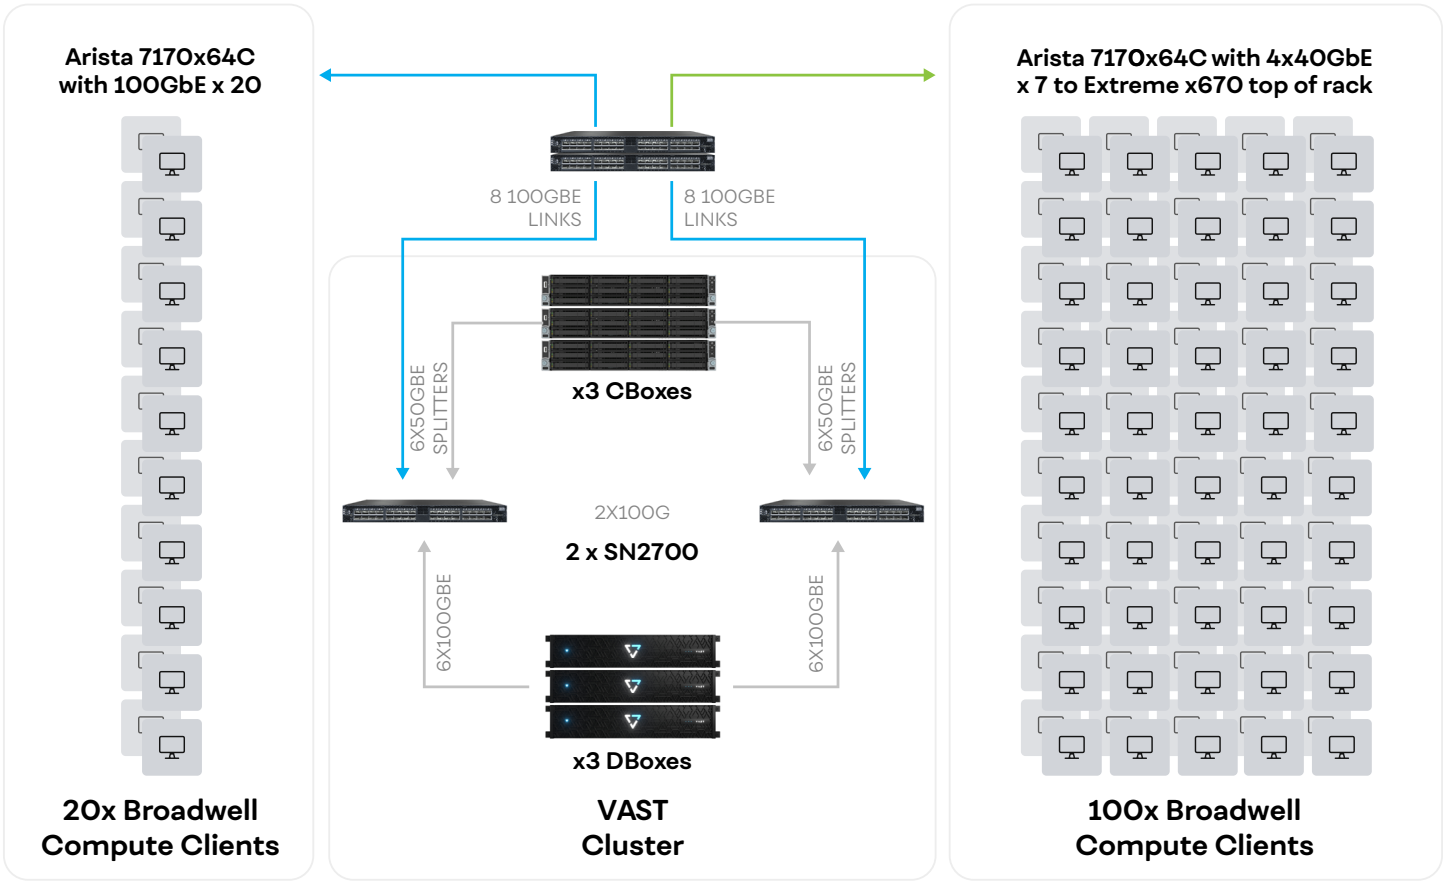

Figure 1: Block Diagram Logical View for Scale Testing

All 100 physical clients are outfitted with 10GbE network cards, while 20 of them have also been equipped with a second 100GbE NIC card.

A shell script-based test harness has been developed to build different test cases easily and execute the automated tests at scale. The test harness uses synthetic benchmarking tools FIO and Elbencho.

By using different subnets for the 10GbE and 100GbE networks, the client cluster can be configured to blast traffic to the VAST Data cluster in two different topologies via the test harness:

a) 100 clients with 10GbE networking

b) 20 clients with 100GbE networking

As both FIO and Elbencho will be running in Client/Server mode, one additional system is needed on the 10GbE network to execute the test harness.

Note that VIP pools created on the VAST Data cluster for 10Gb and 100Gb testing match the client 10Gb and 100Gb subnets respectively, but both need to be different to VAST Data's VMS. The table below shows an example of the IP assignments.

Table 1: 100GbE and 10GbE Subnets and VIP Assignments

Systems	Subnet	IP Examples
Client cluster 10GbE clients + head node	10.A.B.x	10.A.B.1~10.A.B.101
VAST Data VIP pool for 10G Testing	10.A.B.x	10.A.B.110~10.A.B.157
VAST Data VMS	10.A.C.y	10.A.C.D
Client cluster 100GbE clients	192.168.B.z	192.168.B.1~192.168.B.20
VAST Data VIP pool for 100GeE Testing	192.168.B.z	192.168.B.110~192.168.B.157

1.2 VAST Data 3x3 Cluster Hardware Configuration

The 3x3 VAST Data cluster is composed of:

1) 3 CBox enclosures, each enclosure contains:
   - A. a quad node chassis
   - B. 1 dual-port Mellanox MT27800 QSFP28 100GbE primary NIC per node for fast data path
   - C. 1 1GbE/10GbE NIC per node for management
2) 3 DBox enclosures, each enclosure contains:
   - A. a dual node chassis
   - B. 2 dual-port Mellanox MT27800 QSFP28 100GbE primary NIC per node for fast data path
   - C. 2 GbE NIC per node for management
3) 2 Mellanox SN2700 VAST Data switches, connecting the CBox enclosures and DBox enclosures

Each DBox provides 592TB of usable storage, for a total capacity of 1.775PB (before data reduction).

1.2.1 CBOX AND DBOX CONFIGURATIONS

Table 2: CBox and DBox Configurations

Component	CBOX Enclosure Configurations	DBOX Enclosure Configuration
Platform	Intel® Server Board S2600BPB (Cascade Lake)	Viking Enterprise Solutions (VES) NSS2560
# of Nodes	3x4	3x2
# of Sockets	2	2
CPU	Intel® Xeon® Silver 4216 CPU @ 2.10GHz	Intel® Xeon® CPU E5-2630 v4 @ 2.20GHz
Cores per socket / threads per socket	16/32	10/20
Microcode	0x5002f01	0xb000038
Intel® Hyper-Threading Technology	ON	ON
Intel® Turbo Boost Technology	OFF	ON
BIOS version	SE5C620.8 6B.02.01.0008.031920191559	11.03
System DDR memory configuration: slots/ capacity/speed	8 slots/32 GB/2400 MT/s (Micron) 8 slots/ No DRAM modules	8 slots/8 GB/2933 MT/s (Micron) 16 slots/ No DRAM modules
Total memory per node (All DRAM)	256G	64G
Storage	Boot: Intel SSD D3-S4510 960G	6 x Intel® Optane™ SSD P5800 1.6T 22 x Intel SSD D5-P4326 15.36T
Network Interface Card	1 x Mellanox® MT27800 ConnectX®-5 VPI Dual-Port 100 Gigabit Ethernet Adapter / Node 1 x Ethernet Controller 10G X550T / Node	2 x Mellanox® MT27800 ConnectX®-5 VPI Dual-Port 100 Gigabit Ethernet Adapter / Node 2 x Intel® 82574L Gigabit Ethernet / Node
NIC firmware	MT27800: 16.26.4012 (MT_0000000008) X550T: 0x80000a42, 1.1767.0	MT27800: 16.27.6008 (MT_0000000207-1) 82574L: 2.1-3

1.2.2 MELLANOX® SN2700 SWITCH

22 out of the 32-port 100GbE Mellanox® SN2700 switch are being allocated for data path connectivity.

Table 3: Mellanox® SN2700 Port Allocation

Number of Ports	Speed / Port	Connectivity
8	100GbE	Arista
6	100GbE	CBox (12 x 50GbE)
6	100GbE	DBox
2	100GbE	ISL (inter-switch link)

1.3 100 Clients With 10GbE Networking

This section outlines the hardware details used to set up the benchmarking environment of 100 servers with 10GbE networking in-between the client cluster and the VAST Data cluster.

1.3.1 NETWORK CONNECTIVITY

The diagram below depicts the network connectivity between client cluster and VAST Data cluster.

A. 100 clients are connected to 7 Extreme x760v switches, with up to 16 clients / per switch to ensure each client's 10GbE networking interface speed can be fully utilized.

B. Each of the 7 Extreme switches is connected to the Arista 7170-64C switch through a 4 x 40Gbit uplink.

C. The Arista 7170-64C is also connected to Mellanox SN2700 within the VAST Data cluster through 16 x 100Gbit links.

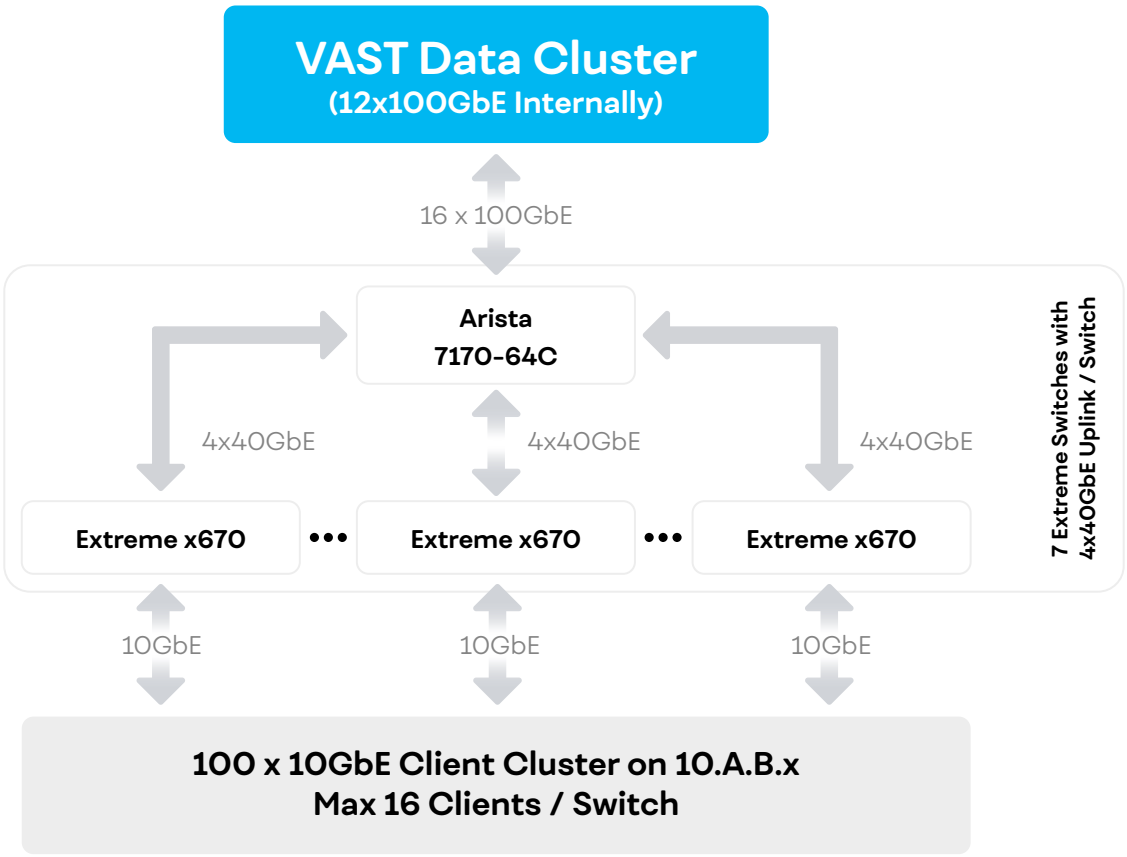

Figure 2: 100 x 10GbE Client Cluster and VAST Data Cluster Connectivity

Note that the head node is not shown in the diagram. Although the head node doesn't blast IO to VAST Data cluster itself, it does need to mount the VIPs as well for test harness utilities. So one additional 10GbE node is used for this purpose.

1.3.2 ARISTA 7170-64C SWITCH

One Arista 7170-64C switch is used as the hub in-between the VAST Data cluster and the client cluster. Please refer to https://www.arista.com/assets/data/pdf/Datasheets/7170-Datasheet.pdf for the datasheet specifications of the Arista 7170-64C switch.

In summary, it has:

A. 64 x 40/100GbE QSFP100 ports

B. 2 x 1GbE SFP+ ports for additional high speed management

Port allocations of the Arista 7170-64C is indicated in the table below:

Table 4: Arista® 7170-64C Port Allocation

Number of Ports	Speed / Port	Connectivity	Total BW Allowed at Arista 7170-64C
28	40GbE (5GB)	7 Extreme Switches	28 x 5 GB = 140 GB
20	100GbE (12.5GB)	20 clients with 100GbE	20 x 12.5 GB = 250 GB
16	100GbE (12.5GB)	3 x 3 VAST Data cluster	16 x 12.5 GB = 200 GB

1.3.3 EXTREME X670V SWITCH

Extreme x670v switch has:

A. 48 ports, 10Gbit / port

B. 4 QSFP+ ports, 40Gbit / port for uplinks

Here we use the 10Gbit ports to connect to client systems, and the QSFP+ ports to connect to the Arista switch.

- Max Bandwidth of the 100 10Gbit client systems: 100 x 10Gbit = 1000 Gbit = 125GB
- Max Bandwidth of a single Extreme Switch: 4 x 40Gbit = 160Gbit = 20GB

Important: although each Extreme x670v switch has 48 ports, the maximum ports that can be used to connect to client systems shall be no more than 16 because each Extreme x670 switch can only supply 160Gbit of uplink bandwidth using the 4 QSFP+ ports. As a result, 7 Extreme switches were required for this exercise.

1.4 20 Clients With 100GbE Networking

This section outlines the hardware details used to set up the benchmarking environment of 20 client systems with 100GbE networking in-between the clients and the VAST Data cluster.

1.4.1 NETWORK CONNECTIVITY

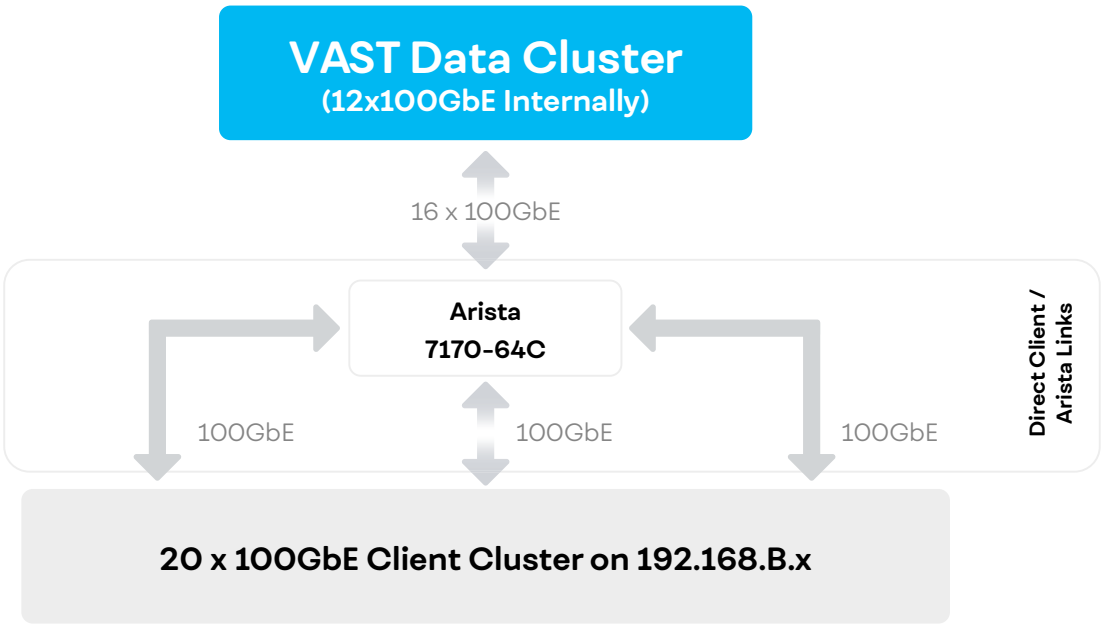

Figure 3: 20 x 100GbE Client Cluster and VAST DATA Cluster Connectivity

1.4.2 ARISTA® 7170-64C SWITCH

The 20 100GbE client systems are connected to Arista 7170-64C 20 100GbE ports directly.

1.4.3 EXTREME® X670V SWITCH

Unlike the 10GbE networking configuration, there is no Extreme x670v in-between the 20 100GbE clients and the Arista 7170-64C.

1.5 Client System Configurations

The table below shows the 10GbE client and 100GbE capable client system configurations. Note that the major difference between the two is simply that the 100GbE capable ones have dual NICs to support both 100GbE and 10GbE.

Table 5: Client System Configurations

Component	Hardware Configuration 100GbE Clients	Hardware Configuration 10GbE Clients
Platform	Intel® Server Board S2600WT2	Intel® Server Board S2600WT2
# of Nodes	20	101 (1 head node + 100 client nodes)
# of Sockets	2	2
CPU	Intel® Xeon® CPU E5-2699 v3 @ 2.30GHz	Intel® Xeon® CPU E5-2699 v3 @ 2.30GHz
Cores per socket / threads per socket	18/36 (72 per system)	18/36 (72 per system)
Microcode	0x46	0x44 / 0x46 (may have different versions)
Intel® Hyper-Threading Technology	ON in BIOS, configurable by test harness, default is OFF	ON in BIOS, configurable by test harness, default is OFF
Intel® Turbo Boost Technology	ON in BIOS, configurable by test harness, default is OFF	ON in BIOS, configurable by test harness, default is OFF
BIOS version	SE5C610.8 6B.01.01.0008.021120151325	SE5C610.8 6B.01.01.0008.021120151325
Local Storage used	None	None
System DDR memory configuration: slots/capacity/ speed	16 slots/8 GB/2133 MT/s (Micron) 8 slots/ No DRAM modules	16 slots/8 GB/2133 MT/s (Micron) 8 slots/ No DRAM modules
Total memory per node (All DRAM)	128 GB	128 GB
Storage	1x Intel® SSD DC S3500 800GB (boot device)	1x Intel® SSD DC S3500 800GB (boot device)
Network Interface Card	1x Intel® Ethernet Adapter E810-C 1 x Ethernet Controller 10-Gigabit X540-AT2	1 x Ethernet Controller 10-Gigabit X540-AT2
NIC firmware	E810-C: 2.15 0x800049c3 1.2789.0 X540-AT2: 0x800003f1	0x800003f1

1.6 Software Configurations

The table below lists the software configurations for VAST Data CNodes, DNodes and client systems.

Table 6: SW Configurations

Software / Benchmarking Tool	Version	Link
Client OS Distribution	CentOS Linux 8.4.2105	
Client Kernel	4.18.0-240.22.1.el8_3.x86_64	
VAST Data Appliance	release-3.6.0-sp7 build 465786	Proprietary, contact support@vastdata.com
VAST Data Cnode OS Distribution	CentOS Linux release 7.7.19088.2003 (Core)	
VAST Data CNode Kernel	3.10.0-1127.18.2.el7.vastos.9.x86_64	
VAST Data DNode OS Distribution	CentOS Linux release 7.8.2003 (Core)	
VAST Data DNode Kernel	3.10.0-1127.18.2.el7.VASTos.9.x86_64	
Multipath and NConnect Bundled Driver for CentOS 4.18.0.240	mlnx-nfsrdma-VAST_3.9.3.for.4.18.0.240.x.centos-kernel_4.18.0_240.22.1.el8_3.x86_64.rpm (built by VAST Data)	
irqbalance	1.3.0	
FIO	3.19	fio-3.19-3.el8.x86_64.rpm
Elbencho	Docker image id: 6ac5b651eaa9	Source Repo: https://github.com/breuner/elbencho Docker Registry: https://hub.docker.com/r/breuner/elbencho
Test Harness	a3bf9ac996062de	https://github.com/intel/scale-testing-for-vastdata.git

2) Testing Prerequisites and Setup

Prior to executing the test harness, the following steps must be performed on both the VAST Data cluster and the client systems:

1. Prepare the VAST Data cluster by creating the VIP Pools and S3 User and Access/Secret Keys for testing. The VIP ranges and S3 User and Access/Secret Keys created shall be added to env.conf.override.
2. Prepare the client cluster

   A. Clone the test harness repo onto the head node

   B. Create env.conf.override with the appropriate VIP pool IPs, client cluster host ranges, S3 User and Access/Secret Keys which are mandatory. It can also optionally contain any tunable configurations that are different to the default ones as in env.conf.

   C. Install dependencies

   D. Update MTU for the client cluster (VAST Data default is > 9000 already)

   E. Sanity check of the network connectivity

3. Run test with the test harness

This remainder of this chapter covers the preparation steps in further detail. Running tests with the test harness using FIO and Elbencho are covered in subsequent chapters.

Note that the focus of this guide is on the essential steps to run benchmarks against VAST Data storage at scale, general instructions of how to install VAST Data cluster and client cluster itself are out of scope for this guide.

2.1 Preparing the VAST Data Cluster

All benchmark tests rely on the interactions with the VIP pools on the VAST Data cluster. The easiest way to create a VIP pool is to use the VAST UI.

In our benchmark tests, 48 VIPs are used for both 10GbE and 100GbE.

2.1.1 CREATE VIP POOL 10.A.B.110-10.A.B.157 FOR 10GBE

Open VAST UI in a browser with VMS IP, e.g. 10.A.D.E, move the cursor to the left side bar, and click "Network Access" -> "Virtual IP Pool" -> "Create VIP Pool" on the upright corner to open the "Add Virtual IP Pool" window. Fill in Name / Start IP / End IP and Subnet CIDR, then click the "Create" button.

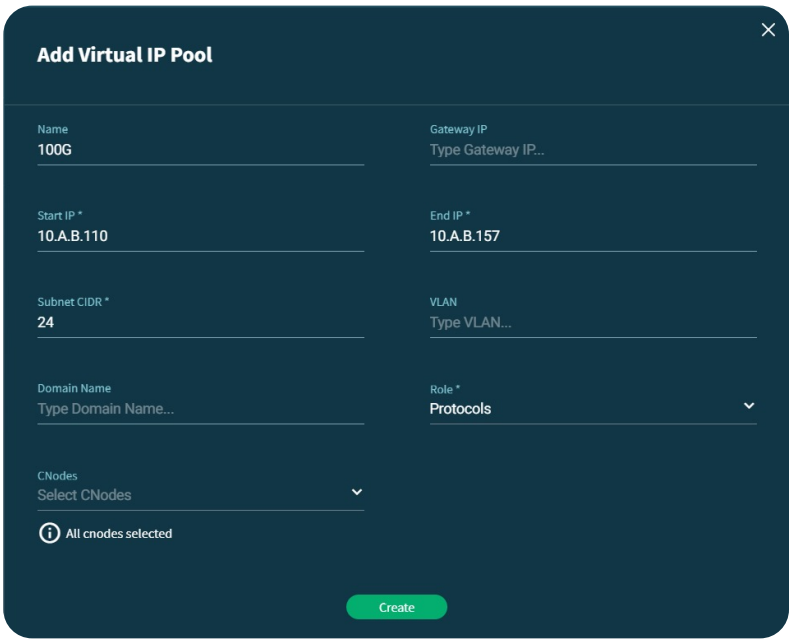

Figure 4: Create VIP Pool for 10GbE

2.1.2 CREATE VIP POOL 192.168.B.110-192.168.B.157 FOR 100GBE

Follow the same procedure to create the VIP pool for 100GbE.

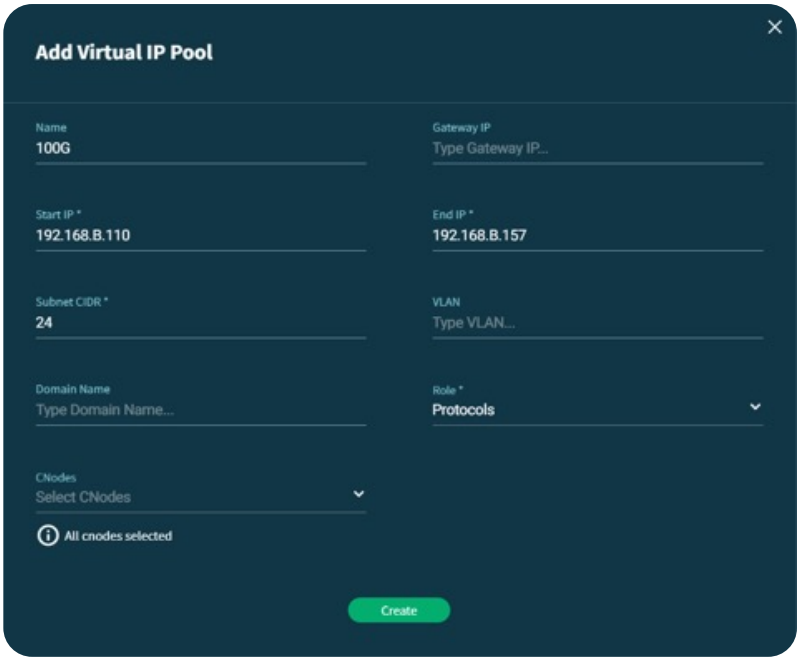

Figure 5: Create VIP Pool for 100GbE

2.1.3 CREATE USER FOR ELBENCHO S3 TEST

S3 test requires Access Key and Access Secret Key to be created for a user.

A. Open VAST UI in the browser with VMS IP, move the cursor to the left side bar, and click "User Management" -> "Create User" on the upright corner to open the "Add User" window. Fill in Name / UID, enable "Allow create bucket" and "Allow delete bucket", then click the "Create" button.

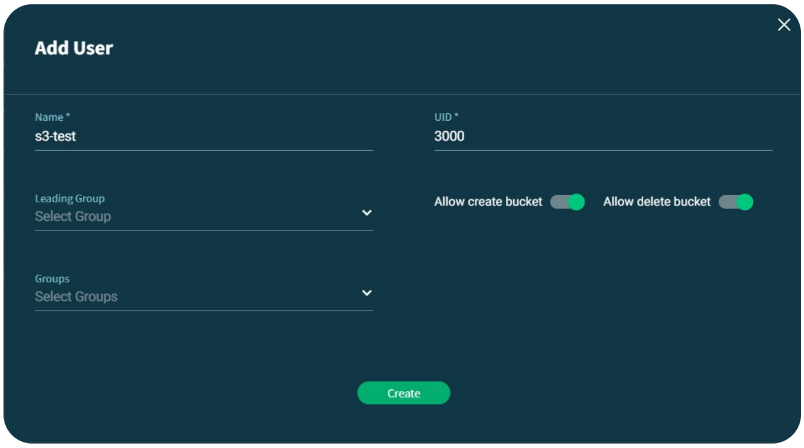

Figure 6: Create S3 User

B. Click the 3 dots below "Actions" in the row of the newly created user in the "User Management" window -> click "Edit" to open the "Update User" window.

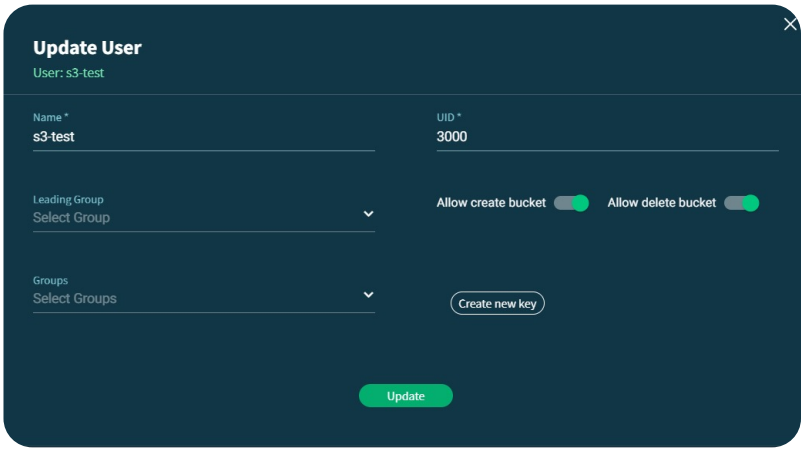

Figure 7: Update S3 User

C. Click "Create new key" button. Access Key and Secret Key will show on the screen. Click the icon in-between the Access Key and Status to copy the Secret Key. And click "Copy key" to copy the Secret Key. Click "Update" after both keys have been copied. Note that once closing the "Update User" window, the Secret Key is not visible anymore when clicking "Edit" again. However, "Create new key" can be used to add new key pairs.

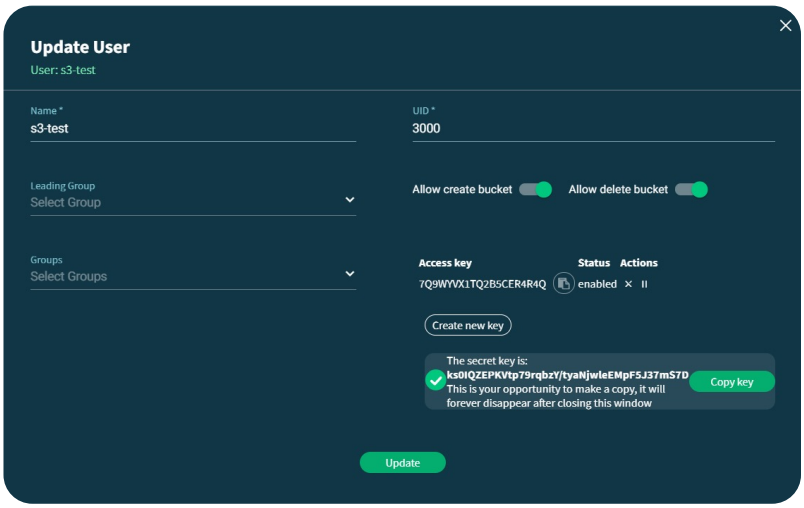

Figure 8: Create S3 Access and Secret Keys

2.2 Preparing the Client Cluster

2.2.1 CLONE THE TEST HARNESS REPO

Both FIO and Elbencho benchmarks are automated, test harness can be cloned from GitHub. Usually this is cloned to a head node which facilitates the test; however no workloads will be running in-between the head node itself and VAST Data cluster.

```
$ git clone https://github.com/intel/scale-testing-for-vastdata.git
```

Figure 9 below illustrates how the test harness interacts with Elbencho/FIO and ClusterShell to run tests against the cluster:

- VAST Data FS is exposed by VIPs mounted to the head (for utilities) and client nodes (for IO) via NFSv3 for file testing.
- VAST Data Objects are also exposed by VIPs for S3 object testing.
- Test Harness uses ClusterShell to configure clients, execute tests, and gather telemetry and results as needed.
- Both FIO and Elbencho run in Client/Server mode, with Client on the head node, and Server on client nodes.
- A separate remote server for result store can also be NFS mounted on the head node for result upload.
- $HOME/vast-scale-testing folder will be created on both the head and client nodes to facilitate the tests.
- A result folder will be created in repo top folder on the head node while the test is running.

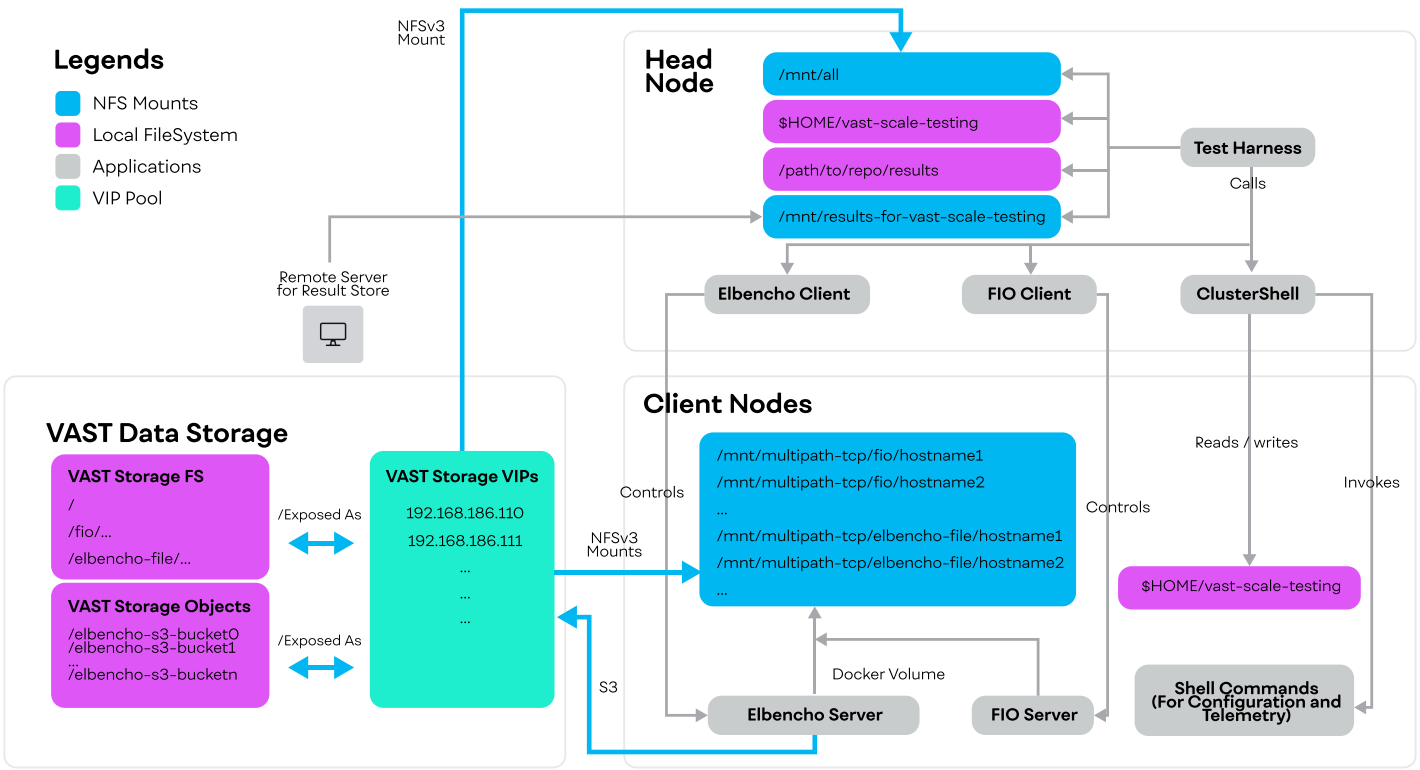

Figure 9: Block Diagram of How Test Harness Works

Figure 10 below depicts the shell script-based Test Harness composition on both the head node and client nodes.

- Repo content on the head node is via git clone.
- Some scripts will be deployed to client nodes via install_dep.sh for faster execution of test harness.
- During test, both the head node and client nodes will generate some folders, i.e. for mounts, results and test facilitation.
- Test Runner run_test_no_runlog.sh is the heart of the test harness
- The test launcher run_test.sh is a wrapper of the runner, and it invokes the runner by passing all command line options transparently.
- The other *.sh files are utilities
- env.conf / common.conf define the global variables needed by the launcher, runner and utility scripts
- common.conf is only directory aware; env.conf is topology aware
- There is also a validation component in the test harness

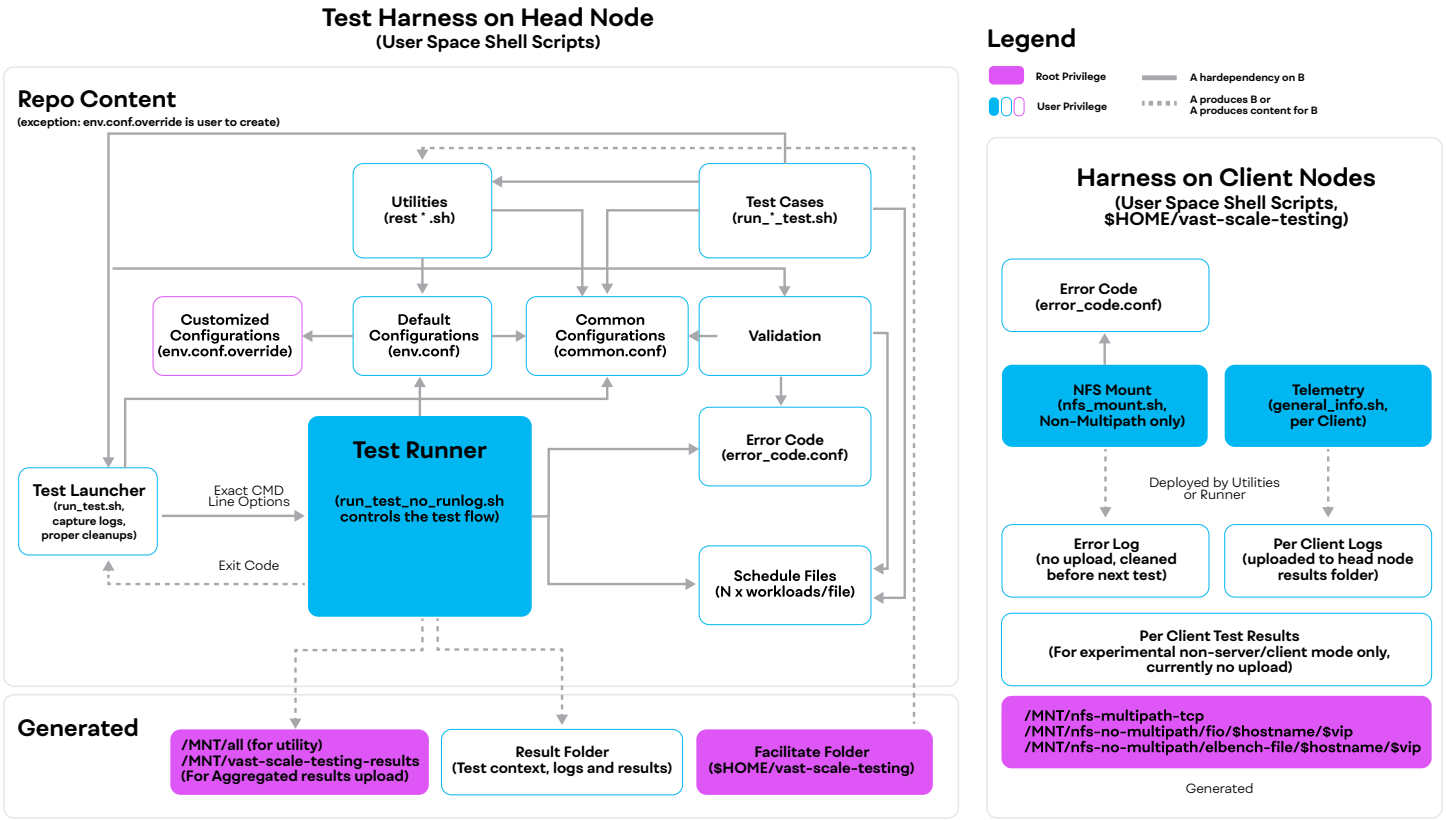

Figure 10: Block Diagram of Test Harness Composition

Figure 11 below shows the flow diagram of the runner (run_test_no_runlog.sh).

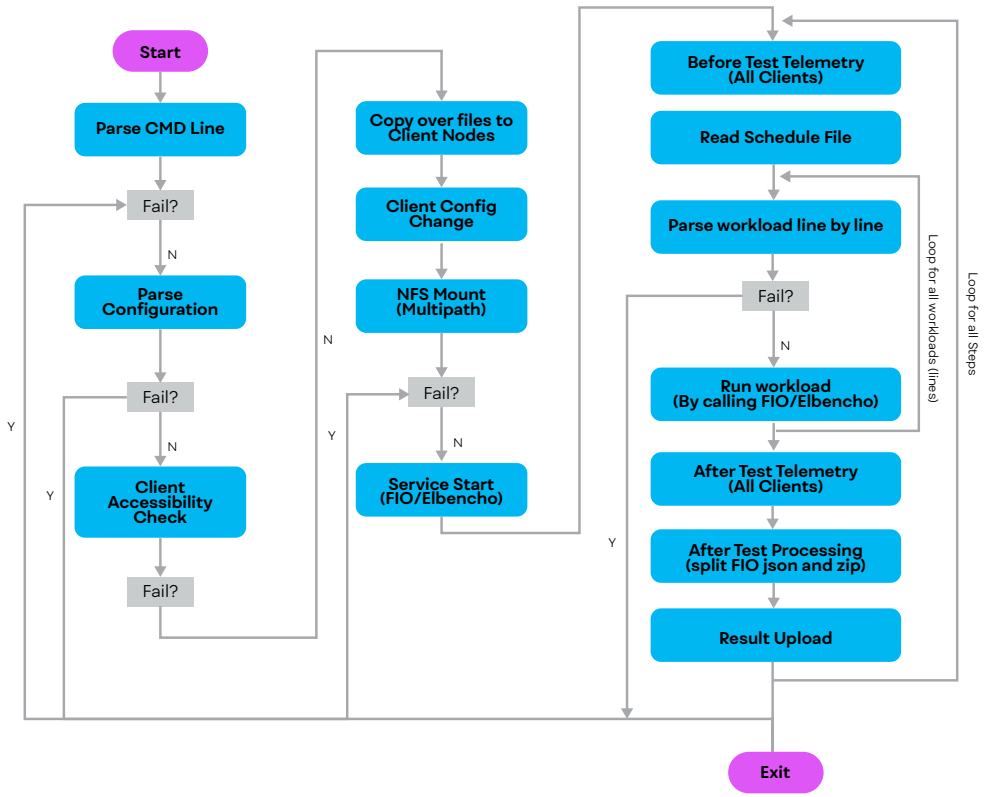

Figure 11: Runner Flow Diagram

There are 2 top level sub folders, namely scripts and workloads in the test harness after a fresh clone. A results folder also will be generated after running the test. A drivers folder can be created by the user in the top-level folder for storing Multipath & NConnect bundled driver required by the dependency installation script. The table below contains the high-level overview of all files.

Table 7: Overview of Test Harness Files

Folder Name	File	Description
Top Level	LICENSE	BSD 3-Clause License
	README.md	Top Level of Readme
	.gitignore	To ignore results folder and env.conf.override for source control
Workloads	fio/template.ini	It contains common fio parameters on which all FIO workloads will be based. Some options e.g. directory, numjobs, bs, iodepth, rw, create_only etc. are place holders only and might be updated by the harness after parsing command line or schedule files. The head node generates the per-client (directory option is different for each client) FIO job file, and supplies them to FIO as the –client parameter.

	fio/*.sch elbencho/file/*.sch elbencho/s3/*.sch fio/README.md elbencho/file/README.md elbencho/s3/README.md	A rich set of predefined schedule files for FIO/Elbencho File and Elbencho S3 testing respectively. Each schedule file defines 1 or more workloads, with 1 line of FIO/Elbencho File or Elbencho S3 parameters per workload. Refer to README.md in each folder to know which parameters are mandatory and which are optional. Most parameters are just native FIO/Elbencho File or Elbencho S3 parameters, with a few additions by the test harness itself. Refer to README.md for details. Suffix "10g" and "100g" indicates which interface each file is for. Parameter –size (FIO) and –file (Elbencho File and S3) for 100GbE are designed to be 5 times of 10GbE's. This is to keep the total data size across all clients of 10GbE and 100GbE testing the same, since there are 100 10GbE clients but only 20 100GbE clients. If the number of clients are different, then adjust the –size and –file accordingly for your own environment. A typical way of running a test would be from scripts folder: $ ./run_test.sh -ns 100 -sf path-to-schedule-file $ ./run_test.sh -ns 10 -sf path-to-schedule-file e.g.: $ ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_100g.sch $ ./run_test.sh -ns 10 -sf ../workloads/fio/bw_test_10g.sch Note that: 1. Schedule file suffix shall match the -ns parameter when running the test. If the suffix doesn't match, either not sufficient data will be generated, or data will be generated 5 times as designed for the scenario on write, thus takes much longer to run and also may not produce the result as designed by the test case. 2. Schedule files must be in these 3 folders for the test harness to detect the workload type, i.e. fio vs Elbencho file vs Elbencho S3 properly. Putting a schedule file outside of these 3 folders or putting them into the folder not matching its type are not supported by the test harness.
scripts	env.conf	Global configuration file. Used by most scripts that need to be topology (Clients and VIP) aware. **Usage:** $ source env.conf [network_speed step_size] **Where:** network_speed: in [10, 100]. Default: 10 step_size: in [0, CLIENT_COUNT]. Default: 0 (no incremental stepping, run on all clients in 1 step) This file is common for all users. For any user specific variables, please define them in env.conf.override instead. Dummy variables are provided in env.conf more as the references, and have to be redefined in env.conf.override. Define tunable ones in env.conf.override only if you want to configure them differently.
	env.conf.override	This is the user specific configuration file, which doesn't exist in fresh clones, but shall be created by the user. This file is ignored by .gitignore and won't be under source control. Refer to Create env.conf.override to know how to configure the variables in the file to suit the needs of a test.
	common.conf	Global configuration file. Used by scripts that only need to be folder structure aware, but not topology aware. **Usage:** $ source common.conf
	error_code.conf	Defines test harness generated error code.

	install_dep.sh	Run this to install the dependencies required by the test harness. **Usage:** $ ./install_dep.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10 Section Install the dependencies explains all dependencies to be installed.
	run_test_with_no_runlog.sh	Script called by run_test.sh to launch the test. Note that the script will do FIO server cleanups and NFS umounts before the test to ensure a clean run, however it will leave the FIO servers and NFS mounts in place after the test. **Usage:** $ ./run_test_with_no_runlog.sh [options] Refer to run_test.sh/run_test_no_runlog.sh command line options for all command line options.
	run_test.sh	Script to launch the test by calling run_test_with_no_runlog.sh and save the console output as run.log in the results folder at the end of the test. For better debuggability, also use this in the command line rather than run_test_with_no_runlog.sh directly. **Usage:** $ ./run_test.sh [options] Options are the same as run_test_with_no_runlog.sh as it just passes all arguments to run_test_with_no_runlog.sh transparently. Refer to run_test.sh/run_test_no_runlog.sh command line options for all command line options.
	split_json.sh	Script used by run_test_no_runlog.sh to split FIO json result per client. **Usage:** $ ./split_json.sh filename Filename must have ending _summary.log, and it must be the FIO json+ output format from FIO Client/Server mode.

	run_iperf_test.sh	Basic networking performance test with iperf3.
		Usage:
		Precondition: iperf3 installed on both client and VAST Data side (install_dep.sh installs that for client side only. If VAST Data side has a different version, then log on to VMS and do "clush -g cnodes sudo yum install iperf3" to install it first. Mismatching iperf, ie. iperf3 on client and iperf2 on VAST Data side won't work).
		Steps to test 10G performance:
		Step 1: log onto VAST VMS node, note down one of the VIP on the VMS node, e.g. 10.A.B.E
		Step 2: start iperf as the service on the VMS node
		$ iperf3 -s -p5900
		Step 3: run iperf test script on the client:
		$ ./run_iperf_test.sh 10 10.A.B.E
		Steps to test 100G performance:
		Step 1: log onto VAST VMS node, note down one of the VIP on the VMS node, e.g. 192.168.B.F
		Step 2: start iperf as the service on the VMS node, but bind iperf3 to the 100GbE interface IP
		$ iperf3 -s -p5900 -B 192.168.B.F
		Step 3: on client 101, run iperf test script:
		$ ./run_iperf_test.sh 100 192.168.B.F
		Refer to Sanity Check of the Network Connectivity with iperf3 for examples.

	mtu_update.sh	Script to update eth0 or specified network interface to 9000 MTU for better performance. The script will be deployed to all client nodes with install_dep.sh. **Usage:** $ ./mtu_update.sh [network_interface] **Where:** network_interface: the interface for which MTU to be changed. Default: eth0 Refer to MTU Update for more details.
	nic_info.sh	Script to query the NIC information. **Usage:** $ ./nic_info.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10
	cleanup_test.sh	Cleanup all processes launched by the test harness. This is for manual FIO server cleanups on all clients, if desired. **Usage:** $ ./cleanup_test.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10
	cleanup_mounts.sh	Cleanup all NFS mounts created by the test harness. This is for manual NFS mount cleanups on all clients, if desired. **Usage:** $ ./cleanup_mounts.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10
	purge_cluster.sh	Mount a VAST Data cluster VIP root to the head node at /mnt/all, and cleanup all test data using .VAST_trash. This assumes all data generated by the test harness are stored in following folders: 1. /mnt/all/$REMOTE_PATH_FIO 2. /mnt/all/$REMOTE_PATH_ELBENCHO_FILE 3. /mnt/all/$REMOTE_PATH_ELBENCHO_S3 4. /mnt/all/elbencho-s3-bucket* It doesn't clean up other data stored outside of these folders. Refer to Create env.conf.override for how the remote paths are defined for FIO / Elbencho File and Elbencho S3.
	nfs_mount.sh	Copied by install_dep.sh to all client nodes to mount VIPs individually, used by run_test_no_runlog.sh only when MULTIPATH_ENABLED is set to 0 in env.conf. **Usage:** $ ./nfs_mount.sh MOUNT REMOTEPATH HOST_SUBNET NCONNECT_NUM VIPs **Where:** MOUNT: Mountpoint parent folder. The actual mountpoint can be a subfolder below it. REMOTEPATH: remote path depending on whether it's FIO, Elbencho file or S3 testing. HOST_SUBNET: host subnet, eg 192.168.B or 10.A.B NCONNECT_NUM: nconnect number to be used VIPs: all VIPs to be mounted **Note:** if this script gets updated, then it must be re-deployed to all clients for run_test_no_runlog.sh to execute properly. Currently install_dep.sh does the deployment; run_test_no_runlog.sh also does the deployment before every test which can be opt out in the future.

	general_info.sh	Script used by run_test_no_runlog.sh to collect general system and hardware information for telemetry purpose before the test, logs will be saved to results/rundir/context folder, where rundir is auto generated based on date, schedule filename, the head node name and prefix etc. **Usage:** $ ./docker_start.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10

	elbencho_deploy.sh	Anonymous and authenticated docker users may run into docker pull limits quickly.
		This script provides a workaround to manually pull once from the head node and deploy the elbencho docker image to all systems.
		Usage:
		Step 1: pull once on the head node
		$ docker pull breuner/elbencho
		Step 2: note down "IMAGE ID" of the pulled image
		$ docker image ls
		Example response:
		REPOSITORY TAG IMAGE ID CREATED SIZE
		$ docker.io/breuner/elbencho latest 6ac5b651eaa9 4 days ago 136 MB
		Step 3: update ELBENCHO_ID and ELBENCHO_TAG accordingly in env.conf. override if different to default. Tag is user defined string to help identify the version in a more convenient way.
		Step 4: run this script to deploy accordingly
		$ ./elbencho_deploy.sh

	container_logs.sh	To query Elbencho container logs on client nodes and print to console. Useful for checking error conditions for abnormal execution. **Usage:** $ ./ container_logs.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10
	docker_start.sh	To start docker on all client systems. This can be useful if the system has not been configured to automatically start docker services after reboot. **Usage:** $ ./docker_start.sh [network_speed] **Where:** network_speed: in [10, 100]. Default: 10
	run_bw_test.sh	Three iterations of BW test with FIO, Elbencho file and Elbencho S3 on both 10GbE and 100GbE clients. **Usage:** $ ./run_bw_test.sh This is a wrapper script for running workloads.
	run_iops_test.sh	Three iterations of iops test with FIO, Elbencho file and Elbencho S3 on both 10GbE and 100GbE clients. **Usage:** $ ./run_iops_test.sh This is a wrapper script for running workloads.

	run_fio_bw_test_short.sh	Three iterations of short BW test with FIO, Elbencho file and Elbencho S3 on 100GbE clients. **Usage:** $ ./run_fio_bw_test_short.sh This is a wrapper script for running workloads.
	run_elbencho_s3_sweep_iosizes.sh	Three iterations of sweeping S3 iosizes test with Elbencho on 100GbE clients. **Usage:** $ ./run_elbencho_s3_sweep_iosizes.sh This is a wrapper script for running workloads.
	run_fio_sweep.sh	Worker, iodepth and iosize sweep on both 100GbE and 10GbE with FIO. **Usage:** $ ./run_fio_sweep.sh This is a wrapper script for running workloads.
	run_elbencho_file_sweep.sh	Threads, iodepth and iosize sweep on both 100GbE and 10GbE with Elbencho File. **Usage:** $ ./run_elbencho_file_sweep.sh This is a wrapper script for running workloads.
	run_elbencho_s3_sweep.sh	Threads and iosize sweep on both 100GbE and 10GbE with Elbencho S3. **Usage:** $ ./run_elbencho_s3_sweep.sh This is a wrapper script for running workloads.
results	Generated after running the test. Refer to Test Results Explained for how to understand and parse the results.	
drivers	Created by user to store kernel dependent Multipath driver (contact support@VASTdata.com for it)	

2.2.2 CREATE ENV.CONF.OVERRIDE

There are 3 types of variable definitions in this env.conf.

A. User defined dummy variables

These are environment dependent variables and updates are usually mandatory except for SHAREPOINT_HOST/SHAREPOINT_FOLDER, which depends on if RESULTS_CIFS_UPLOAD_ENABLE is 1 or not.

B. User defined tunable variables

These variables allow the test to be running in different ways, which may or may not require changes.

C. Derived variables

These are variables populated automatically, mostly by parsing the user defined variables and command line arguments. No changes required for these variables.

Usually users don't need to change env.conf, but do need to create a user specific env.conf.override by providing the actual values for dummy variables + any tunable variables that are different to the default values in env.conf. Here is an example of it:

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


VIP_POOL_10G="Your10GVIPRange"
VIP_POOL_100G="Your100GVIPRange"
CLIENT_NODES_10G="Your10GClientClusterRange"
CLIENT_NODES_100G="Your100GClientClusterRange"


ELBENCHO_S3KEY="ReplaceWithYouOwnKey"
ELBENCHO_S3SECRET="ReplaceWithYouOwnSecret"


FIO_PATH=/usr/local/bin
```

All user defined variables are listed below.

Table 8: User Defined Variables

Variable Types	Variables	Default	Description
Dummy	_VIP_POOL_10G _VIP_POOL_100G	10.A.B.110-10.A.B.157, 192.168.B.110- 192.168.B.157	The test harness assumes the VIPs are always contiguous. Follow Create VIP Pool 10.A.B.110-10.A.B.157 for 10GbE and Create VIP Pool 192.168.B.110-192.168.B.157 for 100GbE to create them and define the actual values in env.conf.override.
	_CLIENT_NODES_10G _CLIENT_NODES_100G	client-clus- ter[01-100], cli- ent-cluster[01-20]	Define the actual client cluster hostnames and range in env.conf.override. Client node IPs don't have to be contiguous, i.e. they can be client-cluster[05,10-20,50-55, ...] The number of client nodes can be any number >= 1. The lead node (FIO client) shall not be in the list. However, do keep in mind that if the number of clients are not 20 for 100GbE and 100 for 10GbE, then the size of data generated by using the predefined schedule files for write workloads would be different, hence may have implications to read performance later depending on where the data is. To deal with different number of hosts: • If total clients are indeed different: adjust the data size in schedule file accordingly, e.g. if there are only 20 10G, then possibly you can just use the 100g version of schedule files for all of the tests, or • If total clients are 20 / 100, but for read test you would like to test a subset of clients: write with the exact 20 / 100 clients, but afterwards reduce the number of read clients as needed for read workloads.
	ELBENCHO_S3KEY ELBENCHO_S3SECRET	ELBENCHO_ S3KEY="YU5... LQ1", ELBENCHO_ S3SECRET="aqb... EtnQ"	Needed for S3 testing. Follow Create User for Elbencho S3 Test to create them and define them in env.conf.override.
	SHAREPOINT_HOST SHAREPOINT_FOLDER	YourSharepointHo- stName, //${SHARE- POINT_HOST}/ YourRemotePathOn- SharepointHost	For result upload to NFS share on remote server. Add it to env.conf.override only if RESULTS_UPLOAD_ENABLE also is defined to 1 in env.conf.override. Refer to the description of RESULTS_UPDATE_ENABLE variable for how to set it up
Tunable	TURBO_BOOST_ENABLE	0	Intel Turbo Boost control. 0: to disable it 1: to enable it 1 only works only if BIOS also has it enabled.
	HYPER_THREADING_ ENABLE	0	Hyper Threading control 0: to disable it 1: to enable it 1 only works only if BIOS also has it enabled.
	IRQBALANCE_ENABLE	0	irqbalance control. 0: to disable it 1: to enable it This requires irqbalance to be installed (covered by in-stall-dep.sh).

Tunable	ORDERED_NODES	1	Nodes can be defined as ordered or unordered. Normally ordered types can be used. However, sometimes it is useful to run in the specified order other than the sorted order, e.g. for incremental stepping test, if you would like to iterate through the faster clients first, then the slower clients. 1. Ordered nodes **CAN** be defined in 2 different formats of: a) node[01-02,30-32,10-12], or b) node{01..02} node{30..32} node{10..12} Nodes will be expanded to "node01 node02 node10 node11 node12 node30 node31 node32" with nodeset --expand 2. Unordered nodes **SHALL** be defined in format of: node{01..02} node{30..32} node{10..12} Nodes will be expanded to "node01 node02 node30 node31 node32 node10 node11 node12" with eval only.
	FIO_SERVER_CLIENT_MODE	1	To enable/disable FIO Server/Client mode. **1: to enable FIO Server/Client mode.** Aggregation: results are aggregated in fio json+ output as "All clients" Where: head node "results" folder **0: to disable FIO Server/Client mode.** Aggregation: no result aggregation of all clients, json+ output is per client. Where: in ~/$VAST_SCALE_TESTING_PATH folder For better total BW/IOPS/Latency tracking, use Server/Client mode.
	MULTIPATH_ENABLE	1	Kernel dependent Multipath driver is needed in order to enable this. Contact support@vastdata.com for it if needed. Enabling Multipath allows simpler VIP mount on client nodes and more balanced load on VAST Data CNodes and DNodes, especially when worker/thread count is low. When the variable is set to 0, VIPs are going to be mounted individually. Note that workload files could be created differently on remote nodes for MULTIPATH_ENABLE=0 and 1, hence for the same read workload, by just changing this variable, it may cause the file to be regenerated.
	MOUNT_ALL	1	Applicable to FIO and MULTIPATH_ENABLE=0 only. 1: One mount per VIP per client node. Overall mounts per client = VIP number. Mounts are created before any workloads start. This is a preferred way of mounting, as it's faster and has no noticeable performance penalty due to over mount, i.e. if mount number > job/worker number. 0: One mount per job/thread per client node. Overall mounts per client = job/thread number. Mounts are created on the fly when running workload. This could be slower.

	MOUNT_STAGGER	5	This is only applicable to MULTIPATH_ENABLE=0 and MOUNT_ALL=0 . This is to avoid creating too many clush sessions when doing individual mounts on the fly at one time.
	MOUNT_MULTIPATH_ENABLED	/mnt/nfs-multipath-tcp	The mountpoint on the client system when Multipath is enabled.
	MOUNT_MULTIPATH_DISABLED	/mnt/nfs-no-multipath	The mountpoint on the client system when Multipath is disabled.
	NCONNECT_NUM	48	NConnect allows multiple TCP connections to be created for each mount, hence improves the throughput effectively. This is independent of Multipath and can be used together with Multipath. NConnect is supported from Linux kernel 5.3. For kernel 4.18, CentOS also backported it. The Multipath driver built by VAST Data can be a bundle of both. NCONNECT_ENABLE must be set to 1 for this to take effect.
	NCONNECT_ENABLE	1	When MULTIPATH_ENABLE=0, and NCONNECT_ENABLE=1, big NCONNECT_NUM means a large number of TCPs will be created per client as the table below shows. This could (1) take a long time to mount before a test can start (2) cause connection aborts if too many TCP contentions are created. Best practice is to use: MULTIPATH_ENABLE=1 NCONNECT_ENABLE=1 Or: MULTIPATH_ENABLE=0 NCONNECT_ENABLE=0
	NCONNECT_NUM	48	NConnect allows multiple TCP connections to be created for each mount, hence improves the throughput effectively. This is independent of Multipath and can be used together with Multipath. NConnect is supported from Linux kernel 5.3. For kernel 4.18, CentOS also backported it. The Multipath driver built by VAST Data can be a bundle of both. NCONNECT_ENABLE must be set to 1 for this to take effect.

MULTIPATH_ENABLED	NCONNECT_ENABLE	Per Client TCPs
1	1	NCONNECT_NUM
1	0	1
0	0	VIP Count
0	1	VIP Count * NCONNECT_NUM

Table 9: Per Client TCP Connections with MOUNT_ALL=1

	REMOTE_PATH_FIO	fio	Folder to be created under VIP root for FIO testing
	REMOTE_PATH_ ELBENCHO_FILE	elbencho-file	Folder to be created under VIP root for Elbencho File testing
	REMOTE_PATH_ ELBENCHO_S3	elbencho-s3	Folder to be created under VIP root for Elbencho S3 Object testing
	FIO_PATH	/usr/local/bin	If FIO is installed in /usr/bin or other path, change FIO_PATH accordingly
	VAST_SCALE_TESTING_ PATH	~/VAST-scale-testing	VAST-scale-testing will be created under home to facilitate the installations and tests. For FIO non-Server/Client mode, per client result folders will also be created there.
	ELBENCHO_ID	6ac5b651eaa9	Use "docker image ls" to find it on the lead node after the docker pull.
	ELBENCHO_TAG	scaletest01	Change the tag after a new deployment to help identify the image.
	ELBENCHO_IMG	elbencho.$ELBEN-CHO_TAG	Full image identifier by tag.
	ELBENCHO_IMAGE_ FOLDER	$VAST_SCALE_TEST-ING_PATH/elbencho	This defines the folder to store the elbencho image on head and client nodes.
	RESULTS_UPDATE_ENABLE	0	Disabled by default. To use to: 1. Ensure CIFSNFS share for result upload is configuredsetup on a remote server and working from the head node 2. UpdateDefine SHAREPOINT_HOST and SHAREPOINT_ FOLDER accordingly in env.conf.override 3. Create /root/.cifs with the username and password for CIFS username=your_cifs_username password=your_cifs_password 4. ChangeDefine RESULTS_CIFS_UPDATE_ENABLE assigned valueto 1 in env.conf.override
	RESULTS_MNTPOINT	/mnt/VAST-scale-testing-results	Mountpoint for results upload on head node. The test harness automatically mount the sharepoint, upload the results to sharepoint, and then umount the sharepoint before ending the test.
	INTER_WORKLOAD_WAIT_ IN_SECONDS	60	Sleep time in-between workloads in seconds.
	INTER_STEP_WAIT_IN_ SECONDS	120	Sleep time in-between tests in seconds.
	CLUSH_MAX_FAN_OUT	120	Fan out for ClusterShell comands CentOS default 64 is not optimal for scale testing with 100 client nodes, as nodes beyond the threshold will be staggered until the previous batch is done, causing undesired synchronization issues.

2.2.3 INSTALL THE DEPENDENCIES

S Prior to installing the dependencies, please check your kernel version with "uname -r" and acquire the Multipath + NConnect buddled driver matching your kernel version from support@vastdata.com and add MULTIPATH_DRIVER to env.conf.override. After that, the driver and all other dependencies can be installed by calling install-dep.sh.

ince 100GbE hosts are just a subset of 10GbE hosts, installing dependencies on 10GbE shall be good enough.

```
$ ./install-dep.sh
```

However, for an environment using different 100GbE and 10GbE clients, make sure both are done.

```
$ ./install-dep.sh 10
$ ./install-dep.sh 100
```

A manual reboot of all clients are needed to activate the Multipath + NConnect bundled driver before running the test. The reboot is commented out in case another convenient time is preferred for reboot.

```
$ source env.conf 10
$ clush -w $CLIENT_NODES_10G, $CLIENT_NODES_100G "sudo reboot"
$ sudo reboot  # To reboot head node too
```

After reboot, a mount with Nconnect and Multipath of the VIPs should be successful. For example, the following command if executed on 10.A.B.5 10GbE client will mount all 48 VIPs as remote ports to local port 10.A.B.5, and it's using 10.A.B.123:/ as the remote address to mount.

```
$ sudo mkdir -p /mnt/nfs-multipath-tcp
$ sudo mount -v -o vers=3,proto=tcp,nconnect=48,port=20048,localports=10.A.B.5,
emoteports=10.A.B.110-10.A.B.157 10.A.B.123:/ /mnt/nfs-multipath-tcp     ===>
here both nconnect and multipath used
$ mount
…
10.A.B.123:/ on /mnt/nfs-multipath-tcp type nfs (rw,relatime,vers=3,rsize=1048576,
wsize=1048576,namlen=255,hard,proto=tcp,nconnect=48,port=20048,timeo=600,retrans=2,
sec=sys,mountaddr=10.A.B.123,mountvers=3,mountport=20048,mountproto=tcp,
local_lock=none,addr=10.A.B.C)    ===>
where C can be any number within 110~157
$ sudo umount /mnt/nfs-multipath-tcp  ===> this will umount the VIPs
```

When Multipath is enabled, the test harness always uses the full range of VIPs defined in env.conf.override for remote ports. The remote address however will be round robinned among all clients, e.g. the 1st client uses 10.A.B.110, 2nd client uses 10.A.B.111 etc.

All these however will be managed by the test harness and is transparent to the user.

If the client systems don't restart docker automatically after reboot, also do the following to start docker.

```
$ source env.conf 10
$ ./docker_start.sh 10
```

For environments that use different 100GbE and 10GbE clients, make sure the following is also done.

```
$ source env.conf 100
$ ./docker_start.sh 100
```

If docker is not started on a client, then the entire test with Elbencho will be aborted, and there might be error like below in console or run.log.

```
$ ....Communication error in preparation phase: Connection refused. Service:
hostname:1611
```

By default, install-dep.sh only pulls Elbencho to the head node, the deployment of Elbencho to all other nodes needs to be done either by using elbencho_deploy.sh after manual steps 2 and 3 as mentioned in its comments (suitable for users with docker pull limt):

```
$./elbencho_deploy.sh
```

Or by a direct pull manually (suitable for users without docker pull limit), eg:

```
$ source env.conf 10
$ clush -w "$CLIENT_NODES" "docker pull breuner/elbencho"
```

In this case, refer to steps 2 and 3 for how to get ELBENCHO_ID and ELBENCHO_TAG added in env.conf. override. then:

```
$ source env.conf 10
$ clush -w "$CLIENT_NODES" "docker image tag $ELBENCHO_ID breuner/elbencho:$ELBENCHO_TAG"
```

Elbencho and FIO services will be started by the test harness automatically on all client nodes, no need to manually do it before the test.

```
$ ps aux | grep fio
bduser     67771  0.0  0.0 440076  5316 ?        Ss   Nov18   0:00 /usr/bin/fio
--server --daemonize=fiopid
$ docker ps
CONTAINER ID   IMAGE                            COMMAND                  CREATED
STATUS          PORTS     NAMES
d099ebdffb8f   breuner/elbencho:scaletest01   "/usr/bin/elbencho -…"   50 seconds ago
Up 50 seconds             elbencho-server
```

The services won't be cleaned up automatically after the test though, but manual cleanup is possible if desired.

```
$ ./cleanup_test.sh 10
$ ./cleanup_test.sh 100  # if 100GbE is not a subset of 10GbE hosts
```

Table 9 lists all the dependencies to be installed and why each dependency is needed.

Table 9: Dependencies to be Installed

Name	Description
ClusterShell	The harness uses clush command extensively, so that all operations can be issued from the shell of the head node without logging on to other systems.
Multipath & NConnect Driver	Refer to descriptions of MULTIPATH_ENABLE and NCONNECT_NUM in Table 8: User Defined Variables for why Multipath and NConnect are needed. This shall be installed on all client systems running FIO server or Elbencho as the service.
FIO	Server/Client mode is used to run FIO, for easier result aggregation. To understand FIO Server/Client mode more, please refer to https://linux.die.net/man/1/fio section CLIENT/SERVER. In summary: FIO Client: this is the lead node which will facilitate the FIO execution on FIO Servers, it doesn't generate the IOs to the VAST Data cluster storage. FIO Server: the 100 10GbE clients/20 100GbE clients that are performing read/write from/to the VAST Data cluster storage through NFSv3 mounts. As such, on top of the 100 client systems, another system is needed to be used as the FIO client lead node. The test harness repo is only required to be cloned on the lead node, however FIO shall be installed on all nodes including the lead node and clients.
jq	A json parser used by the harness to split FIO json results to per client files. This could make it easier to further process the json results by some tools i.e. ElasticSearch, if the smaller size of per client json result is desired.
iperf3	Utility to help with basic networking performance tests.

elbencho	Elbencho is a distributed storage benchmark for file systems, object stores & block devices with support for GPUs, developed by VAST Data.
	Source repo: https://github.com/breuner/elbencho
	Docker container registry: https://hub.docker.com/r/breuner/elbencho
	Script install-dep.sh pulls the docker container with following command to the head node only (assume docker has been installed on the system).
	$ docker pull breuner/elbencho
	The deployment of Elbencho to client nodes is through elbencho_deploy.sh. For users that don't have a pulling limit, docker pull can be used directly on all client nodes.
	To know how to use Elbencho, refer to help, eg:
	$ docker run --net=host -it breuner/elbencho –help
	$ docker run --net=host -it breuner/elbencho –help-s3

nfs_mount.sh	A NFS mount helper script from the harness to be deployed to all client nodes, not a true external dependency.
mtu_update.sh	A script from the test harness to be deployed to all client nodes for more convenient MTU update, not a true external dependency. Refer to MTU Update for more details.
general_info.sh	A script from the test harness to be deployed to all client nodes to collect system telemetry as the running context, not a true external dependency.

2.2.4 MTU UPDATE

For best performance, MTU 9000 is recommended for both the 10GbE and 10GbE interfaces. This can be done by following steps, supposing the network interface is eth0 for 10GbE and eth1 for 100GbE.

Important:

1. Please adapt to the proper interfaces if they are different
2. Ensure switch side connected to this interface supports 9000 too, otherwise the system may not be reachable after this update

```
$ source env.conf  # Needed by variable $VAST_SCALE_TESTING_PATH
$ clush -w $ CLIENT_NODES_10G "~/ $VAST_SCALE_TESTING_PATH/mtu_update.sh”
$ clush -w $ CLIENT_NODES_10G “cat /etc/sysconfig/network-scripts/ifcfg-eth0”
$ clush -w $ CLIENT_NODES_100G "~/ $VAST_SCALE_TESTING_PATH/mtu_update.sh eth1”
$ clush -w $ CLIENT_NODES_100G “cat /etc/sysconfig/network-scripts/ifcfg-eth1”
```

2.3 Sanity Check of the Network Connectivity with iperf3

Before running any workloads, it is recommended to verify the proper network connectivity and speed using iperf3.

2.3.1 CHECK THE 10GBE CONNECTIVITY

Step 1: log onto VAST VMS node, note down one of the VIPs on the VMS node, eg 10.A.B.E

Step 2: start iperf3 as the service on the VMS node

```
$ iperf3 -s -p5900
```

Step 3: run iperf3 test script on the head node:

```
$ ./run_iperf_test.sh 10 10.A.B.E
```

This will iterate through all the 10GbE clients as defined in env.conf.override and run iperf3 test one by one using clush.

```
client-cluster-02: Connecting to host 10.A.B.E, port 5900
client-cluster-02: [  5] local 10.A.B.2 port 43376 connected to 10.A.B.E port 5900
```

```
client-cluster-02: [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
client-cluster-02: [  5]   0.00-1.00   sec  1.15 GBytes  9.89 Gbits/sec    0   1.94 MBytes
client-cluster-02: [  5]   1.00-2.00   sec  1.15 GBytes  9.89 Gbits/sec    0   1.94 MBytes
client-cluster-02: [  5]   2.00-3.00   sec  1.15 GBytes  9.91 Gbits/sec    0   2.07 MBytes
client-cluster-02: [  5]   3.00-4.00   sec  1.15 GBytes  9.90 Gbits/sec    0   2.07 MBytes
client-cluster-02: [  5]   4.00-5.00   sec  1.15 GBytes  9.90 Gbits/sec    0   2.07 MBytes
client-cluster-02: - - - - - - - - - - - - - - - - - - - - - - - - -
client-cluster-02: [ ID] Interval           Transfer     Bitrate         Retr
client-cluster-02: [  5]   0.00-5.00   sec  5.76 GBytes  9.90 Gbits/sec    0
sender
client-cluster-02: [  5]   0.00-5.00   sec  5.76 GBytes  9.89 Gbits/sec
receiver
client-cluster-02:
client-cluster-02: iperf Done.
client-cluster-03: Connecting to host 10.A.B.E, port 5900
client-cluster-03: [  5] local 10.A.B.3 port 45376 connected to 10.A.B.E port 5900
client-cluster-03: [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
client-cluster-03: [  5]   0.00-1.00   sec  1.15 GBytes  9.87 Gbits/sec    0   2.15 MBytes
client-cluster-03: [  5]   1.00-2.00   sec  1.14 GBytes  9.76 Gbits/sec    0   2.15 MBytes
client-cluster-03: [  5]   2.00-3.00   sec   975 MBytes  8.18 Gbits/sec    0   2.81 MBytes
client-cluster-03: [  5]   3.00-4.00   sec  1.03 GBytes  8.84 Gbits/sec    0   2.81 MBytes
client-cluster-03: [  5]   4.00-5.00   sec  1.15 GBytes  9.90 Gbits/sec    0   2.81 MBytes
client-cluster-03: - - - - - - - - - - - - - - - - - - - - - - - - -
```

2.3.2 CHECK THE 100GBE CONNECTIVITY

Step 1: log onto VAST VMS node, note down one of the VIP on the VMS node, eg 192.168.B.F

Step 2: start iperf3 as the service on the VMS node, but bind iperf3 to the 100GbE interface IP

```
$ iperf3 -s -p5900 -B 192.168.B.F
```

Step 3: run iperf3 test script on the head node:

```
$ ./run_iperf_test.sh 100 192.168.B.F
```

This will iterate through all the 100GbE clients as defined in env.conf.override and run iperf test 1 by 1 using clush.

```
client-cluster-05: Connecting to host 192.168.B.F, port 5900
client-cluster-05: [  5] local 192.168.D.5 port 55085 connected to 192.168.B.F port 5900
client-cluster-05: [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
client-cluster-05: [  5]   0.00-1.00   sec  1.88 GBytes  16.2 Gbits/sec    0    673 KBytes
client-cluster-05: [  5]   1.00-2.00   sec  1.98 GBytes  17.0 Gbits/sec    0    778 KBytes
client-cluster-05: [  5]   2.00-3.00   sec  2.06 GBytes  17.7 Gbits/sec    0    778 KBytes
client-cluster-05: [  5]   3.00-4.00   sec  2.03 GBytes  17.4 Gbits/sec    0    778 KBytes
client-cluster-05: [  5]   4.00-5.00   sec  2.04 GBytes  17.5 Gbits/sec    0    778 KBytes
client-cluster-05: - - - - - - - - - - - - - - - - - - - - - - - - -
client-cluster-05: [ ID] Interval           Transfer     Bitrate         Retr
client-cluster-05: [  5]   0.00-5.00   sec  10.0 GBytes  17.2 Gbits/sec    0
sender
client-cluster-05: [  5]   0.00-5.00   sec  9.99 GBytes  17.2 Gbits/sec
receiver
client-cluster-05:
client-cluster-05: iperf Done.
client-cluster-06: Connecting to host 192.168.B.F, port 5900
client-cluster-06: [  5] local 192.168.D.6 port 43215 connected to 192.168.B.F port 5900
client-cluster-06: [ ID] Interval           Transfer     Bitrate         Retr  Cwnd
client-cluster-06: [  5]   0.00-1.00   sec  1.78 GBytes  15.3 Gbits/sec    0    900 KBytes
client-cluster-06: [  5]   1.00-2.00   sec  1.93 GBytes  16.6 Gbits/sec    0    900 KBytes
client-cluster-06: [  5]   2.00-3.00   sec  1.94 GBytes  16.7 Gbits/sec    0    935 KBytes
client-cluster-06: [  5]   3.00-4.00   sec  2.03 GBytes  17.5 Gbits/sec    0    987 KBytes
client-cluster-06: [  5]   4.00-5.00   sec  2.02 GBytes  17.3 Gbits/sec    0   1.01 MBytes
client-cluster-06: - - - - - - - - - - - - - - - - - - - - - - - - -
client-cluster-06: [ ID] Interval           Transfer     Bitrate         Retr
client-cluster-06: [  5]   0.00-5.00   sec  9.70 GBytes  16.7 Gbits/sec    0
sender
client-cluster-06: [  5]   0.00-5.00   sec  9.69 GBytes  16.7 Gbits/sec
receiver
client-cluster-06:
```

For 100GbE, multiple streams need to be used in order to see the full bandwidth. This is not supported by the script yet.

2.4 Checking Client Performance With Short Bandwidth Test

It might be a good idea to run short BW constant size stepping test on 100GbE client nodes one by one, to ensure each is indeed getting close to 12.5GB of throughput.

```
$ ./run_test.sh -ns 100 -sf ../workload/fio/bw_test_short_100g.sch -co
$ ./run_test.sh -ns 100 -sf ../workload/fio/bw_test_short_100g.sch -ss 1 -st 1
```

The 1st command will write to create the files but no actual read. If the file has been created already before, it will not rewrite. This is equivalent to using "–create_only" parameter in iops_test_100g.sch for some of the workloads. fio_template.ini has "#create_only" and" #" will be removed if either "-co" is defined in the command line, or "–create_only" is defined in the schedule file.

The 2nd command is the actual read BW test.

Read throughput results like below indicate some 100GbE are not getting the expected throughput. Check the client, NIC card and switch configurations as well as cables to resolve the problem first. For example, a few items might impact the BW are:

- Loose installation of the NIC card.
- Client BIOS configuration "CPU Power and Performance Policy": "Performance" is preferred over "Performance Balanced" or other options.
- Client BIOS configuration "Fan Profile": "Performance" is preferred over "Acoustic".
- DRAM size and speed. For 100GbE, higher DRAM speed works better than bigger DRAM size, eg for some CPUs, 192G of DRAM may run at lower frequency than 128G DRAM, hence produce lower performance.

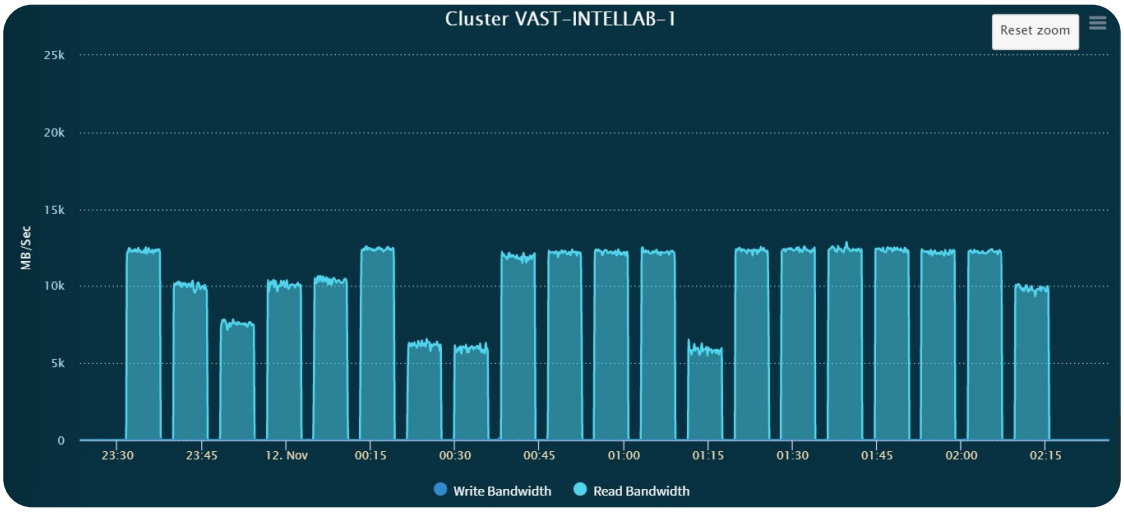

Figure 12: Constant Size Stepping Example

Of course, the same one by one stepping test can also be done on 10GbE to ensure each is getting close to 1.25GB throughput, but with 100 nodes, that will take much longer:

```
$ ./run_test.sh -ns 10 -sf ../workload/fio/bw_test_short_10.sch -co
$ ./run_test.sh -ns 10 -sf ../workload/fio/bw_test_short_10.sch -ss 1 -st 1
```

Doing incremental stepping tests may allow us to generate a quicker view of how the performance scales with the number of nodes. The command below will run a short BW test on 100GbE in steps of 2, 4, 6...20 clients at a time.

```
$ ./run_test.sh -ns 100 -sf ../workload/fio/bw_test_short_100.sch -ss 2
```

This could be helpful to narrow down the slow nodes if the performance doesn't scale as expected. For example, in the screenshot below, we observe step 4 throughput growth is extremely low, and that turns out to be at the step of adding the 6th and 7th slowest client nodes.

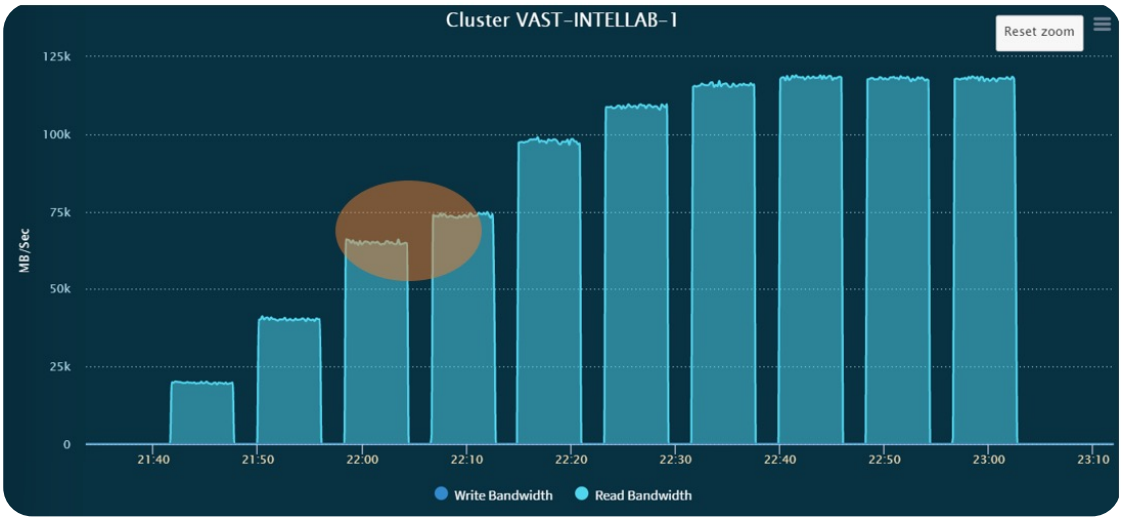

Figure 13: Incremental Stepping Example

However, keep in mind that the performance may not scale linearly when approaching the full BW anyways, even if the nodes are all getting close to full BW throughput individually. Refer to Minimum 100GbE Clients to Saturate for the actual performance scaling trend as the reference on 100GbE.

3) Performance Expectations and Test Results

3.1 Understanding Data Path Bandwidth

To eliminate misleading test results due to architectural bottlenecks, it is important to understand the bandwidth of all components within a given data path. Here are some major factors to be aware of for the hardware configurations used:

Table 10: Data Path Bandwidth

Configuration	Clients		Extreme Switch <-> Arista Switch		Arista Cluster <-> VAST		VAST Cluster		Overall BW	
	GBit	GB	GBit	GB	GBit	GB	GBit	GB	GBit	GB
100 Clients With 10GbE	100 x 10 = 1000	100 x1.25 = 125	1120 (7s x 4p x 40 Gbit/p)	140 (7s x 4p x 5 GB/s)	1600 (16p x 100 Gbit/p)	1600 (16p x 100 Gbit/p)	960	120	960	120
20 Clients With 100GbE	20 x100 = 2000	20 x 12.5 = 250	N/A	N/A	1600 (16p x 100 Gbit/p)	1600 (16p x 100 Gbit/p)	960	120	960	120

Notes: s – Switch, p – Switch Port

The maximum bandwidth will be limited by the minimum value of the entire data path.

3.2 Performance Test Results

In this section, we'll showcase some BW/IOPS and Latency test results using the predefined workloads from the test harness. All results in this section are averaged across 3 separate runs. The exact command lines to produce the results are listed. To parse the result, please refer to Test Results Explained for more details.

3.3 Bandwidth (1024K) by Data Size

For benchmarking, bandwidth can be data size dependent as it determines where the data will be located, i.e. in Intel® Optane™ write buffer or backend capacity QLC NAND drives. To understand VAST Data's architecture, please refer to UNIVERSAL STORAGE EXPLAINED.

1. Three different data sizes (total data written) were used for the tests:
2. 1.6TB
3. 9.6TB or 10TB
4. 24TB

Command line:

```
$ ./run_bw_test.sh
```

Table 11: BW Summary (1.6TB)

Benchmarking Tool	Workload Type	100 Clients with 10GbE				20 Clients with 100GbE			
		Write		Read		Write		Read	
		GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s
FIO	Sequential	21.1	19.7	57	53.1	20.1	18.7	53.7	50
	Random	22	20.5	60.5	56.3	22.4	20.9	60.2	56.1
Elbencho File	Sequential	22.3	20.8	57.2	53.3	20.7	19.3	53.3	49.6
	Random	22	20.5	61	56.8	20.9	19.5	58.4	54.4
Elbencho S3	Sequential	21.7	20.2	55.5	51.7	19.1	17.8	53.6	49.9
	Random	N/A	N/A	55	51.2	N/A	N/A	52.7	49.1

Table 12: BW Summary (9.6TB/10TB)

Benchmarking Tool	Workload Type	100 Clients with 10GbE				20 Clients with 100GbE			
		Write		Read		Write		Read	
		GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s
FIO (9.6T)	Sequential	18.8	17.5	110.9	103.2	18.9	17.6	114.4	106.6
	Random	21.3	19.8	110.9	103.3	21.2	19.7	117.7	109.6
Elbencho File (9.6T)	Sequential	20.4	19	108.7	101.3	19.8	18.4	114.1	106.2
	Random	21.4	19.9	108.7	101.4	21.7	20.2	115	107.1
Elbencho S3 (9.6T)	Sequential	19.8	18.4	106.3	99	18.2	17	110	102.5
	Random	N/A	N/A	106.6	99.3	N/A	N/A	111	103.4

Table 13: BW Summary (24TB)

Benchmarking Tool	Workload Type	100 Clients with 10GbE				20 Clients with 100GbE			
		Write		Read		Write		Read	
		GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s	GB/s	GiB/s
FIO	Sequential	18.5	17.3	108.7	101.2	18.4	17.2	115.5	107.5
	Random	20.7	19.3	109.4	101.9	21.1	19.7	116.5	108.5
Elbencho File	Sequential	19.5	18.2	108.5	101	19.8	18.5	109.7	104.5
	Random	21	19.6	108.9	101.5	21.2	19.7	110.1	104.9
Elbencho S3	Sequential	18.2	17	105.8	98.5	17.3	16.1	110.7	103.1
	Random	N/A	N/A	104	96.9	N/A	N/A	110.2	102.6

Notes:

The bar graph below shows the comparison of random read BW by data size of 100GbE and 10GbE, with 3 different benchmarking methods: FIO, Elbencho File store over NFSv3 protocol, and Elbencho S3 Object store over S3 protocol.

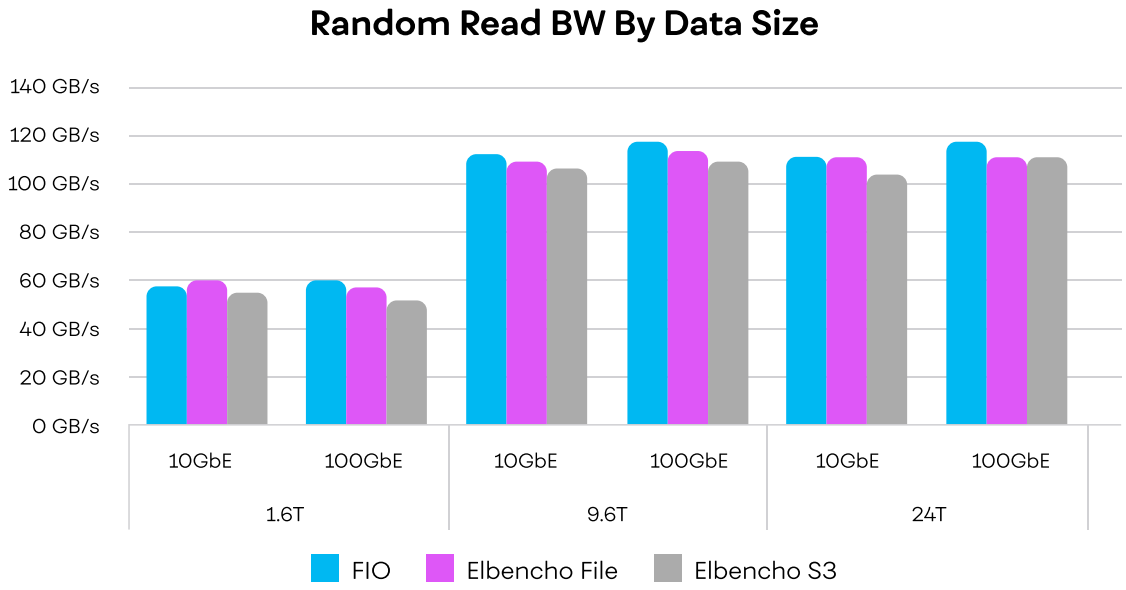

Figure 14: Random Read BW by Data size

Highlights:

- File Store:
  - 10G: FIO and Elbencho results are comparable.
  - 100G: Elbencho results are a little lower.
- S3 access: only slightly lower (~2-3%) than NFS-based access.
- Efficiency using 9.6TB FIO result:
  - 20 x 100GbE clients: 117.7GB/s (max theoretical of 120GB/s) = 98.0%.
  - 100 x 10GbE clients: 110.9GB/s (max theoretical of 120GB/s) = 92.4%.
- Read bandwidth is largely determined by where the data is persisted:
  - 1.6T: data is still in Intel® Optane™, bandwidth is lower due to the small Intel® Optane™: QLC NAND ratio (2.9%).
  - 9.6T/24T: data are mostly in capacity QLC NAND. Bandwidth is much higher with high concurrency.

Important: in real-world scenarios, the lower bandwidth of small data set sizes is rarely a concern as subsequent writes keep pushing previously written data to QLC NAND.

3.4 Bandwidth Saturation Test: 100GbE Clients

Command line:

```
$ ./run_fio_bw_test_short.sh
```

The theoretical maximum bandwidth of a single 100GbE is 12.5GB/s, which means that less than 20 clients are needed to saturate the maximum theoretical bandwidth (120GB/s) of this network configuration.

Table 14: 1024KB Random Read BW by Number of 100GbE Clients

Client	2	4	6	8	10	12	14	16	18	20
Actual BW	24.5	48.5	71.6	92.2	107.7	113.7	116.5	117.4	117.6	117.3
Theoretical BW	25	50	75	100	120	120	120	120	120	120

Figure 15 shows the plot of scaling trend from data in Table 14.

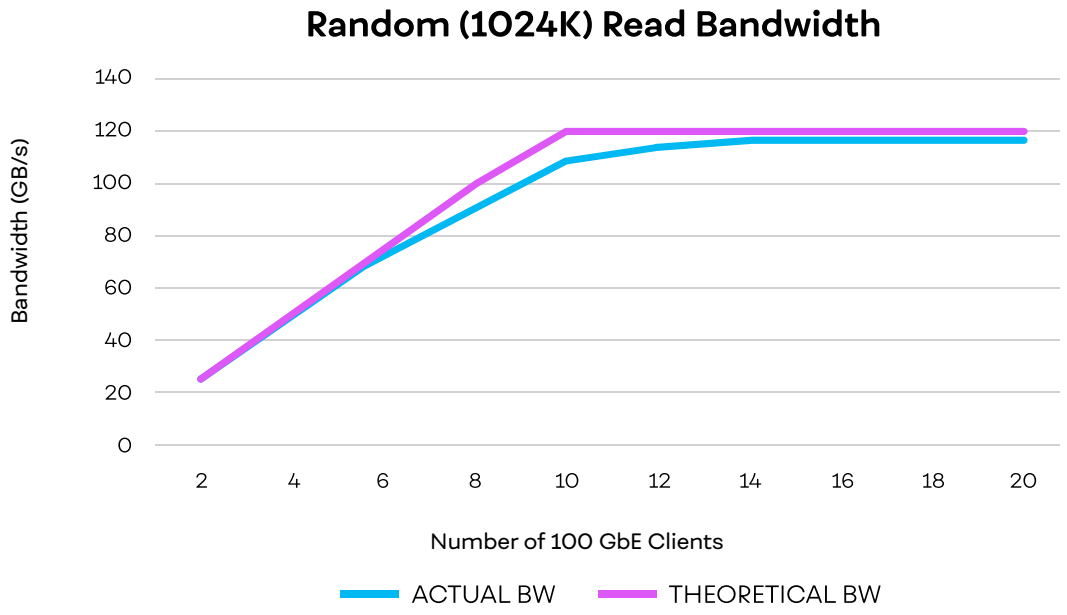

Figure 15: Random (1024K) Read BW by Number of 100GbE Clients

Highlights:

- Bandwidth scales linearly up to 8 clients, as it approaches max theoretical bandwidth of the cluster.
- It takes ~14-16 clients to saturate this hardware configuration.

3.5 Bandwidth Saturation Test: 10GbE Clients

Command line:

```
$ ./run_fio_bw_test_short.sh
```

Likewise, the theoretical maximum bandwidth of a single 10GbE is 1.25GB/s.

Table 15: Random (1024K) Read BW by Number of 10GbE Clients

Client	10	20	30	40	50	60	70	80	90	100
Actual BW (GB/s)	12.3	24.7	37	48.9	61.3	73.6	80.9	92.6	102.9	110.5
Theoretical BW (GB/s)	12.5	25	37.5	50	62.5	75	87.5	100	112.5	120

Figure 16 shows the plot of scaling trend from data in Table 15.

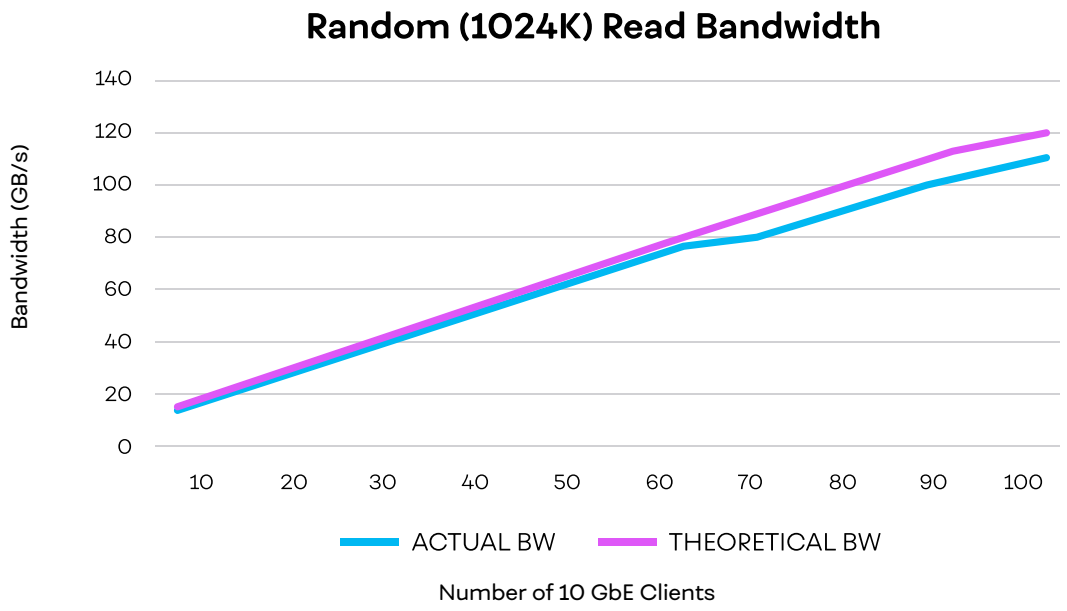

Figure 16: Random (1024K) Read BW by Number of 10GbE Clients

Highlights:

- Bandwidth scales linearly from 10 to 100 clients (in steps of 10).
- With 100 10GbE clients, total throughput reaches 92% utilization of max theoretical bandwidth.

3.6 S3 Random Read Bandwidth by Object Size

In addition to NFS-based performance testing, a S3 bandwidth test across various object sizes was also conducted. The results are detailed in Figure 17.

Command line:

```
$ ./run_elbencho_s3_sweep_iosizes.sh
```

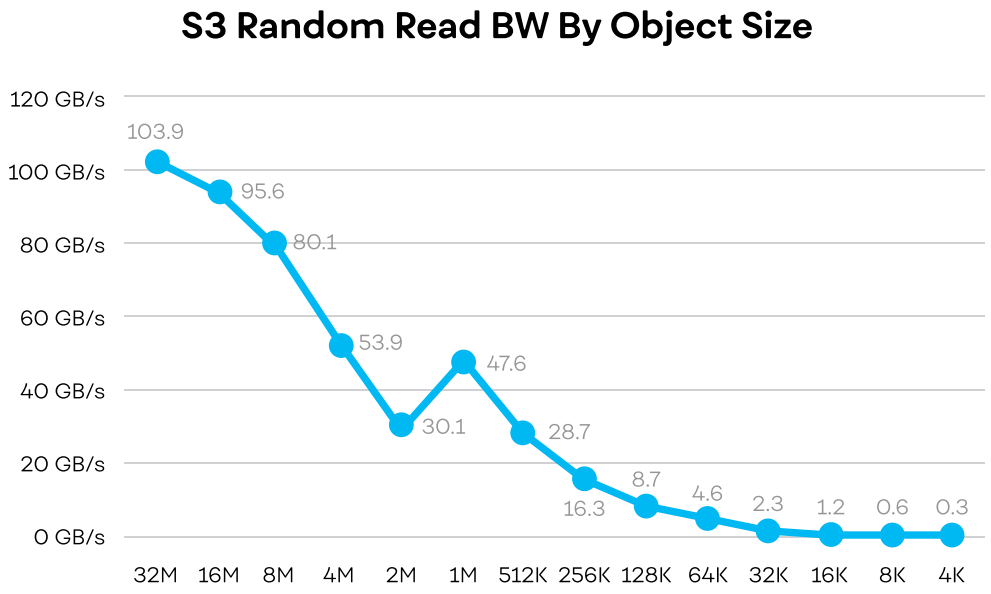

Figure 17: S3 Random Read BW by Object size

Highlights:

- Bandwidth goes up as object size increases in general, except for a dip at 2MB object size.

Notes

- Special tuning by VAST Data may resolve this dip using 2MB objects. Please contact support@vastdata.com if better performance for 2MB is needed.

3.7 Random (4K) IOPS Test

Command line:

```
$ ./run_iops_test.sh
```

Table 16: Random IOPS Summary

Benchmarking Tool	BS	Workers	iodepth	100 Clients with 10GbE			100 Clients with 10GbE		
				Write IOPS (K)	Read IOPS (K)		Write IOPS (K)	Read IOPS (K)	
					Before Flush	After Flush		Before Flush	After Flush
FIO	4K	4	64	309	895	493	347	944	477
		4	16	301	903	532	301	830	473
Elbencho File	4K	4	64	300	866	613	299	780	558
		4	16	283	875	614	263	773	529
Elbencho S3	4K: read 5M: write	4	N/A	4	120	95	2	31	24

Figure 18 shows the plot of the Random Read IOPS from Table 16.

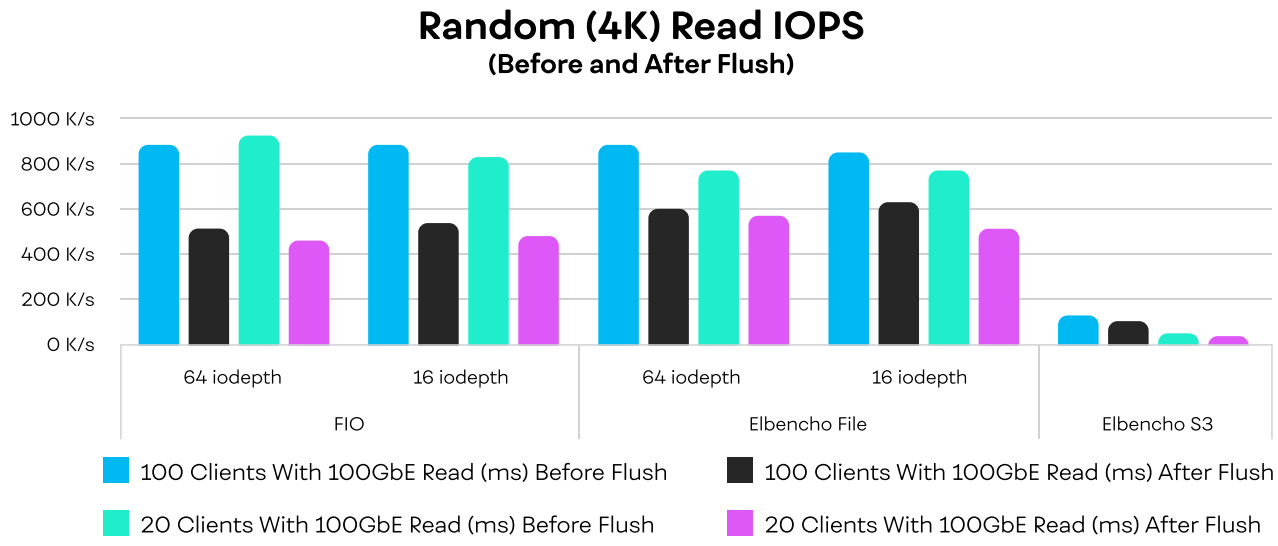

Figure 18: Random (4K) Read IOPS Before and After Flush

Notes:

- Before flush: data in Intel® Optane™ (data size < write buffer threshold)
- After flush: data in QLC NAND (data size > write buffer threshold)

Highlights:

- Random Read IOPS is 69%~97% higher before flush (FIO), due to Intel® Optane™ SSD's high IOPS.

3.8 Random (4K) Read Latency Test

Latency results are collected with the same IOPS test command in the previous section.

Table 17: Random Average Latency Summary

Benchmarking Tool	BS	Workers	iodepth	100 Clients with 10GbE			100 Clients with 10GbE		
				Write IOPS (K)	Read IOPS (K)		Write IOPS (K)	Read IOPS (K)	
					Before Flush	After Flush		Before Flush	After Flush
FIO	4K	4	64	82.5	28.5	51.8	15	5.4	10.7
		4	16	21.1	7	12	4.2	1.5	2.6
Elbencho File	4K	4	64	84.1	29.6	41.5	17.3	6.5	9.1
		4	16	22.5	7.2	10.4	4.8	1.6	2.4
Elbencho S3	4K: read 5M: write	4	N/A	97.3	3.3	4.1	42.9	2.5	3.2

Figure 19 shows the plot of the latencies from Table 17.

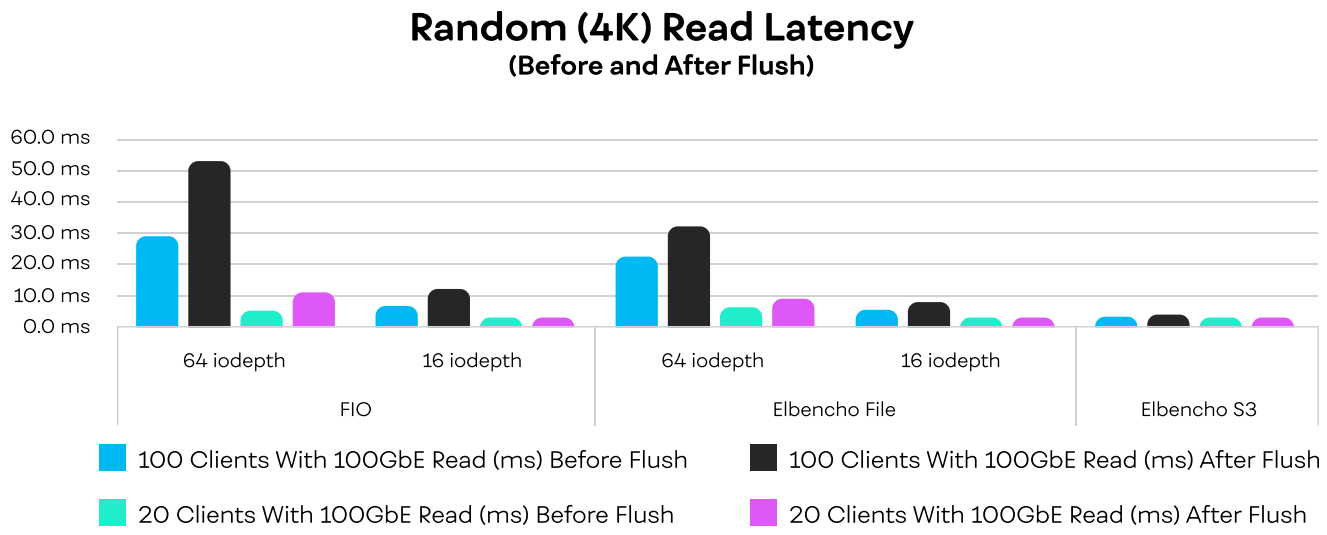

Figure 19: Random (4K) Read Latency Before and After Flush

Notes:

- Before flush: data in Intel® Optane™ (data size < write buffer threshold)
- After flush: data in QLC NAND (data size > write buffer threshold)

Highlights:

- Latency is ~41.6% ~49.5% lower before flush (FIO), due to the low latency characteristics of Intel® Optane™ SSD's

3.9 Performance Impact of Intel® Optane™

Intel® Optane™ SSD is used as a write buffer in-between CNode and QLC NAND on DNodes in VAST Data's solution, and is a key component of the architecture due to its high BW/IOPS, low latency, high endurance and persistence:

- It stores all metadata and speeds up user data access to QLC NAND, as metadata indicates where the user data is located. Slower metadata access would slow user data access to QLC NAND, hindering the backend access performance.
- It enables VAST Data's Similarity-based data reduction: algorithm requires many fast lookups in hash tables that are too big for DRAM and needs to be persistent.
- It enables data center use-cases for QLC NAND by extending flash lifetime from 5 years to 10 years, providing customers with a dramatically improved TCO.

4) Schedule Files and Running Customized Tests

Schedule files usually fall into 4 different categories:

- BW tests with 1024KB
- IOPS tests with 4KB
- Workers/threads & iodepth sweep tests with 1024KB
  - These are intended to find the optimal configuration for the BW test.
  - Our test result shows 16 workers / 8 iodepth is the optimal configuration
- iosize sweep tests
  - These are intended to do BW tests with the optimal workers/threads * iodepth configuration for iosizes in power of 2.
  - Elbencho File and S3 focus more on 100% read test
  - FIO also has iosize sweep test for:
    - 70% read + 30% write per host test, e.g. to run 70/30 mixed workload on each 100GbE client, use the command below:

```
$ ./run_test.sh -ns 100 -sf ../workloads/fio/sweep_iosizes_70rd_random_100g.sch -co
$ ./run_test.sh -ns 100 -sf ../workloads/fio/sweep_iosizes_70rd_random_100g.sch
```

    - 70% host read + 30% host write test, e.g. to run read workload on 70% 100GbE clients, write workload on 30% of the 100GbE clients, use the command below:

```
$ ./run_test.sh -ns 100 -sf ../workloads/fio/sweep_iosizes_70hostrd_random_100g.sch -co
$ ./run_test.sh -ns 100 -sf ../workloads/fio/sweep_iosizes_70hostrd_random_100g.sch
```

  - Running FIO tests on 100GbE and 10GbE simultaneously is also possible by opening 2 consoles and starting 1 test in each console, providing CLIENT_NODES_10G and CLIENT_NODES_100G defined do not overlap.

Example wrapper scripts of how to use the schedule files are provided in the scripts folder. Check.

5) Appendix

All test results in this Guide are tested with changeset listed in Software Configurations. For quick referencing, the actual schedule file contents of that particular changeset are also listed here.

5.1 References

UNIVERSAL STORAGE EXPLAINED, by VAST Data

VAST_Data-Overview.pdf, by VAST Data, Feb 2021

VAST Data Documentation, by VAST Data

VAST Data Knowledge Base, by VAST Data

Source Repo: https://github.com/breuner/elbencho, by Sven Breuner, VAST Data

Docker Registry: https://hub.docker.com/r/breuner/elbencho, by Sven Breuner, VAST Data

https://www.arista.com/assets/data/pdf/Datasheets/7170-Datasheet.pdf, by Arista Networks, Inc, 2021

https://www.ibm.com/cloud/blog/object-vs-file-vs-block-storage, By IBM Cloud Education, October 2021

5.2 Acronyms and Glossaries

Name	Description
NFS	Network File System
S3	Simple Storage Service
VIP	Virtual IP
VMS	VAST Management System
SSD	Solid State Drive
Intel® Optane™	Intel's revolutionary memory and storage innovation, built with a ground-breaking combination of endurance, consistent high performance, and low latency is designed to bring new computing possibilities to a variety of markets.
CIFS	Common Internet File System (CIFS) is a network file system protocol used for providing shared access to files and printers between machines on the network.
File Store	File storage is when all the data is saved together in a single file with a file extension type that's determined by the application used to create the file or file type, such as .jpg, .docx or .txt. For example, when you save a document on a corporate network or your computer's hard drive, you are using file storage. Files may also be stored on a network-attached storage (NAS) device. These devices are specific to file storage, making it a faster option than general network servers. Other examples of file storage devices include cloud-based file storage systems, network drives, computer hard drives and flash drives.
Object Store	Object storage is a system that divides data into separate, self-contained units that are re-stored in a flat environment, with all objects at the same level. There are no folders or sub-directories like those used with file storage. Additionally, object storage does not store all data together in a single file. Objects also contain metadata, which is information about the file that helps with processing and usability. Users can set the value for fixed-key metadata with object storage, or they can create both the key and value for custom metadata associated with an object.
Block Store	Block storage is when the data is split into fixed blocks of data and then stored separately with unique identifiers. The blocks can be stored in different environments, such as one block in Windows and the rest in Linux. When a user retrieves a block, the storage system reassembles the blocks into a single unit. Block storage is the default storage for both hard disk drive and frequently updated data. You can store blocks on Storage Area Networks (SANs) or in cloud storage environments.

5.3 run_test.sh/run_test_no_runlog.sh command line options

Command Line Option	Description
-col--create-only	For creating the FIO files only, no actual read access. Default: 0.
-drl--dryrun	Do not run workload. Everything else will be executed the same as a normal run including nfs mounts and starting fio and Elbencho services on clients.
-nsl--network-speed	GbE speed, valid values ["10", "100"]. Default: 10
-pfl--prefix	Extra prefix to be added to the test name
-sfl--schedule-file	Specify the schedule file to run.
-sml--skip-mount	Specify this to skip the mount step. This can speed up the test if the mount is the same as the previous test.
-sql--skip-all-client-config-query	Do not do query of all client configurations before the test. As query takes a few minutes, this allows the test to be started sooner for quicker benchmarking.
--ssl--step-size	Specify the number of clients to increase to run the test. Default: max clients. Examples: when st==0: if max clients = 100, then "-ss 10" will test clients in steps of 10, 20, 30...90, 100 clients if max clients = 100, then "-ss 15" will test clients in steps of 15, 30, 45...75, 90, 100 clients if max clients = 100, then "-ss 0" or "-ss 100" or omitting -ss option will test all 100 clients in 1 step when st==1: if max clients = 10, then "-ss 1 -st 1" will test the 10 clients 1 at a time, in 10 steps if max clients = 10, then "-ss 2 -st 1" will test the 10 clients 2 at a time, in 5 steps
-stl--step-type	0 - incremental (default), 1 - constant. Refer to -ssl--step-size examples to see the differences between the two types
--runtime	FIO runtime. Default: as specified in template FIO job file
--ramp_time	FIO ramptime. Default: as specified in template FIO job file
-hl--help	To print this help

5.4 Test Results Explained

5.4.1 EXAMPLE FIO RESULT

```
results_2021-11-13-21.41.22_20Clients_bw_test_short_100g_headnode
├── bw_test_short_100g.zip         => folder containing actual fio job files per client,
 |                                    zipped to reduce size
├── bw_test_short_100g.sch         => schedule file for this test
├── context.zip                    => running contexts, zipped to reduce size
├── env.conf                       => env.conf for this test
├── env.conf.override              => env.conf.override for this test
├── results_2021-11-13-21.41.22_20Clients_bw_test_short_100g_headnode_Logs.zip
 |                  => Per client FIO json+ results.
 |                  For aggregated results, check the "All Clients" sub folder.
 |                  The json+ file names match their FIO job file names,
 |                  except for extension (.ini for job file and .json for json+ result.)
 |                  Identify each workload by the matching job filename.
└── run.log         => console output for this test
```

FIO job filename is auto generated based on mandatory parameters defined by the workload, eg:

a) --rw=randread --numjobs=16 --iodepth=8 --bs=1024k --size=30g --runtime=300 --ramp_time=60

This will generate: rd_rnd_qd_128_1024k_16w.ini, where:

128: equals to numjobs 16 x iodepth 8.

1024k: the block size.

16: equals to numjobs.

rd_rnd: random read.

b) --rw=read --numjobs=16 --iodepth=8 --bs=1024k --size=30g --runtime=300 --ramp_time=60

This will generate: rd_qd_128_1024k_16w.ini, where:

rd: sequential read (as no rnd before qd).

c) --rw=read --numjobs=16 --iodepth=8 --bs=1024k --size=30g --runtime=300 --ramp_time=60 –tag=1

This may generate: rd_qd_128_1024k_16w_5tag.ini, where:

5: the workload number in the .sch, starting from 1.

5 means this is the 5th workload in the schedule file.

Note that –name parameter only affects the FIO generated data filename, not FIO job filename.

5.4.2 EXAMPLE ELBENCHO FILE RESULT

Here is an example result folder for Elbencho file testing:

```
results_2021-11-10-07.11.45_20Clients_bw_test_100g_headnodename
├── bw_test_100g.sch           => schedule file for this test
├── context.zip                => running context for this test, zipped to reduce size
├── elbencho
│   └── file
│       ├── rd-1024kbs-16t-8iodepth_rnd.csv          => result csv files
│       ├── wr-1024kbs-16t-8iodepth.csv
│       └── wr-1024kbs-16t-8iodepth_rnd.csv
├── env.conf                         => config file for this test
├── env.conf.override                => env.conf.override for this test
└── run.log                          => console output
```

Note that the results of multiple workloads in a schedule file might collapse to the same .csv file. Identify the workload by the attributes of the workloads, eg threads, size, bs size, random or not etc as defined in the schedule file for each workload. If 2 workloads have the same attributes, then identify by the order of the workload and the order or runtime in the .csv.

5.4.3 EXAMPLE ELBENCHO S3 RESULT

```
results_2021-11-11-01.04.48_100Clients_iops_test_10g_redwood-bdw106_run3
├── context.zip     => running context for this test, zipped to reduce size
├── elbencho
│   └── s3
│       ├── rd-4kbs-4t-iodepth.csv               => result .csv files
│       ├── wr-32mbs-16t-iodepth.csv
│       └── wr-5mbs-4t-iodepth.csv
├── env.conf                    => configuration file for this test
├── env.conf.override           => env.conf.override for this test
├── iops_test_10g.sch           => schedule file for this test
└── run.log                     => console output for this test
```

The same as the Elbencho file result, the results of multiple S3 workloads in a schedule file might collapse to the same .csv file as well. Identify the workload by the attributes of the workloads, e.g. threads, size, bs size, random or not etc as defined in the schedule file for each workload. If 2 workloads have the same attributes, then identify by the order of the workload and the order or runtime in the .csv.

5.5 Test Harness Top Folder

5.5.1 LICENSE

```
SPDX-License-Identifier: BSD-3-Clause


Intel Copyright © 2022, Intel Corporation.


Redistribution and use in source and binary forms, with or without modification,
are permitted provided that the following conditions are met:


1. Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
t
2. Redistributions in binary form must reproduce the above copyright notice,
this list of conditions and the following disclaimer in the documentation and/or
other materials provided with the distribution.


3. Neither the name of the copyright holder nor the names of its contributors may be
used to endorse or promote products derived from this software without specific prior
written permission.


THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS “AS IS” AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED
WARRANTIES
OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT
SHALL
THE COPYRIGHT HOLDER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL,
EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE
```

```
GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND
ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE
OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE
POSSIBILITY OF SUCH DAMAGE.
```

5.5.2 README.MD

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# This is the test harness developed to make the scale testing of VAST Data:
#    1. effortlessly
#    2. consistently
#    3. easy to customize
#
# For how to use the harness, please refer to “Test Preparations” of: [link to be updated]
```

5.5.3 .GITIGNORE

```
results_*
env.conf.override
```

5.6 Test Harness Shell Scripts (scale-testing-for-VASTdata/scripts)

5.6.1 CLEANUP_MOUNTS.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


network_speed=${1:-”10”}

source env.conf $network_speed


MOUNT=$MOUNT_MULTIPATH_ENABLED

clush -w $CLIENT_NODES “for i in \$(seq 1 \$(mount | grep multipath-tcp | wc -l | cut
-d’ ‘ -f 2)); do sudo umount $MOUNT > /dev/null 2>/dev/null; done”


MOUNT=$MOUNT_MULTIPATH_DISABLED

remote_path_all=($REMOTE_PATH_FIO $REMOTE_PATH_ELBENCHO_FILE $REMOTE_PATH_ELBENCHO_S3)

for REMOTE_PATH in ${remote_path_all[@]}

do

    clush -w $CLIENT_NODES -f $CLUSH_MAX_FAN_OUT “sudo umount ${MOUNT}/${REMOTE_PATH}/*
> /dev/null 2>/dev/null &” &

done

wait
```

5.6.2 CLEANUP_TEST.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause
```

```
network_speed=${1:-"10"}
source env.conf $network_speed

clush -w $HEAD_NODE,$CLIENT_NODES "sudo pkill -9 fio"
sudo pkill -9 fio
sudo pkill -9 run_test_no_runlog.sh
sudo pkill -9 run_test.sh

docker container kill elbencho-client > /dev/null 2>/dev/null
docker container rm elbencho-client > /dev/null 2>/dev/null
clush -w $CLIENT_NODES -f $CLUSH_MAX_FAN_OUT "docker container kill elbencho-server > /
dev/null 2>/dev/null;

docker container rm elbencho-server > /dev/null 2>/dev/null;
                                                    sudo pkill -9 nfs_mount.sh"
```

5.6.3 CONTAINER_LOGS.SH

```
#!/bin/bash
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

network_speed=${1:-"10"}

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" "$network_speed"
clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "docker logs elbencho-server"
```

5.6.4 DOCKER_START.SH

```
#!/bin/bash
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause
```

```
network_speed=${1:-"10"}


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" "$network_speed"


clush -w "$HEAD_NODE","$CLIENT_NODES" "sudo systemctl start docker"
clush -w "$HEAD_NODE","$CLIENT_NODES" "sudo chmod 666 /var/run/docker.sock"


# Direct docker pull is disabled for Anonymous and authenticated docker users
# to use elbencho_deploy.sh as the workaround instead.
#clush -w "$HEAD_NODE","$CLIENT_NODES" "docker pull breuner/elbencho"
```

5.6.5 ELBENCHO_DEPLOY.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Anonymous and authenticated docker users may run into docker pull limit quickly.
# This script provides a workaround to manually pull once and deploy the elbencho
# docker image to all systems.


# How to use it:
#
# Step 1: pull once
#     docker pull breuner/elbencho
# Step 2: note down "IMAGE ID" of the pulled image
#     docker image ls
# Example response:
#     REPOSITORY                  TAG          IMAGE ID      CREATED      SIZE
```

```
# docker.io/breuner/elbencho  latest       6ac5b651eaa9 4 days ago   136 MB
# Step 3: update ELBENCHO_ID and ELBENCHO_TAG accordingly in env.conf.
# Tag is user defined string to help identify the version in a more convenient way.
# Step 4: run this script to deploy accordingly

# For more information of how to use elbencho, please refer to:
#     https://github.com/breuner/elbencho
#     https://hub.docker.com/r/breuner/elbencho
#
# To see the help of elbencho:
#     docker run --net=host -it breuner/elbencho --help-all

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" 10 # 10 used here because all 100G nodes are 10 nodes as
well. This may not be general enough. To be improved.

clush -w "$HEAD_NODE","$CLIENT_NODES" "sudo systemctl start docker"
clush -w "$HEAD_NODE","$CLIENT_NODES" "sudo chmod 666 /var/run/docker.sock"

mkdir -p "$ELBENCHO_IMAGE_FOLDER"
if [ ! -f "$ELBENCHO_IMAGE_FOLDER"/"$ELBENCHO_IMG" ]; then
    docker image tag "$ELBENCHO_ID" breuner/elbencho:$ELBENCHO_TAG
    docker image save -o "$ELBENCHO_IMAGE_FOLDER"/"$ELBENCHO_IMG" "$ELBENCHO_ID"
fi
clush -w "$CLIENT_NODES" "mkdir -p $ELBENCHO_IMAGE_FOLDER"
clush -w "$CLIENT_NODES" -c $ELBENCHO_IMAGE_FOLDER/$ELBENCHO_IMG
clush -w "$CLIENT_NODES" "docker image load -i $ELBENCHO_IMAGE_FOLDER/$ELBENCHO_IMG"
clush -w "$CLIENT_NODES" "docker image tag $ELBENCHO_ID breuner/elbencho:$ELBENCHO_TAG"
```

5.6.6 COMMON.CONF

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Convenient variables for files that:
#  - need to know the folder structures
#  - however don’t need to be client and VIP aware directly

# Harness folder structure
source=”${BASH_SOURCE}”
SCRIPT_DIR=”$(cd “$(dirname “$source”)” || exit; pwd)”
TOP_DIR=${SCRIPT_DIR/%\/scripts/}
WORKLOAD_DIR=”$TOP_DIR/workloads”
RESULTS_DIR=”$TOP_DIR/results”

# vast-scale-testing will be created under $HOME/ to facilitate the installations and
tests.
# For FIO non-Server/Client mode, per client result folders will also be created there.
VAST_SCALE_TESTING_PATH=”$HOME/vast-scale-testing”

# Generated by run_test_no_runlog.sh to pass the running dir
# to parent run_test.sh. Removed after the test.
RUN_DIR_TEMP_FILE=”$VAST_SCALE_TESTING_PATH/run_dir.txt”

# Generated by run_test.sh to capture console output.
# Moved to result folder after the test.
RUN_LOG_FILE=”$VAST_SCALE_TESTING_PATH/run.log”

# Generated by rUn_test_no_runlog.sh to capture stderr with clush,
```

```
# in order to exit the test harness properly on failure.
# Moved to result folder after the test.
RUN_ERROR_LOG_FILE="$VAST_SCALE_TESTING_PATH/run_error.log"

# Generated by nfs_mount.sh to capture stderr nfs_mount.sh,
# in order to exit the test harness properly on failure.
# Moved to result folder after the test.
NFS_MOUNT_LOG_FILE="$VAST_SCALE_TESTING_PATH/nfs_mount.log"

# Mountpoint for results upload on head node
# The test harness automatically mount the sharepoint, upload the results to
sharepoint,
# and then umount the sharepoint before ending the test.
RESULTS_MNTPOINT="/mnt/vast-scale-testing-results"
```

5.6.7 ENV.CONF

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

network_speed=${1:-"10"}
step_size=${2:-"0"}

#================== Input validation ================================================
============#
source="${BASH_SOURCE}"
scripts_dir="$(cd "$(dirname "$source")" || exit; pwd)"
source "$scripts_dir/common.conf"
if [ "$network_speed" != "10" ] && [ "$network_speed" != "100" ]; then
```

```
  echo “Invalid network speed specified, must be in [10, 100].”
    exit 22 # Standard EINVAL for Invalid ArgumentL
fi

#==================== User Defined Dummy Variables, Update Accordingly
==========================#

# Update with the actual IPs and client cluster hostnames.

# The test harness assume the VIPs are always contiguous.
#
# Client node IPs don’t have to be contiguous, ie it can be client-
cluster[05,10-20,50-55]
# The number of client nodes can be any number >= 1. The lead node(FIO client) shall
not be
# in the list.

_VIP_POOL_10G=”10.A.B.110-10.A.B.157”
_VIP_POOL_100G=”192.168.B.110-192.168.B.157”
_CLIENT_NODES_10G=”client-cluster[01-100]”
_CLIENT_NODES_100G=”client-cluster[01-20]”

# Replace the key and secret with the valid ones
ELBENCHO_S3KEY=”YU5...LQ1”
ELBENCHO_S3SECRET=”aqb...EtnQ”

# For result upload to NFS share on remote server,
# update needed only if RESULTS_UPLOAD_ENABLE is 1.
SHAREPOINT_HOST=”YourSharepointHostName”
SHAREPOINT_FOLDER=”//${SHAREPOINT_HOST}/YourRemotePathOnSharepointHost”
```

```
#==================== User Defined Tunable Variables, Configure Accordingly
==========================#

# Intel Turbo Boost control
TURBO_BOOST_ENABLE=0

# Hyper Threading control
HYPER_THREADING_ENABLE=0

# irqbalance control
IRQBALANCE_ENABLE=0

# Nodes can be defined as ordered or unordered.
#
# 1. Ordered nodes can be defined in format of:
#    a) node[01-02,30-32,10-12], or
#    b) node{01..02} node{30..32} node{10..12}
# Nodes will be expanded to "node01 node02 node10 node11 node12 node30 node31 node32"
with nodeset --expand

# 2. Unordered nodes SHALL be defined in format of:
#    node{01..02} node{30..32} node{10..12}
# Nodes will be expanded to "node01 node02 node30 node31 node32 node10 node11 node12"
with eval only
ORDERED_NODES=1

# To enable/disable FIO Server/Client mode.
# 1: to enable FIO Server/Client mode.
#     Aggregation: results are aggregated in fio json+ output as "All_Clients"
#     Where:       head node "results" folder
```

```
# 0: to disable FIO Server/Client mode.
#     Aggregation: no result aggregation of all clients, json+ output is per client.
#     Where:       in $HOME/$VAST_SCALE_TESTING_PATH folder
# For better total BW/IOPS/Latency tracking, use Server/Client mode.

FIO_SERVER_CLIENT_MODE=1

# Mount configurations
# Kernel dependent Multipath driver is needed in order to enable this.
# Contact support@vastdata.com for it if needed.
# Enabling Multipath allows more balanced load on VAST Data CNodes and DNodes,
# especially when worker/thread count is low.
#
#When the variable is set to 0, VIPs are going to be mounted individually.
MULTIPATH_ENABLE=1

# Applicable to FIO and  MULTIPATH_ENABLE=0 only.
# 1: 1 mount per VIP per client node. Overall mounts per client = VIP number.
#    Mounts are created before any workloads start.
#
#    This is a preferred way of mounting, as it's faster and no noticeable
#    performance penalty due to over mount, i.e. mount number > job/worker number.
#
# 0: 1 mount per job/thread per client node. Overall mounts per client = job/thread
number.
#    Mounts are created on the fly when running workload. This could be slower.
MOUNT_ALL=1

# This is only applicable to MULTIPATH_ENABLE=0 and MOUNT_ALL=0 .
# This is to avoid creating too many clush sessions when doing individual mounts on
the fly.
```

```
MOUNT_STAGGER=5


# The mountpoint on client system when Multipath is enabled.
MOUNT_MULTIPATH_ENABLED=”/mnt/nfs-multipath-tcp”


# The mountpoint on client system when Multipath is disabled.
MOUNT_MULTIPATH_DISABLED=”/mnt/nfs-no-multipath”


# NConnect allows multiple TCP connections to be created for each mount,
# hence improves the BW effectively. This is independent of Multipath,
# and can be used together with Multipath.
# NCONNECT_ENABLE must be set to 1 for this to take effect.
NCONNECT_NUM=48


# When MULTIPATH_ENABLE=0, and NCONNECT_ENABLE=1, big NCONNECT_NUM means large number
of TCPs
# will be created per client as the table below shows. This could:
#    (1) take a long time to mount before a test can start
#    (2) cause connection aborts if too many contentions. Best practice is to use:
#            MULTIPATH_ENABLE=1
#            NCONNECT_ENABLE=1
#        Or:
#            MULTIPATH_ENABLE=0
#            NCONNECT_ENABLE=0
#
#     Per Client TCP Connections
#     ================================================================
#     MULTIPATH_ENABLE     | NCONNECT_ENABLE | Per Client TCPs
#     ----------------------------------------------------------------
#     1                  |1                   |NCONNECT_NUM
```

```
#     1                   |0                    |1
#     0                   |0                    |VIP Count
#     0                   |1                    |VIP Count * NCONNECT_NUM
#     ----------------------------------------------------------------
NCONNECT_ENABLE=1

# Folder to be created under VIP root for FIO, elbencho file and object testing
REMOTE_PATH_FIO="fio"
REMOTE_PATH_ELBENCHO_FILE="elbencho-file"
REMOTE_PATH_ELBENCHO_S3="elbencho-s3"

# If FIO is installed in /usr/local/bin or other path, change FIO_PATH accordingly
FIO_PATH="/usr/bin"

# Elbencho definitions
ELBENCHO_ID="6ac5b651eaa9"  # Use "docker image ls" to find it on lead node after the
docker pull
ELBENCHO_TAG="scaletest01"  # Change the tag after a new deployment to help identify
the image
ELBENCHO_IMG="elbencho.$ELBENCHO_TAG" # Full image identifier by tag
# This defines the folder to store the elbencho image on head and client nodes.
ELBENCHO_IMAGE_FOLDER="$VAST_SCALE_TESTING_PATH/elbencho"

# Disabled by default.
# To enable this, make sure the sharepoint is set up properly
RESULTS_UPLOAD_ENABLE=0
# NFS share by default. The entire mount command can be redefined in case other type
of
# mount is preferred. Note that the command must start with "sudo mount".
RESULTS_UPLOAD_MOUNT_CMD="sudo mount -t nfs ${SHAREPOINT_FOLDER} ${RESULTS_MNTPOINT}"
```

```
# General test controls
INTER_WORKLOAD_WAIT_IN_SECONDS=60
INTER_STEP_WAIT_IN_SECONDS=120

# CentOS default 64 is not optimal for scale testing with 100 client nodes,
# as nodes beyond the threshold will be staggered until the previous batch is done,
# causing undesired synchronization issues.
CLUSH_MAX_FAN_OUT=120

# Put your own values in env.conf.override to:
#   - replace the dummy variables (must have)
#   - replace any tunable variables if desired (optional).
# Ensure all dependencies of the variable above this shall be updated as well.
#
# env.conf.override is ignored on commit by .gitignore,
# you don't want to push them to the public repo
_customized_config="$SCRIPT_DIR/env.conf.override"
if [ -f "$_customized_config" ]; then
    # shellcheck source=/dev/null
    source "$_customized_config"
fi

#==================== Derived Variables, No Changes Required ========================
=========#

HEAD_NODE="$(hostname -s)"

CLUSTER_DOMAIN_NAME=$(hostname | sed "s/$HEAD_NODE.//g")
```

```
_HOST_SUBNET_FIRST_VIP_10G=$(echo "$_VIP_POOL_10G" | cut -d'-' -f 1)
_HOST_SUBNET_LAST_VIP_10G=$(echo "$_VIP_POOL_10G" | cut -d'-' -f 2)
_HOST_SUBNET_FIRST_VIP_LAST_FIELD_10G=$(echo "$_HOST_SUBNET_FIRST_VIP_10G" | cut -d'.' -f 4)
_HOST_SUBNET_LAST_VIP_LAST_FIELD_10G=$(echo "$_HOST_SUBNET_LAST_VIP_10G" | cut -d'.' -f 4)
_HOST_SUBNET_10G=$(echo "$_HOST_SUBNET_FIRST_VIP_10G" | cut -d'.' -f 1-3)

_HOST_SUBNET_FIRST_VIP_100G=$(echo "$_VIP_POOL_100G" | cut -d'-' -f 1)
_HOST_SUBNET_LAST_VIP_100G=$(echo "$_VIP_POOL_100G" | cut -d'-' -f 2)
_HOST_SUBNET_FIRST_VIP_LAST_FIELD_100G=$(echo "$_HOST_SUBNET_FIRST_VIP_100G" | cut -d'.' -f 4)
_HOST_SUBNET_LAST_VIP_LAST_FIELD_100G=$(echo "$_HOST_SUBNET_LAST_VIP_100G" | cut -d'.' -f 4)
_HOST_SUBNET_100G=$(echo "$_HOST_SUBNET_FIRST_VIP_100G" | cut -d'.' -f 1-3)

NETWORK_SPEED="$network_speed"
_HOST_SUBNET_NAME=_HOST_SUBNET_"${NETWORK_SPEED}"G
eval HOST_SUBNET="\$$_HOST_SUBNET_NAME"

_HOST_SUBNET_FIRST_VIP_NAME=_HOST_SUBNET_FIRST_VIP_"${NETWORK_SPEED}"G
eval HOST_SUBNET_FIRST_VIP="\$$_HOST_SUBNET_FIRST_VIP_NAME"

_HOST_SUBNET_LAST_VIP_NAME=_HOST_SUBNET_LAST_VIP_"${NETWORK_SPEED}"G
eval HOST_SUBNET_LAST_VIP="\$$_HOST_SUBNET_LAST_VIP_NAME"

_HOST_SUBNET_FIRST_VIP_LAST_FIELD_NAME=_HOST_SUBNET_FIRST_VIP_LAST_FIELD_"${NETWORK_SPEED}"G
eval HOST_SUBNET_FIRST_VIP_LAST_FIELD="\$$_HOST_SUBNET_FIRST_VIP_LAST_FIELD_NAME"

_HOST_SUBNET_LAST_VIP_LAST_FIELD_NAME=_HOST_SUBNET_LAST_VIP_LAST_FIELD_"${NETWORK_
```

```
SPEED}”G
eval HOST_SUBNET_LAST_VIP_LAST_FIELD=”\$$_HOST_SUBNET_LAST_VIP_LAST_FIELD_NAME”

ALL_VIPS=()
for (( ip = “$HOST_SUBNET_FIRST_VIP_LAST_FIELD”; ip <= “$HOST_SUBNET_LAST_VIP_LAST_
FIELD”; ip++ ))
do
    ALL_VIPS+=(“$ip”)
done

_CLIENT_NODES_COMPACTED_NAME=”_CLIENT_NODES_${NETWORK_SPEED}G”
eval _CLIENT_NODES_COMPACTED=”\$$_CLIENT_NODES_COMPACTED_NAME”

if [ “$_CLIENT_NODES_COMPACTED” == “” ]; then
    _CLIENT_NODES_STRING=””
else
    _CLIENT_NODES_STRING=$(eval echo “$_CLIENT_NODES_COMPACTED”)
    if [ “$ORDERED_NODES” == “1” ]; then
        _CLIENT_NODES_STRING=$(nodeset --expand “$_CLIENT_NODES_STRING”)
    fi
fi

CLIENT_NODES_ARRAY=($_CLIENT_NODES_STRING)

CLIENT_NODES=${_CLIENT_NODES_STRING// /,}

_ALL_NODES_STRING=”$HEAD_NODE $_CLIENT_NODES_STRING”
ALL_NODES_ARRAY=($_ALL_NODES_STRING)
```

```
CLIENT_COUNT=$(echo "$_CLIENT_NODES_STRING" | wc --word)


if [ "$step_size" -le "0" ] || [ "$step_size" -gt "$CLIENT_COUNT" ]; then
    step_size="$CLIENT_COUNT"
fi
STEP_SIZE="$step_size"


if [ "$CLIENT_COUNT" -eq "0" ]; then
    export STEPS=0
else
    export STEPS=$(((CLIENT_COUNT + STEP_SIZE - 1) / STEP_SIZE))
fi
```

5.6.8 ERROR_CODE.CONF

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Most scripts just use the standard shell error code.

# Some defined their own error codes which are included in this file.


# Error code generated by run_test_no_runlog.sh

ERROR_CODE_MULTIPATH_MOUNT_FAILURE=80

ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_ALL=81

ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL=82

ERROR_CODE_INVALID_NETWORK_SPEED=83

ERROR_CODE_INVALID_SCHEDULE_FILE=84

ERROR_CODE_INVALID_STEPPING_TYPE=85

ERROR_CODE_INVALID_OPTION=86

ERROR_CODE_VIP_POOL_NOT_DEFINED=87

ERROR_CODE_CLIENTS_NOT_DEFINED=88

ERROR_CODE_MISSING_FIO_WORKLOAD_OPTIONS=89

ERROR_CODE_MISSING_ELBENCHO_FILE_WORKLOAD_OPTIONS=90

ERROR_CODE_MISSING_ELBENCHO_S3_WORKLOAD_OPTIONS=91

ERROR_CODE_FAIL_TO_GOTO_DIR=92

ERROR_CODE_FAIL_TO_EXIT_DIR=93

ERROR_CODE_CLIENT_NOT_REACHABLE=94


# Error code generated by nfs_mount.sh

ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE=110 # For hard mount, sadly timeo doesn't work,
so it takes a few minutes to error out

ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NUMBER_OF_INPUTS=111
```

```
ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_MOUNT_BASE=112
ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_SUBNET=113
ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NCONNECT_NUM=114
ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_RANDOMNIZE=115
ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_VIPS=116
```

5.6.9 GENERAL_INFO.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

HOST_SUBNET=$1   # eg: 192.168.B
NUMA_TOPO_OUTPUT_DIR=$2

if [ "$(ifconfig | grep "$HOST_SUBNET")" == "" ]; then
    echo "Invalid host subnet: $HOST_SUBNET"
    exit 22 # Standard EINVAL for Invalid Argument
fi

if [ ! -d "$NUMA_TOPO_OUTPUT_DIR" ]; then
    echo "$NUMA_TOPO_OUTPUT_DIR not found!"
    exit 22 # Standard EINVAL for Invalid Argument
fi

date
hostname
uname -r
cat /etc/centos-release
```

```
# Query memory information
echo "Memory information:"
cat /proc/meminfo

# Query CPU information
echo "CPU information by lscpu:"
lscpu

echo "CPU information by /proc/cpuinfo:"
cat /proc/cpuinfo

# More detailed turbo boost information
cpupower frequency-info

# Query NIC card info
echo "Network information:"
sudo lshw -class network

# Query the networking configuration
echo "ifconfig information:"
ifconfig

# Query the NIC FW version information, i.e. driver and FW version
# Strip off the color control charactors at the beginning with sed
echo "NIC information:"
network_interface=$(sed -r "s:\x1B\[[0-9;]*[mK]::g" <<< "$(sudo ip -br -c addr show |
grep --color=never "$HOST_SUBNET" | cut -d' ' -f 1)")
ethtool -i "$network_interface"

# Query Numa topology
echo "Numa topology:"
```

```
lstopo -l


if [ "$NUMA_TOPO_OUTPUT_DIR" != "" ]; then
    lstopo --logical --output-format png > "$NUMA_TOPO_OUTPUT_DIR/numa.png"
fi


# Query Numa statistics
echo "Numa stats:"
numastat -c -z -m -n
```

5.6.10 INSTALL_DEP.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


network_speed=${1:-"10"}


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" "$network_speed"


# Install clustershell on current node.
sudo install clustershell -y


# Install clustershell on all nodes.
# clustershell and jq installation are only required by harness on the FIO client node,
# However, installation on all nodes allows any node to be used as FIO client node.
clush -w "$CLIENT_NODES" "sudo yum install clustershell"
```

```
# Install other tools needed by the harness.
# lm_sersors not used by harness yet, for further enhancements, it can be used to query temperature during test with "sensors" command.
clush -w "$HEAD_NODE,$CLIENT_NODES" "sudo yum install fio iperf3 jq nfs-utils hwloc-libs hwloc-gui lshw ethtool lm_sensors irqbalance -y"

# Pull from head node only and use elbencho_deploy.sh to deploy instead.
# For users with unlimited docker pull, direct pull from all nodes can be used.

# clush -w $HEAD_NODE,$CLIENT_NODES "docker pull breuner/elbencho"
docker pull breuner/elbencho

# Copy nfs_mount.sh to all clients.
# This is required for nfs mount with "MULTIPATH_ENABLED=0" in env.conf.

clush -w "$CLIENT_NODES" -c nfs_mount.sh --dest "$VAST_SCALE_TESTING_PATH/nfs_mount.sh"
clush -w "$CLIENT_NODES" -c error_code.conf --dest "$VAST_SCALE_TESTING_PATH/error_code.conf"

# Copy mtu_update.sh to all clients.

clush -w "$CLIENT_NODES" -c mtu_update.sh --dest "$VAST_SCALE_TESTING_PATH/mtu_update.sh"

# Copy general_info.sh to head_node and all clients.

clush -w "$CLIENT_NODES" -c general_info.sh --dest "$VAST_SCALE_TESTING_PATH/general_info.sh"

# Install the multipath driver.
# This is required for nfs mount with "MULTIPATH_ENABLED=1" in env.conf.
#
```

```
# Preconditions:
#    1. The driver is kernel version dependent.
#       Obertain the version matching your kernel version from customer.support@vastdata.com if needed.
#    2. "mkdir ../drivers" and place the driver in the folder
#    3. Update MULTIPATH_DRIVER below with the driver name
#
MULTIPATH_DRIVER="mlnx-nfsrdma-vast_3.9.3.for.4.18.0.240.x.centos-kernel_4.18.0_240.22.1.el8_3.x86_64.rpm"

if [ -f "../drivers/$MULTIPATH_DRIVER" ]; then

    clush -w "$HEAD_NODE,$CLIENT_NODES" "mkdir -p $VAST_SCALE_TESTING_PATH/drivers"
    clush -w "$HEAD_NODE,$CLIENT_NODES" -c "../drivers/$MULTIPATH_DRIVER" --dest "$VAST_SCALE_TESTING_PATH/drivers"
    clush -w "$HEAD_NODE,$CLIENT_NODES" "sudo rpm -i $VAST_SCALE_TESTING_PATH/drivers/$MULTIPATH_DRIVER; sudo dracut -f"

    # The new driver will only take effect after the reboot
    # The reboot is commented out in case another convenient time if preferred for reboot.

    # clush -w $CLIENT_NODES "sudo reboot"
    # sudo reboot
fi
```

5.6.11 MTU_UPDATE.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause
```

```
# If the interface is not eth0, specify it in command line argument
#**************
# Important   *
#**************
#  (1)   Please adapt to the proper interfaces if they are different
#  (2)   Ensure switch side connected to this interface supports 9000 too,
#       otherwise the system may not be reachable after this update

network_interface=${1,-"eth0"}

if [ ! -f /etc/sysconfig/network-scripts/ifcfg-"$network_interface" ]; then
    echo "/etc/sysconfig/network-scripts/ifcfg-$network_interface not found!"
    exit 22 # Standard EINVAL for Invalid Argument
fi

echo "MTU=9000" | sudo tee -a /etc/sysconfig/network-scripts/ifcfg-"$network_interface"
sudo systemctl restart NetworkManager

# Commented out as connection could be broken for clush when NetworkManager restarts
#cat /etc/sysconfig/network-scripts/ifcfg-"$network_interface"
```

5.6.12 NFS_MOUNT.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/error_code.conf" # This file shall be copied to all clients, in the
same folder as nfs_mount.sh
```

```
MOUNT=$1
REMOTE_PATH=$2
HOST_SUBNET=$3
NCONNECT_NUM=$4
randomnize=$5
nfs_mount_log_file=$6
all_vips=("${@:7}")  # This is array
NCONNECT_NUM_MIN=0
NCONNECT_NUM_MAX=48
uid=$UID

# Validate the inputs
return_code=0
if [ "$#" -le 6 ]; then
    echo "NFS mount failure, invalid number of inputs. Valid parameters: [MOUNT REMOTE_PATH HOST_SUBNET NCONNECT_NUM RANDOMNIZATION VIPs]. Exiting..."
    return_code=$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NUMBER_OF_INPUTS
elif [[ "$MOUNT" != "/mnt/"* ]]; then
    echo "NFS mount failure, mount base not in /mnt/* format, exiting..."
    return_code=$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_MOUNT_BASE
elif [[ "$HOST_SUBNET" != *"."*"."* ]]; then
    echo "NFS mount failure, invalid subnet, must be in format A.B.C, exiting..."
    return_code=$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_SUBNET
elif [ "$NCONNECT_NUM" -lt $NCONNECT_NUM_MIN ] || [ "$NCONNECT_NUM" -gt $NCONNECT_NUM_MAX ]; then
    echo "NFS mount failure, invalid nconnect number, must be in range [$NCONNECT_NUM_MIN, $NCONNECT_NUM_MAX], exiting..."
    return_code=$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NCONNECT_NUM
elif [ "$randomnize" != "yes" ] && [ "$randomnize" != "no" ]; then
    echo "NFS mount failure, invalid randomization, must be in [yes, no], exiting..."
```

```
    return_code=$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_RANDOMNIZE
fi

if [ "$return_code" != "0" ]; then
    echo "Error code: $return_code" > "$nfs_mount_log_file"
    exit "$return_code"
fi

HN=$(hostname -s)
if [ "$randomnize" == "yes" ]; then
    vip_list=$(shuf -e "${all_vips[@]}" | xargs)
    IFS=" " read -r -a all_vips <<< "$vip_list"
fi

if [[ ! -L "$nfs_mount_log_file" ]] && [ -f "$nfs_mount_log_file" ]; then
    rm -rf "$nfs_mount_log_file"
fi

for MOUNTVIP in "${all_vips[@]}"
do
    remote_address="${HOST_SUBNET}.${MOUNTVIP}:/${REMOTE_PATH}"
    mountpoint_parent="${MOUNT}/${REMOTE_PATH}/${HN}"
    mountpoint="${mountpoint_parent}/${HOST_SUBNET}.${MOUNTVIP}"
    sudo mkdir -p "$mountpoint_parent"
    sudo chown $uid:$uid "$mountpoint_parent"
    sudo mkdir -p "$mountpoint"
    sudo chown $uid:$uid "$mountpoint"

    if [ "$NCONNECT_NUM" == "0" ]; then
        sudo mount -t nfs -o vers=3,tcp "$remote_address" "$mountpoint" 2> "$nfs_
mount_log_file"
```

```
    else
        sudo mount -t nfs -o vers=3,tcp,nconnect="${NCONNECT_NUM}" "$remote_address"
"$mountpoint" 2> "$nfs_mount_log_file"
    fi

    if [[ ! -L "./$nfs_mount_log_file" ]] && [ -s "./$nfs_mount_log_file" ]; then
        if grep -q "exited with exit code\|timed out\|Failed\|failed" "./$nfs_mount_
log_file"; then
            echo "Mount of $remote_address to $mountpoint failed:"
            cat "$nfs_mount_log_file" # No cleanup of this file until the next run
            exit "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE"
        fi
    fi
done
```

5.6.13 NIC_INFO.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

network_speed=${1:-"10"}

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" "$network_speed"

query_network_interface_cmd="sed -r 's:\x1B\[[0-9;]*[mK]::g' <<< \$(sudo ip -br -c addr
show | grep --color=never $HOST_SUBNET | cut -d' ' -f 1)"

echo "Querying NIC type:"
clush -w "$CLIENT_NODES" "network_interface=\$($query_network_interface_cmd); sudo lshw
-class network | grep -B 10 \$network_interface | grep product" | sort
```

```
echo "Querying NIC Firmware Version"
clush -w "$CLIENT_NODES" "network_interface=\$($query_network_interface_cmd); ethtool
-i \$network_interface | grep firmware-version" | sort
echo "Querying MTU:"
clush -w "$CLIENT_NODES" "network_interface=\$($query_network_interface_cmd); ifconfig |
grep \$network_interface" | sort
echo "NUMA binding:"  # BIOS determines if enabled or not, -1 means disabled.
Installation determines which NUMA is at, if enabled in BIOS.
clush -w "$CLIENT_NODES" "network_interface=\$($query_network_interface_cmd); cat /sys/
class/net/\$network_interface/device/numa_node" | sort
```

5.6.14 PURGE_CLUSTER.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" 10 # Always run from 10g interface

sudo umount /mnt/all > /dev/null 2>/dev/null
sudo mount "$_HOST_SUBNET_FIRST_VIP_10G:"/ /mnt/all -t nfs
pushd /mnt/all || exit

file .vast_trash
if [ "$REMOTE_PATH_FIO" != "" ]; then
    if [ -d "/mnt/all/$REMOTE_PATH_FIO" ]; then
        sudo mv /mnt/all/"${REMOTE_PATH_FIO}"/* /mnt/all/.vast_trash                > /
dev/null 2>/dev/null
    fi
```

```
    # Todo: MULTIPATH_ENABLE=0 currently requires remote path to exist first,
    # so use this as the workaround to always ensure that's true.

    # Should consider mount differently and remove this requirement later
    sudo mkdir -p /mnt/all/$REMOTE_PATH_FIO
fi

if [ "$REMOTE_PATH_ELBENCHO_FILE" != "" ]; then
    if [ -d "/mnt/all/$REMOTE_PATH_ELBENCHO_FILE" ]; then
        sudo mv /mnt/all/"${REMOTE_PATH_ELBENCHO_FILE}"/* /mnt/all/.vast_trash      > /
dev/null 2>/dev/null
    fi
    sudo mkdir -p /mnt/all/$REMOTE_PATH_ELBENCHO_FILE
fi

if [ "$REMOTE_PATH_ELBENCHO_S3" != "" ]; then
    if [ -d "/mnt/all/$REMOTE_PATH_ELBENCHO_S3" ]; then
        sudo mv /mnt/all/"${REMOTE_PATH_ELBENCHO_S3}"/* /mnt/all/.vast_trash        > /
dev/null 2>/dev/null
    fi
    sudo mkdir -p /mnt/all/$REMOTE_PATH_ELBENCHO_S3
fi

# For easier cleanup, it is recommended that schedule files in workloads/elbencho/s3
# to create s3 bucket with elbencho-s3-bucket* naming convention.
sudo mv /mnt/all/elbencho-s3-bucket* /mnt/all/.vast_trash  > /dev/null 2>/dev/null

popd || exit
```

5.6.15 RUN_BW_TEST.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Ensure the cluster is 99% free and Auxiliary is o or close to 0 before start ( < 4T
perferred)
# The test will write ~35TB of data, and it takes 10~15 minutes to be truly purged
# Sleep after the purge to ensure all tests start with the same condition

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"

iterations=3
sleep_in_seconds=900 # Need longer wait?

purge_and_wait() {
    "$SCRIPT_DIR"/purge_cluster.sh
    sleep "$sleep_in_seconds"
}

purge_only() {
    "$SCRIPT_DIR"/purge_cluster.sh
}

for i in $(seq 1 $iterations)
do
    prefix="run$i"

    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/bw_test_100g.sch -pf
```

```
"$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/file/bw_test_100g.sch
-pf "$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/bw_test_100g.sch
-pf "$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/bw_test_10g.sch -pf
"$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/file/bw_test_10g.sch
-pf "$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/s3/bw_test_10g.sch
-pf "$prefix"

    if [ "$i" == "$iterations" ]; then
        purge_only
    else
        purge_and_wait
    fi
done
```

5.6.16 RUN_ELBENCHO_FILE_SWEEP.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
```

```
source "$script_dir/common.conf"


"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/file/sweep_random_
threads_and_iodepth_1024k_100g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/file/sweep_random_
iosizes_100g.sch


"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/file/
sweep_random_threads_and_iodepth_1024k_10g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/file/sweep_random_
iosizes_10g.sch
```

5.6.17 RUN_ELBENCHO_S3_SWEEP.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"


"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
```

```
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_threads_
random_32m_100g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_iosizes_
random_100g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_threads_
random_4k_100g.sch


"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_threads_
random_32m_10g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_iosizes_
random_10g.sch
"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_threads_
random_4k_10g.sch
```

5.6.18 RUN_ELBENCHO_S3_SWEEP_IOSIZES.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"
```

```
iterations=3
sleep_in_seconds=600

purge_and_wait() {
    "$SCRIPT_DIR"/purge_cluster.sh
    sleep "$sleep_in_seconds"
}

purge_only() {
    "$SCRIPT_DIR"/purge_cluster.sh
}

for i in $(seq 1 $iterations)
do

    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/sweep_iosizes_
random_100g.sch -pf "run$i"
    purge_and_wait

    if [ "$i" == "$iterations" ]; then
        purge_only
    else
        purge_and_wait
    fi

done
```

5.6.19 RUN_FIO_BW_TEST_SHORT.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
```

```
# SPDX-License-Identifier: BSD-3-Clause


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"


"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/bw_test_short_100g.sch -co


# Run it 3 times
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/bw_test_short_100g.sch -ss 2
-pf run1
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/bw_test_short_100g.sch -ss 2
-pf run2
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/bw_test_short_100g.sch -ss 2
-pf run3
"$SCRIPT_DIR"/purge_cluster.sh
```

5.6.20 RUN_FIO_SWEEP.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"


"$SCRIPT_DIR"/purge_cluster.sh
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/sweep_workers_and_iodepth_
random_1024k_100g.sch -co
sleep 300
"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/sweep_workers_and_iodepth_
```

```
random_1024k_100g.sch

"$SCRIPT_DIR"/purge_cluster.sh

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/sweep_iosizes_100rd_
random_100g.sch -co

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/sweep_iosizes_100rd_
random_100g.sch


"$SCRIPT_DIR"/purge_cluster.sh

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/sweep_workers_and_iodepth_
random_1024k_10g.sch -co

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/sweep_workers_and_iodepth_
random_1024k_10g.sch

"$SCRIPT_DIR"/purge_cluster.sh

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/sweep_iosizes_100rd_random_10g.
sch -co

sleep 300

"$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/sweep_iosizes_100rd_random_10g.
sch
```

5.6.21 RUN_IOPS_TEST.SH

```
#!/bin/bash


# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Ensure the cluster is 99% free and Auxiliary is o or close to 0 before start ( < 4T
perferred)
```

```
# The test will write ~12TB of data, and it takes 5~10 minutes to be truly purged
# Sleep after the purge to ensure all tests start with the same condition

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"

iterations=3
sleep_in_seconds=600

purge_and_wait() {
    "$SCRIPT_DIR"/purge_cluster.sh
    sleep "$sleep_in_seconds"
}

purge_only() {
    "$SCRIPT_DIR"/purge_cluster.sh
}

for i in $(seq 1 $iterations)
do
    prefix="run$i"

    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/fio/iops_test_100g.sch -pf
"$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/file/iops_test_100g.
sch -pf "$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 100 -sf "$WORKLOAD_DIR"/elbencho/s3/iops_test_100g.
sch -pf "$prefix"
    purge_and_wait
```

```
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/fio/iops_test_10g.sch -pf
"$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/file/iops_test_10g.sch
-pf "$prefix"
    purge_and_wait
    "$SCRIPT_DIR"/run_test.sh -ns 10 -sf "$WORKLOAD_DIR"/elbencho/s3/iops_test_10g.sch
-pf "$prefix"

    if [ "$i" == "$iterations" ]; then
        purge_only
    else
        purge_and_wait
    fi
done
```

5.6.22 RUN_IPERF_TEST.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# How to use the script
# Precondition: iperf3 installed on both client and VAST Data side
# Steps to test 10G performance:
#   Step 1: log onto VAST VMS node, note down one of the VIP on the VMS node, eg
10.A.B.E
#   Step 2: start iperf as the service on the VMS node
#      iperf3 -s -p5900
#   Step 3: run iperf test script on the client:
#      ./run_iperf_test.sh 10 10.A.B.E
```

```
# Steps to test 100G performance:
#   Step 1: log onto VAST VMS node, note down one of the VIP on the VMS node, eg 192.168.B.F
#   Step 2: start iperf as the service on the VMS node, but bind iperf3 to the 100GbE interface IP
#      iperf3 -s -p5900 -B 192.168.B.F
#   Step 3: on client 101, run iperf test script:
#      ./run_iperf_test.sh 100 192.168.B.F

network_speed=${1:-"10"}

# Replace the IP accordingly or provide it in command line
vip=${2:-"10.A.B.E"}

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/env.conf" "$network_speed"

echo "$HOST_SUBNET"

for i in "${CLIENT_NODES_ARRAY[@]}"; do clush -w "$i" "iperf3 -c $vip -B \$(ifconfig | grep inet | grep $HOST_SUBNET | xargs | awk '{print \$2}') -p5900 -t 5"; done
```

5.6.23 RUN_TEST.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Any inputs will be passed to run_test_no_runlog.sh transparently hence validated there
```

```
echo "$0" "$@"

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"

# script_dir and SCRIPT_DIR are the same, but SCRIPT_DIR is preferred after sourcing
common.conf
runner="$SCRIPT_DIR/run_test_no_runlog.sh"

# Ctrl + C abort handler
function abort_handler() {
    echo "Test aborted with CTRL C, cleaning up..."
    end_of_test
}

# Abort takes the same steps to end the test as normal execution
function end_of_test() {

    if [[ ! -L "$RUN_DIR_TEMP_FILE" ]] && [ -f "$RUN_DIR_TEMP_FILE" ]; then

        while IFS= read -r run_dir
        do
            if [ -d "$run_dir" ]
            then
                if [[ ! -L "$RUN_LOG_FILE" ]] && [ -f "$RUN_LOG_FILE" ]; then
                    # Abort case will leave run log in staging folder, move it to the
result folder
                    mv "$RUN_LOG_FILE" "$run_dir"
                fi
                if [[ ! -L "$RUN_ERROR_LOG_FILE" ]] && [ -f "$RUN_ERROR_LOG_FILE" ];
then
```

```
                    # Abort case will leave run error log in staging folder, move it to
the result folder
                    mv "$RUN_ERROR_LOG_FILE" "$run_dir"
                fi
            fi
        done < "$RUN_DIR_TEMP_FILE"

        rm -rf "$RUN_DIR_TEMP_FILE"
    fi

    [ -n "$return_code" ] && [ "$return_code" -eq "$return_code" ] 2>/dev/null
    status="$?"
    if [ "$status" -eq 0 ]; then
       # return_code is a number
       exit "$return_code"
    fi
}

# Trap ctrl-c to cleanup before exit
trap abort_handler INT

# Remove the temporary file first
if [[ ! -L "$RUN_DIR_TEMP_FILE" ]] && [ -f "$RUN_DIR_TEMP_FILE" ]; then
    rm -rf "$RUN_DIR_TEMP_FILE"
fi

# Execute the runner only if it's not a symbolic link and if it exists
if [[ ! -L "$runner" ]] && [ -f "$runner" ]; then
    # Error code from runner not handled other than being provided as the return code,
    # this is because we always exit if anything goes wrong.
    set -o pipefail
```

```
    $runner "$@" 2>&1 | tee -a "$RUN_LOG_FILE"
    return_code=$?
fi

# Disable ctrl + c, so that we always do proper cleanup
trap '' INT
end_of_test
```

5.6.24 RUN_TEST_NO_RUNLOG.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Enable xtrace mode, which logs every command as it runs.
set -x

script_dir="$(cd "$(dirname "$0")" || exit; pwd)"
source "$script_dir/common.conf"
source "$script_dir/error_code.conf"

network_speed="10"
step_size=""
cmdline_create_only=0
runtime=""
ramptime=""
cmdline_runtime=""
cmdline_ramptime=""
skip_mount=0
```

```
common_prefix="results_"$(date +%Y-%m-%d-%H.%M.%S)
schedule_file=""
test_type="fio"
REMOTE_PATH=""
fio_template="fio_template.ini"
extra_prefix=""
STEP_TYPE_INCREMENTAL=0
STEP_TYPE_CONSTANT=1
step_type=$STEP_TYPE_INCREMENTAL
dry_run=0
inject_error_num=0
skip_all_client_config_query=0
uid="$UID"

# Additional globals defined by run_test() function
# schedule_name schedule_name_ext hostname client_nodes client_nodes_array result_prefix run_dir

#=====================================================================
# Print the help message.
#=====================================================================
print_help () {
    echo " Usage: $0 [ options ]"
    echo " Valid options:"
    echo
    echo " -co|--create-only: when specified, create the FIO files only for read workload, no actual read access."
    echo " -dr|--dryrun: do not run workload. Everything else will be executed the same as a normal run including nfs mounts and starting fio and Elbencho services on clients."
    echo " -ns|--network-speed: GbE speed in [\"10\", \"100\"]. Default: 10"
    echo " -pf|--prefix: extra prefix to be added to the test name"
```

```
    echo " -sf|--schedule-file: specify the schedule file to run."
    echo " -sm|--skip-mount: specify this to skip the mount step. This can speed up the test if the mount is the same as the previous test."
    echo " -sq|--skip-all-client-config-query: do not do query of all client configurations before the test. As query takes a few minutes, this allows the test to be started sooner for quicker benchmarking."
    echo " -ss|--step-size: specify the number of clients to increase to run the test."
    echo "    Default: max clients. Any value outside of [1..max client count] will have step size of max client count."
    echo "    Examples:"
    echo "      when st==0:"
    echo "          if max clients = 100, then \"-ss 10\" will test clients in steps of 10, 20, 30...90, 100 clients"
    echo "          if max clients = 100, then \"-ss 15\" will test clients in steps of 15, 30, 45...75, 90, 100 clients"
    echo "          if max clients = 100, then \"-ss 0\" or \"-ss 100\" or omitting -ss option will test all 100 clients in 1 step"
    echo "      when st==1:"
    echo "          if max clients = 10,  then \"-ss 1 -st 1\" will test the 10 clients 1 at a time, in 10 steps"
    echo "          if max clients = 10,  then \"-ss 2 -st 1\" will test the 10 clients 2 at a time, in 5 steps"
    echo " -st|--step-type: 0 - incremental (default), 1 - constant. Refer to -ss|--step-size examples to see the differences between the two types"
    echo " --runtime:  FIO runtime. Default: as specified in template FIO job file"
    echo " --ramptime: FIO ramptime. Default: as specified in template FIO job file"
    echo " -h|--help: to print this help"
    echo
    echo " Examples:"
    echo "     # Run FIO BW test on 100G"
    echo "     ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_100g.sch"
    echo "     # Create FIO BW short test files only on 100G"
    echo "     ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_short_100g.sch -co"
```

```
    echo "     # Run FIO BW short test with 120s of runtime and 0s of ramptime, instead
of the values as specified in the schedule file"
    echo "     ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_short_100g.sch
--runtime 120 --ramptime 0"
    echo "     # Run constant type of stepping test, with step size of 1, meaning test
each client 1 by 1"
    echo "     ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_short_100g.sch -ss 1
-st 1"
    echo "     # Run incremental stepping test, with step size of 2"
    echo "     ./run_test.sh -ns 100 -sf ../workloads/fio/bw_test_short_100g.sch -ss 2"
}

#=================================================================
# Parse the command line and set global variables accordingly.
#=================================================================
parse_command_line () {
    # Each client will add its own timestamp as suffix, which could be different from
client to client

    # Supply a common timestamp as the prefix so that we can easily identify if the
results are from
    # the same run with the clush command

    while [ $# -gt 0 ]; do
      case "$1" in
        -co|--create_only)
            cmdline_create_only=1
            shift 1
        ;;
        -dr|--dryrun)
            dry_run=1
            shift 1
```

```
        ;;
        -ei|--error-injection)
            inject_error_num="$2" # This shall be the error code for injection
            shift 2
        ;;
        -ns|--network-speed)
            network_speed="$2"
            if [ "$network_speed" != "10" ] && [ "$network_speed" != "100" ]; then
                echo "Invalid value for network speed, must be in [10, 100]"
                exit "$ERROR_CODE_INVALID_NETWORK_SPEED"
            fi
            shift 2
        ;;
        -pf|--prefix)
            extra_prefix="$2"
            shift 2
        ;;
        -sf|--schedule-file)
            schedule_file_dir="$(cd "$(dirname "$2")" || exit; pwd)"
            schedule_file="$schedule_file_dir/$(basename "$2")" # Always convert to
absolution path
            if [[ -L "$schedule_file" ]] || [ ! -f "$schedule_file" ]; then
                echo "$schedule_file not found or it's a symbolic link, or not in
$WORKLOAD_DIR!"
                exit "$ERROR_CODE_INVALID_SCHEDULE_FILE"
            fi
            shift 2
        ;;
        -sq|--skip-all-client-config-query)
            skip_all_client_config_query=1
            shift 1
```

```
        ;;
        -sm|--skip-mount)
            skip_mount=1
            shift 1
        ;;
        -ss|--step-size)
            step_size="$2"
            # Range bounded by env.conf
            shift 2
        ;;
        -st|--step-type)
            step_type="$2"
            if [ "$step_type" != "0" ] && [ "$step_type" != "1" ]; then
                echo "Invalid value for step type, must be in [0, 1]"
                exit "$ERROR_CODE_INVALID_STEPPING_TYPE"
            fi
            shift 2
        ;;
        --runtime)
            cmdline_runtime="$2"
            shift 2
        ;;
        --ramptime)
            cmdline_ramptime="$2"
            shift 2
        ;;
        -h|--help)
            print_help "$0"
            echo " Usage: $0 [ options ]"
            exit
        ;;
```

```
        *)
          echo "Unrecongized option: $1!"
          print_help "$0"
          exit $ERROR_CODE_INVALID_OPTION
      esac
    done
}

#=====================================================================
# Set additional globals based on the configuration file env.conf.
#=====================================================================
parse_env_and_set_globals () {

    source "$SCRIPT_DIR/env.conf" "$network_speed" "$step_size"

    # Sanity check of the configurations

    if [ "$HOST_SUBNET" == "" ]; then
        echo "VIP Pool not defined"
        exit "$ERROR_CODE_VIP_POOL_NOT_DEFINED"
    fi

    if [ ${#CLIENT_NODES_ARRAY[@]} -eq 0 ]; then
        echo "Clients not defined"
        exit "$ERROR_CODE_CLIENTS_NOT_DEFINED"
    fi

    MOUNT="$MOUNT_MULTIPATH_ENABLED"
    if [ "$MULTIPATH_ENABLE" == "0" ]; then
        MOUNT="$MOUNT_MULTIPATH_DISABLED"
    fi
```

```
    if [[ "$schedule_file" == *"elbencho/file"* ]]; then
        test_type="elbencho-file"
        REMOTE_PATH="$REMOTE_PATH_ELBENCHO_FILE"
    elif [[ "$schedule_file" == *"elbencho/s3"* ]]; then
        test_type="elbencho-s3"
        REMOTE_PATH="$REMOTE_PATH_ELBENCHO_S3"
    else
        test_type="fio"
        REMOTE_PATH="$REMOTE_PATH_FIO"
    fi

    echo "Test type: $test_type"



}

#====================================================================
# Enable/disable irqbalance
#====================================================================
function check_client_status_and_deploy_files () {

    echo "Sanity check of the client nodes' presence..."

    rm -rf "$RUN_ERROR_LOG_FILE"
    clush -w "$CLIENT_NODES" "hostname -s" > /dev/null 2> "$RUN_ERROR_LOG_FILE"

    # This is the first clush, and currently the only place write to error log.
    if [[ ! -L "$RUN_ERROR_LOG_FILE" ]] && [ -s "$RUN_ERROR_LOG_FILE" ]; then
        if grep -q 'exited with exit code\|timed out\|Failed\|failed' "$RUN_ERROR_LOG_
FILE"; then
            cat "$RUN_ERROR_LOG_FILE"
```

```
            exit "$ERROR_CODE_CLIENT_NOT_REACHABLE"
        fi;
    fi

    # Utility shall deploy them, but in case we forget to do that after changing these files.
    clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "mkdir -p $VAST_SCALE_TESTING_PATH"
    clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" -c "$SCRIPT_DIR/nfs_mount.sh" "$SCRIPT_DIR/general_info.sh" "$SCRIPT_DIR/error_code.conf" --dest "$VAST_SCALE_TESTING_PATH"
}

#=====================================================================
# Enable/disable irqbalance
#=====================================================================
function irqbalance_config() {

    if [ "$IRQBALANCE_ENABLE" == "1" ]; then
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "sudo systemctl enable irqbalance; sudo systemctl start irqbalance"
    else
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "sudo systemctl stop irqbalance; sudo systemctl disable irqbalance"
    fi
}

#=====================================================================
# Enable/disable Intel Turbo Boost
#=====================================================================
function turbo_boost_config () {
```

```
    if [ "$TURBO_BOOST_ENABLE" == "1" ]; then
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "echo 0 | sudo tee /sys/devices/system/cpu/intel_pstate/no_turbo"
    else
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "echo 1 | sudo tee /sys/devices/system/cpu/intel_pstate/no_turbo"
    fi
}

#====================================================================
# Enable/disable Hyper Threading
#====================================================================
function hyper_threading_config() {

    if [ "$HYPER_THREADING_ENABLE" == "1" ]; then
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "echo on | sudo tee /sys/devices/system/cpu/smt/control"
    else
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "echo off | sudo tee /sys/devices/system/cpu/smt/control"
    fi
}

#====================================================================
# Config irqbalance/turbo boost and hyper threading
#====================================================================
function set_irqb_tb_ht_controls() {
    irqbalance_config > /dev/null
    turbo_boost_config > /dev/null
    hyper_threading_config > /dev/null
}
```

```
#=================================================================
# Query irqbalance/turbo boost and hyper threading status
#=================================================================
function get_all_client_config() {

    # Todo: not optimized as it does dmidecode / lscpu etc multiple times

    echo "Kernel:"
    clush -w "$CLIENT_NODES" "uname -r" | sort

    echo "OS:"
    clush -w "$CLIENT_NODES" "cat /etc/centos-release" | sort

    echo "Platform:"
    clush -w "$CLIENT_NODES"  "sudo dmidecode | grep -m 1 Product" | sort

    echo "BIOS:"
    clush -w "$CLIENT_NODES" "sudo dmidecode | grep -A 20 'BIOS Information' | grep
Version" | sort
    echo "BIOS grouped:"
    clush --diff -w "$CLIENT_NODES" "sudo dmidecode | grep -A 20 'BIOS Information' |
grep Version" | sort

    echo "CPU information:"
    clush -w "$CLIENT_NODES" "grep -m 3 'stepping\|model name\|microcode' /proc/
cpuinfo" | sort
    echo "Cores per socket:"
    clush -w "$CLIENT_NODES" "lscpu | grep 'Core(s)'" | sort
    echo "Total sockets"
    clush -w "$CLIENT_NODES" "lscpu | grep 'Socket(s)'" | sort
    echo "Threads per core"
```

```
    clush -w "$CLIENT_NODES" "lscpu | grep 'Thread(s)'" | sort
    echo "Virtualization"
    clush -w "$CLIENT_NODES" "lscpu | grep 'VT-x'" | sort
    echo "Total CPU(s):" # = Cores per socket x threads per core x sockets. Threads per
core = 1 if Hyper Threading is disabled.
    clush -w "$CLIENT_NODES" "getconf _NPROCESSORS_ONLN" | sort


    echo "Total memory:"
    clush -w "$CLIENT_NODES" "free -m | grep Mem:" | sort
    echo "Memory speed:"
    clush -w "$CLIENT_NODES" "sudo dmidecode | grep -C 20 DDR | grep  Speed: | tail -1"
| sort
    echo "Memory speed grouped:"
    clush --diff -w "$CLIENT_NODES" "sudo dmidecode | grep -C 20 DDR | grep  Speed: |
tail -1" | sort


    echo "irqbalance status:"
    clush -w "$CLIENT_NODES" "echo irqb:     \$(systemctl status irqbalance | grep
Active)" | sort
    echo "Turbo Boost status:"
    clush -w "$CLIENT_NODES" "echo no_turbo: \$(cat /sys/devices/system/cpu/intel_
pstate/no_turbo)" | sort
    echo "Hyper Threading status"
    clush -w "$CLIENT_NODES" "echo ht:       \$(cat /sys/devices/system/cpu/smt/
active)" | sort


    echo "Power & Performance policy (core 0 only):"
    # Query CPU performance mode
    #
    # Quotes http://manpages.ubuntu.com/manpages/bionic/man8/x86_energy_perf_
policy.8.html:
    #
```

```
    #   The following table shows the mapping from the value strings above to actual
MSR  values.
    #   This mapping is defined in the Linux-kernel header, msr-index.h.
    #
    #       VALUE STRING        EPB  EPP
    #       performance         0    0
    #       balance-performance 4    128
    #       normal, default     6    128
    #       balance-power       8    192
    #       power               15   255
    clush -w "$CLIENT_NODES" "sudo x86_energy_perf_policy | grep -m 1 cpu" | sort
    echo "Per CPU Power & Performance policy (core 0 only):"
    clush -w "$CLIENT_NODES" "cat /sys/devices/system/cpu/cpu0/cpufreq/scaling_
governor" | sort   # This is per CPU, usually they are the same

    echo "Network interface and MTU"
    "$SCRIPT_DIR/nic_info.sh" "$network_speed"
}

#=====================================================================
# Mount the VIPs individually
# NConnect is available is also available in CentOS Linux release 8.3.2011.
# Warning: this can be slow and the connection might be dropped if
#          the VIP NConnect sizes are big as:
#          total tcp connection = VIP size x NConnect size
#=====================================================================
mount_nfs_no_multipath () {

    if [ "$1" == "all" ]; then
        clients="$CLIENT_NODES"
        vips_to_mount=("${ALL_VIPS[@]}")
```

```
    else
        clients="$1"
        vips_to_mount=("$2")
    fi

    if [ "$NCONNECT_ENABLE" == "0" ]; then
        nconnect_num=0
    else
        nconnect_num="$NCONNECT_NUM"
    fi

    randomnize="no"
    if [ "$1" == "all" ]; then
        clush -w "$clients" -f "$CLUSH_MAX_FAN_OUT" "sudo umount ${MOUNT}/${REMOTE_
PATH}/\$(hostname -s)/* > /dev/null 2>/dev/null &" &

        wait
        if [ "$MOUNT_ALL" == "1" ] || [ "$test_type" == "elbencho-file" ];
        then
            if [ "$test_type" == "elbencho-file" ]; then
                randomnize="yes"
            fi

            if [ "$inject_error_num" == "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_ALL" ];
then
                # Mess up randomnize to force the error
                randomnize="yesno"
            fi

            # Mount may fail, but clush return code is 0. So capture the stderr result
to know the actual error instead.
```

```
            clush -w "$clients" -f "$CLUSH_MAX_FAN_OUT" "cd $VAST_SCALE_TESTING_PATH;
./nfs_mount.sh $MOUNT $REMOTE_PATH $HOST_SUBNET $nconnect_num $randomnize $NFS_MOUNT_
LOG_FILE ${vips_to_mount[*]} &" &

            wait

            # If file is not empty, then something went wrong
            clush -S -w "$clients" -f "$CLUSH_MAX_FAN_OUT" "if [[ ! -L $NFS_MOUNT_LOG_
FILE ]] && [ -s $NFS_MOUNT_LOG_FILE ]; then if grep -q 'exited with exit code\|timed
out\|Failed\|failed' $NFS_MOUNT_LOG_FILE; then cat $NFS_MOUNT_LOG_FILE; exit $ERROR_
CODE_NO_MULTIPATH_MOUNT_FAILURE_ALL; fi; fi"
            status=$?
            if [ "$status" -gt 0 ]; then
                echo "Mount failed, exiting..."
                exit "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_ALL"
            fi
        fi
    else
        if [ "$inject_error_num" == "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL" ];
then
            # Mess up randomnize to force the error
            randomnize="yesno"
        fi

        # Mount may fail, but clush return code is 0. So capture the stderr result to
know the actual error instead.
        clush -w "$clients" -f "$CLUSH_MAX_FAN_OUT" "cd $VAST_SCALE_TESTING_PATH; ./
nfs_mount.sh $MOUNT $REMOTE_PATH $HOST_SUBNET $nconnect_num $randomnize $NFS_MOUNT_LOG_
FILE ${vips_to_mount[*]} &" &
        # No wait to complete here for better performance, as check of error condition
is also deferred
    fi
}
```

```
#================================================================
# Mount the VIPs using multipath and nconnect
#================================================================
mount_nfs_multipath () {

    # Create node index and vip mappings in round robin fashion
    #
    # Although each client will set remoteports to all VIPs,
    # each client is going to mount 1 VIP to the mountpoint,
    # and that VIP will be round robinned.

    i=0
    node_index=()
    node_vip=()
    vip_last_field=$HOST_SUBNET_FIRST_VIP_LAST_FIELD

    for node in "${CLIENT_NODES_ARRAY[@]}"
    do
        node_index+=("${node}:${i}")
        node_vip+=("${node}:${vip_last_field}")
        i=$((i + 1))
        vip_last_field=$((vip_last_field + 1))
        if [ $vip_last_field -gt "$HOST_SUBNET_LAST_VIP_LAST_FIELD" ]
        then
            vip_last_field="$HOST_SUBNET_FIRST_VIP_LAST_FIELD"
        fi
    done

    # Now mount it.
    # localports need to match the HOST_SUBNET
```

```
    # chown is necessary otherwise FIO cannot write to the folder.

    if [ "$inject_error_num" == "$ERROR_CODE_MULTIPATH_MOUNT_FAILURE" ]; then
        # Mess up HOST_SUBNET to force the error
        host_subnet_org="$HOST_SUBNET"
        HOST_SUBNET="A.B.C"
    fi

    if [[ ! -L "$NFS_MOUNT_LOG_FILE" ]] && [ -f "$NFS_MOUNT_LOG_FILE" ]; then
        rm -rf "$NFS_MOUNT_LOG_FILE"
    fi

    clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "sudo umount ${MOUNT} > /dev/null
2>/dev/null;
                                                    sudo mkdir -p ${MOUNT};
                                                    sudo mount -v -o
vers=3,proto=tcp,nconnect=$NCONNECT_NUM,port=20048,localports=\$(ifconfig | grep
inet | grep $HOST_SUBNET | xargs | awk '{print \$2}'),remoteports=$HOST_SUBNET_
FIRST_VIP-$HOST_SUBNET_LAST_VIP $HOST_SUBNET.\$(echo ${node_vip[*]} | grep -o
[^[:space:]]*\$(hostname -s):[^[:space:]]* | cut -d':' -f 2):/ ${MOUNT};
                                                    io_dir=${MOUNT}/${REMOTE_
PATH}/\$(hostname -s);
                                                    sudo mkdir -p \${io_dir}; sudo
chown $uid:$uid \${io_dir}" 2> "$NFS_MOUNT_LOG_FILE"

    if [ "$inject_error_num" == "$ERROR_CODE_MULTIPATH_MOUNT_FAILURE" ]; then
        # Mess up HOST_SUBNET to force the error
        HOST_SUBNET="$host_subnet_org"
    fi

    # If file is not empty, then something went wrong
    if  [[ ! -L "$NFS_MOUNT_LOG_FILE" ]] && [ -s "$NFS_MOUNT_LOG_FILE" ]; then
```

```
        if grep -q "exited with exit code\|timed out\|Failed\|failed" "$NFS_MOUNT_LOG_
FILE"; then
            echo "Mount failed: exiting..."
            exit "$ERROR_CODE_MULTIPATH_MOUNT_FAILURE"
        fi
    fi
}

#====================================================================
# Mount VIPs based on test type and multipath configuration
#====================================================================
mount_vips () {

    # Mount VIP pools
    # S3 is not using file system, so no NFS mount needed
    if [ "$skip_mount" == "0" ] && [ "$test_type" != "elbencho-s3" ]; then
        if [ "$MULTIPATH_ENABLE" == "1" ]; then
            mount_nfs_multipath
       else
            mount_nfs_no_multipath "all"
        fi
    fi

}

#====================================================================
# Start FIO server or elbencho server on clients
# Precondition: VIPs are mounted already
#====================================================================
start_services () {
```

```
    if [ "$test_type" == "fio" ]; then
        # Start fio on all clients
        if [ "$FIO_SERVER_CLIENT_MODE" == "1" ]; then
            clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "sudo pkill fio; sudo rm
fiopid; ${FIO_PATH}/fio --server --daemonize=fiopid &" &
        fi
    elif [ "$test_type" == "elbencho-file" ] || [ "$test_type" == "elbencho-s3" ]; then
        clush -w "$CLIENT_NODES" -f "$CLUSH_MAX_FAN_OUT" "docker container kill
elbencho-server > /dev/null 2>/dev/null;
                                                    docker container rm elbencho-
server > /dev/null 2>/dev/null;
                                                    docker run --rm --name
elbencho-server -v ${MOUNT}/${REMOTE_PATH}/\$(hostname -s):/data/${REMOTE_PATH}
--net=host -dt breuner/elbencho:$ELBENCHO_TAG --service --foreground --zones 0,1"
    fi

    wait
}

#=================================================================
# Split the FIO json summary files to per client record + All_Clients.
#
# For aggregated results such as total BW/IOPS/Latency, check All_Clients record.
#
# This is to make it easier for Elastic Search (out of scope of this harness),
# otherwise the size could be to big when creating the index
#=================================================================
split_server_log_file () {
    for file in *_summary.log
    do
        "${SCRIPT_DIR}/split_json.sh" "$file"
    done
```

```
}


#====================================================================
# ZIP the logs to keep the size of result smaller
#====================================================================
zip_logs () {
    zip -r "$(basename "$run_dir")_Logs.zip" ./*summary.inf ./*summary.json ./redwood-*
All_Clients
    rm -rf ./*summary.log ./*summary.inf ./*summary.json ./redwood-* All_Clients
    zip -r context.zip ./context
    zip -r "$schedule_name.zip" "$schedule_name"
    rm -rf context "$schedule_name"
}


#====================================================================
# Save the configurations for this test to result folder
# Todo: save the VAST Data Appliance version as well using public API?
#====================================================================
before_test () {

    # Set the xtrace prefix to show a timestamp and line number
    PS4="[$(date +%T)] $LINENO: "

    # Enable extended globbing
    shopt -s extglob

    # Save the general telemetry information
    mkdir -p "$run_dir/context"

    if [ "$skip_all_client_config_query" == "0" ]; then
        # Save irqbalance/turbo boost and hyper threading information
```

```
        get_all_client_config | tee "$run_dir/context/all_client_config.log"
    fi

    # Save the context and hardware telemetry of head node.
    # It may not have 100G, but it does have 10G
    echo "$FIO_PATH/fio version: $($FIO_PATH/fio --version)" > "$run_dir/context/general_info.log"
    "${SCRIPT_DIR}/general_info.sh" "$_HOST_SUBNET_10G" "$run_dir/context" >> "$run_dir/context/general_info.log"   # Head node may not have 100G, but does have 10G, so use 10G subnet
    sudo dmidecode | tee "$run_dir/context/hardware_info.log" > /dev/null

    # Do the same for client nodes as well.
    # For client nodes, every run will wipe out the context from previous runs
    local context_path="$VAST_SCALE_TESTING_PATH/context"
    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "mkdir -p $context_path;
                                                      echo $FIO_PATH/fio version: $($FIO_PATH/fio --version) > $context_path/general_info.log;
                                                      $VAST_SCALE_TESTING_PATH/general_info.sh $HOST_SUBNET $context_path >> $context_path/general_info.log;
                                                      sudo dmidecode | tee $context_path/hardware_info.log > /dev/null" &
    wait

    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" --rcopy "$context_path/hardware_info.log" --dest "$run_dir/context"
    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" --rcopy "$context_path/general_info.log" --dest "$run_dir/context"
    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" --rcopy "$context_path/numa.png" --dest "$run_dir/context"

    for file in "$run_dir"/context/numa.png.*
    do
```

```
        local client_name
        client_name="$(basename "$file" | cut -d'.' -f 3)"
        mv "$run_dir/context/numa.png.$client_name" "$run_dir/context/numa.$client_
name.png"
    done

    # Save the configuration and schedule files
    cp "${SCRIPT_DIR}/env.conf" "$run_dir"
    if [ -f "${SCRIPT_DIR}/env.conf.override" ]; then
        cp "${SCRIPT_DIR}/env.conf.override" "$run_dir"
    fi
    cp "$schedule_file" "$run_dir"

    # Only FIO test needs to save the job files in the folder
    if [ "$test_type" == "fio" ]; then
        mkdir "$schedule_name"
    fi
}

#=====================================================================
# Automatically upload the results to NFS share after the test
# This can be enabled/disabled by RESULTS_UPLOAD_ENABLE.
#=====================================================================
upload_results () {

    if [ "$RESULTS_UPLOAD_ENABLE" == "0" ]; then
        return
    fi

    local cluster_result_path
    cluster_result_path="$(basename "${run_dir}")"
```

```
    local retry=0
    while [[ $retry -eq 0 || ( "$(ls "${RESULTS_MNTPOINT}")" != "" && $retry -lt 3 ) ]]
    do
        sudo umount -lf "${RESULTS_MNTPOINT}" >/dev/null 2>/dev/null
        sudo rm -rf "${RESULTS_MNTPOINT}/results*"
        sudo mkdir -p "${RESULTS_MNTPOINT}"
        if [ "$RESULTS_UPLOAD_MOUNT_CMD" != "" ] && [[ "$RESULTS_UPLOAD_MOUNT_CMD" ==
"sudo mount"* ]]; then
            $RESULTS_UPLOAD_MOUNT_CMD
        else
            break
        fi
        sleep 1
        sudo mkdir -p "${RESULTS_MNTPOINT}/${test_type}/${cluster_result_path}"
        sudo cp -r "${run_dir}"/* "${RESULTS_MNTPOINT}/${test_type}/${cluster_result_
path}"
        sleep 1
        sudo umount "${RESULTS_MNTPOINT}"
        sleep 1
        retry=$((retry+1))
    done
}


#=================================================================
# Post processing of logs and result upload
#=================================================================
after_test () {


    # Save of the end of test dmesg for both head node and client nodes
    dmesg > "$run_dir/context/End_of_test_dmesg.log"
```

```
    local context_path="$VAST_SCALE_TESTING_PATH/context"
    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "mkdir -p $context_path;
                                                      dmesg > $context_path/End_of_
test_dmesg.log &" &
    wait

    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" --rcopy "$context_path/End_of_
test_dmesg.log" --dest "$run_dir/context" &

    wait

    if [ "$test_type" == "fio" ] && [ "$FIO_SERVER_CLIENT_MODE" == "1" ]; then
        split_server_log_file
        zip_logs
    fi

    # Move normal log only.
    # Error log shall cause the harness to exit immediately and won't be here.
    # Error log will be moved by the run_test.sh
    if [[ ! -L "$RUN_LOG_FILE" ]] && [ -f "$RUN_LOG_FILE" ]; then
        mv "$RUN_LOG_FILE" "$run_dir"
    fi

    upload_results
}

#=================================================================
# Run FIO workload
# Todo: so many passed in parameters, compact all to an array instead?
#=================================================================
run_fio_workload () {
```

```
    local wl_name=$1
    local wl_num=$2
    local jobname=$3
    local client_count=$4
    local rw=$5
    local numjobs=$6
    local iodepth=$7
    local bs=$8
    local size=$9
    local read_perc=${10}
    local random_type=${11}
    local create_on_open=${12}
    local fallocate=${13}
    local create_serialize=${14}
    local fill_device=${15}
    local create_only=${16}

    echo "Running #$wl_num workload: $wl_name"

    # Always generate the json+ output
    local fio_cmd="${FIO_PATH}/fio --output-format=json+ --output ${wl_name}_summary.log"

    if [ "$FIO_SERVER_CLIENT_MODE" == "0" ]; then
        clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "mkdir -p $VAST_SCALE_TESTING_
PATH/${result_prefix}_${schedule_name}_${hostname}/${schedule_name}"
    fi

    local i jobfile fio_server
    i=0
    jobfile="$run_dir/$schedule_name/${wl_name}.ini"
    for fio_server in "${client_nodes_array[@]}"
```

```
    do
        i=$((i + 1))
        if [ $i -gt "$client_count" ]; then break; fi
        mkdir -p "$run_dir/$schedule_name/$fio_server"

        # Generate the actual per client FIO job file based on template and the passed
in parameters defined by schedule file

        if [ "$create_only" == "1" ]; then
            sed "/^#create_only=/c create_only=$create_only" -i "$jobfile"
        fi

        if [ "$runtime" != "" ]; then
            sed "/^runtime=/c runtime=$runtime" -i "$jobfile"
        fi

        if [ "$ramptime" != "" ]; then
            sed "/^ramp_time=/c ramp_time=$ramptime" -i "$jobfile"
        fi

        if [ "$create_on_open" != "invalid" ]; then
            sed "/^#create_on_open=/c create_on_open=$create_on_open" -i "$jobfile"
        fi

        if [ "$fallocate" != "invalid" ]; then
            sed "/^#fallocate=/c fallocate=$fallocate" -i "$jobfile"
        fi

        if [ "$create_serialize" != "invalid" ]; then
            sed "/^#create_serialize=/c create_serialize=$create_serialize" -i
"$jobfile"
```

```
fi

if [ "$fill_device" == "1" ]; then
    sed "/^#fill_device=/c fill_device=1" -i "$jobfile"
    sed '/^runtime=/d' -i "$jobfile"
    sed '/^ramp_time=/d' -i "$jobfile"
    sed '/^time_based/d' -i "$jobfile"
fi

sed "/^rw=/c rw=$rw" -i "$jobfile"
sed "/^numjobs=/c numjobs=$numjobs" -i "$jobfile"
sed "/^iodepth=/c iodepth=$iodepth" -i "$jobfile"
sed "/^bs=/c bs=$bs" -i "$jobfile"

if [ "$size" != "invalid" ]; then
    sed "/^size=/c size=$size" -i "$jobfile"
fi

if [ "$jobname" != "invalid" ]; then
    sed "/\[jobname]/c [$jobname]" -i "$jobfile"
else
    sed "/\[jobname]/c [$wl_name]" -i "$jobfile"
fi

# Two different ways of running mixed workloads
# 1. rwmixread:     mix workload on every client
# 2. rwmixreadhost: read on some clients, but write on other clients
if [ "$read_perc" != "invalid" ]; then
    # update rwmixread
    if [ "$random_type" == "rwmixread" ]; then
        sed "/^#rwmixread=/c rwmixread=$read_perc" -i "$jobfile"
```

```
            elif [ "$random_type" == "rwmixreadhost" ]; then
                local random=$((1 + RANDOM % 100))
                if [ "$random" -gt "$read_perc" ]; then
                    if [[ "$wl_name" == *"_rnd_"* ]]; then
                        sed "/^rw=/c rw=randwrite" -i "$jobfile"
                    else
                        sed "/^rw=/c rw=write" -i "$jobfile"
                    fi
                else
                    if [[ "$wl_name" == *"_rnd_"* ]]; then
                        sed "/^rw=/c rw=randread" -i "$jobfile"
                    else
                        sed "/^rw=/c rw=read" -i "$jobfile"
                    fi
                fi
            fi
        fi

        if [ "$MULTIPATH_ENABLE" == "1" ]; then
            sed "s/hostname/${fio_server}/g" "$jobfile"  > "$run_dir/$schedule_name/$fio_
server/${wl_name}.ini"
        else
            local vip_list vips
            vip_list=$(shuf -e "${ALL_VIPS[@]}" | xargs)
            IFS=" " read -r -a vips <<< "$vip_list"

            local fio_directory=""
            local vip_count=0
            local vips_to_mount=()
            for MOUNTVIP in "${vips[@]}"; do
                if [ "$vip_count" == "$numjobs" ] && [ "$MOUNT_ALL" == "0" ]; then
```

```
                    # Number of vips to be mounted == number of jobs
                    break;
                fi

                if [ "$fio_directory" != "" ]; then
                    fio_directory+=":"
                fi
                fio_directory+="${MOUNT}/${REMOTE_PATH}/${fio_server}/${HOST_
SUBNET}.${MOUNTVIP}"   # hostname is placeholder, and will be replaced later

                vips_to_mount+=("$MOUNTVIP")
                vip_count=$((vip_count+1))
            done
            sed "/directory=/c directory=$fio_directory" "$jobfile" > "$run_
dir/$schedule_name/$fio_server/${wl_name}.ini"
        fi

        if [ "$FIO_SERVER_CLIENT_MODE" == "1" ]; then
            fio_cmd+=" --client $fio_server $run_dir/$schedule_name/$fio_server/${wl_
name}.ini"
        else
            clush -w "$fio_server" -c "$run_dir/$schedule_name/$fio_server/$wl_
name.ini" --dest "$VAST_SCALE_TESTING_PATH/${result_prefix}_${schedule_
name}_${hostname}/${schedule_name}" &
        fi

        # Mount vips for this server if they haven't been mounted already
        # Applicable to individual mount only (multipath mount is always done upfront,
and not affected by MOUNT_ALL flag)
        if [ "$MOUNT_ALL" == "0" ];
        then
            mount_nfs_no_multipath "$fio_server" "${vips_to_mount[@]}"
```

```
            # Ensure mount is done every a few servers.
            # This is to avoid overwhelming clush
            if [ $(( i % MOUNT_STAGGER )) -eq 0 ]; then
                wait

                # If file is not empty, then something went wrong
                clush -S -w "$fio_server" -f "$CLUSH_MAX_FAN_OUT" "if [[ ! -L $NFS_
MOUNT_LOG_FILE ]] && [ -s $NFS_MOUNT_LOG_FILE ]; then if grep -q 'exited with exit
code\|timed out\|Failed\|failed' $NFS_MOUNT_LOG_FILE; then cat $NFS_MOUNT_LOG_FILE;
exit $ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL; fi; fi"
                local status=$?
                if [ "$status" -gt 0 ]; then
                    echo "Mount failed, exiting..."
                    exit "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL"
                fi
            fi
        fi
    done

    if [ "$dry_run" == "0" ]; then
        # Run test
        echo "Run test..."
        if [ "$FIO_SERVER_CLIENT_MODE" == "1" ]; then
            $fio_cmd
        else
            wait

            # If file is not empty, then something went wrong
            clush -S -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "if [[ ! -L $NFS_
MOUNT_LOG_FILE ]] && [ -s $NFS_MOUNT_LOG_FILE ]; then if grep -q 'exited with exit
code\|timed out\|Failed\|failed' $NFS_MOUNT_LOG_FILE; then cat $NFS_MOUNT_LOG_FILE;
```

```
exit $ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL; fi; fi"
            status=$?
            if [ "$status" -gt 0 ]; then
                echo "Mount failed, exiting..."
                exit "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE_PARTIAL"
            fi

            clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "cd $VAST_SCALE_TESTING_
PATH/${result_prefix}_${schedule_name}_${hostname}; $fio_cmd $VAST_SCALE_TESTING_
PATH/${result_prefix}_${schedule_name}_${hostname}/${schedule_name}/${wl_name}.ini"
            wait
        fi
    fi
}

#=====================================================================
# Parse FIO workload arguments and run it
#=====================================================================
run_test_fio() {
    local wl_num=$1
    local client_count=$2

    local fio_args_str fio_args
    fio_args_str=$(echo "$3" | xargs)
    IFS=" " read -r -a fio_args <<< "$fio_args_str"

    local rw="invalid"
    local numjobs="invalid"
    local iodepth="invalid"
    local bs="invalid"
    local jobname="invalid"
```

```
local size="invalid"
local read_perc="invalid"
local random_type="invalid"
local create_on_open="invalid"
local fallocate="invalid"
local create_serialize="invalid"
local fill_device="invalid"
local tag="invalid"
local create_only="invalid"

# All supported options in schedule file shall be added in this for loop

for arg in "${fio_args[@]}"
do
   case "$arg" in
     --rw=*)
         rw="${arg#*=}"
     ;;
     --numjobs=*)
         numjobs="${arg#*=}"
     ;;
     --iodepth=*)
         iodepth="${arg#*=}"
     ;;
     --bs=*)
         bs="${arg#*=}"
     ;;
     --size=*)
         size="${arg#*=}"
     ;;
```

```
        --runtime=*)
            if [ "$cmdline_runtime" == "" ]; then
                runtime="${arg#*=}"
            else
                runtime=$cmdline_runtime
            fi
        ;;
        --ramp_time=*)
            if [ "$cmdline_ramptime" == "" ]; then
                 ramptime="${arg#*=}"
            else
                 ramptime=$cmdline_ramptime
            fi
        ;;
        --name=*)
            jobname="${arg#*=}"
        ;;
        --fill_device=*)
            fill_device="${arg#*=}"
        ;;
        --rwmixread=*)
            random_type="rwmixread"
            read_perc="${arg#*=}"
        ;;
        --rwmixhostread=*)
            random_type="rwmixhostread"
            read_perc="${arg#*=}"
        ;;
        --fallocate=*)
            fallocate="${arg#*=}"
        ;;
```

```
            --create_on_open=*)
                create_on_open="${arg#*=}"
            ;;
            --create_serialize=*)
                create_serialize="${arg#*=}"
            ;;
            --tag=*)
                tag="${arg#*=}"
            ;;
            --create_only)
                create_only=1
            ;;
        esac
    done

    if [ "$cmdline_create_only" == "1" ]; then
        create_only=1
    fi

    # Generate the summary log file name for the workload
    local qd=$((numjobs * iodepth))

    case "$rw" in

        randread)
            op="rd_rnd"
        ;;

        randwrite)
            op="wr_rnd"
        ;;
```

```
        randrw)
            op="rw_rnd"
        ;;

        read)
            op="rd"
        ;;

        write)
            op="wr"
        ;;

        rw|readwrite)
            op="rw"
        ;;
    esac

    # Sanity check of the mandatory options in schedule file
    if [ "$rw" == "invalid" ] || [ "$numjobs" == "invalid" ] || [ "$iodepth" ==
"invalid" ] || [ "$bs" == "invalid" ]; then
         echo "Please specify --rw/--numjobs/--iodepth/--bs in schedule file"
         exit "$ERROR_CODE_MISSING_FIO_WORKLOAD_OPTIONS"
    fi

    # Need to differentiate the 2 mix read types in naming
    if [ "$read_perc" != "invalid" ]; then
        if [ "$random_type" == "rwmixreadhost" ]; then
            wl_name="${op}_qd_${qd}_${bs}_${read_perc}hostrd_${numjobs}w"
        else
            wl_name="${op}_qd_${qd}_${bs}_${read_perc}rd_${numjobs}w"
```

```
        fi
    else
        wl_name="${op}_qd_${qd}_${bs}_${numjobs}w"
    fi

    # For workloads that cannot be differentiated by the mandatory options,
    # append workload number as the unique identifier for the log filename
    if [ "$tag" == "1" ]; then
        if [ "$jobname" == "invalid" ]; then
            jobname="$wl_name"
        fi
        wl_name+="_${wl_num}tag"
    fi

    cp "$WORKLOAD_DIR/fio/$fio_template" "$schedule_name/${wl_name}.ini"
    if [ "$wl_num" -gt 1 ]; then
        sleep "$INTER_WORKLOAD_WAIT_IN_SECONDS"
    fi

    # Run workload
    echo -e "\nBegin #$wl_num workload: $wl_name"
    run_fio_workload "$wl_name" "$wl_num" "$jobname" "$client_count" "$rw" "$numjobs"
"$iodepth" "$bs" "$size" "$read_perc" "$random_type" "$create_on_open" "$fallocate"
"$create_serialize" "$fill_device" "$create_only"
    echo -e "End of #$wl_num workload: $wl_name\n"
}

#=====================================================================
# Parse elbencho file workload arguments and run it
#=====================================================================
run_test_elbencho_file() {
```

```
    local wl_num=$1
    local client_count=$2
    local elbencho_args elbencho_args_array
    elbencho_args=$(echo "$3" | xargs)
    IFS=" " read -r -a elbencho_args_array <<< "$elbencho_args"

    local threads="invalid"
    local iodepth="invalid"
    local bs="invalid"
    local op="invalid"
    local jobname="invalid"
    local tag="invalid"

    # Parse the options in schedule file for log file naming
    local arg
    for arg in "${elbencho_args_array[@]}"
    do
       case "$arg" in
         --write)
             op="wr"
         ;;
         --read)
             op="rd"
         ;;
         --iodepth=*)
             iodepth="${arg#*=}"
         ;;
         --threads=*)
             threads="${arg#*=}"
         ;;
         --block=*)
```

```
                bs="${arg#*=}"
            ;;
            --rand) # This is optional
                random=1
            ;;
            --name=*) # This is not from Elbencho
                jobname="${arg#*=}"
            ;;
            --tag=*)
                tag="${arg#*=}"
            ;;
        esac
    done

    # Remove elbencho unsupported options
    local elbencho_args_stripped elbencho_args_stripped_array
    elbencho_args_stripped=$(echo "$elbencho_args" | sed -e 's/[^ ]*--name=[^ ]*//ig' |
sed -e 's/[^ ]*--tag=[^ ]*//ig')
    IFS=" " read -r -a elbencho_args_stripped_array <<< "$elbencho_args_stripped"

    # Sanity check of the mandatory options in schedule file
    if [ "$op" == "invalid" ] || [ "$threads" == "invalid" ] || [ "$iodepth" ==
"invalid" ] || [ "$bs" == "invalid" ]; then
            echo "Please specify --rw/--threads/--iodepth/--bs in schedule file"
            exit "$ERROR_CODE_MISSING_ELBENCHO_FILE_WORKLOAD_OPTIONS"
    fi

    local ELBENCHO_FILE_RESULT_SUBDIR_ON_HOST="elbencho/file"
    local ELBENCHO_FILE_RESULT_DIR_ON_CONTAINER="/vast-scale-testing/file"
    local result_filename="${op}-${bs}bs-${threads}t-${iodepth}iodepth"
    if [ "$random" == "1" ]; then
```

```
        result_filename+="_rnd"
    fi

    # For workloads that cannot be differentiated by the mandatory options,
    # append workload number as the unique identifier for the log filename
    if [ "$tag" == "1" ]; then
        result_filename+="_${wl_num}tag"
    fi

    wl_name="$result_filename"

    if [ "$wl_num" -gt 1 ]; then
        sleep "$INTER_WORKLOAD_WAIT_IN_SECONDS"
    fi

    echo -e "\nBegin #$wl_num workload: $wl_name"
    mkdir -p "$run_dir/$ELBENCHO_FILE_RESULT_SUBDIR_ON_HOST"
    docker container kill elbencho-client > /dev/null 2>/dev/null;
    docker container rm elbencho-client > /dev/null 2>/dev/null;

    local data_directory=()
    if [ "$MULTIPATH_ENABLE" == "1" ]; then
        if [ "$jobname" == "invalid" ]; then
            data_directory=("/data/${REMOTE_PATH}")
        else
            data_directory=("/data/${REMOTE_PATH}/$jobname")
            clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "io_
dir=${MOUNT}/${REMOTE_PATH}/\$(hostname -s)/$jobname;
            sudo mkdir -p \${io_dir}; sudo chown $uid:$uid \${io_dir}"
        fi
    else
```

```
        local vip_count=0
        local vips_to_mount=()

        for MOUNTVIP in "${ALL_VIPS[@]}"; do
            if [ "$jobname" != "invalid" ]; then
                data_directory+=("/data/${REMOTE_PATH}/${HOST_
SUBNET}.${MOUNTVIP}/$jobname")
            else
                data_directory+=("/data/${REMOTE_PATH}/${HOST_SUBNET}.${MOUNTVIP}")
            fi

            vips_to_mount+=("$MOUNTVIP")
            vip_count=$((vip_count+1))
        done

        # Mount vips for this server if they haven't been mounted already
        # Applicable to individual mount only (multipath mount is always done upfront,
and not affected by MOUNT_ALL flag)
        if [ "$MOUNT_ALL" == "0" ];
        then
            echo "MOUNT_ALL==0 is not supported for elbencho file testing"
            exit
        else
            # Need to create the folder if jobname is specified.
            # For no jobname scenario, the folder should have been created when
mounting.
            if [ "$jobname" != "invalid" ]; then
                for MOUNTVIP in "${vips_to_mount[@]}"; do
                    clush -w "$client_nodes" -f "$CLUSH_MAX_FAN_OUT" "io_
dir=${MOUNT}/${REMOTE_PATH}/\$(hostname -s)/${HOST_SUBNET}.${MOUNTVIP}/$jobname;
                                                                    sudo mkdir -p
\${io_dir}; sudo chown $uid:$uid \${io_dir}"
```

```
                done
            fi
        fi
    fi

    if [ "$dry_run" == "0" ]; then
        docker run --rm --name elbencho-client --net=host -v "$run_dir/$ELBENCHO_
FILE_RESULT_SUBDIR_ON_HOST:$ELBENCHO_FILE_RESULT_DIR_ON_CONTAINER" "breuner/
elbencho:$ELBENCHO_TAG" "${data_directory[@]}" "${elbencho_args_stripped_array[@]}"
--hosts "$client_nodes" --csvfile "$ELBENCHO_FILE_RESULT_DIR_ON_CONTAINER/${result_
filename}.csv"
    fi

    echo -e "End of #$wl_num workload: $wl_name\n"
}

#=====================================================================
# Parse elbencho S3 workload arguments and run it
#=====================================================================
run_test_elbencho_s3() {
    local wl_num=$1
    local client_count=$2

    local elbencho_args elbencho_args_array
    elbencho_args=$(echo "$3" | xargs)
    IFS=" " read -r -a elbencho_args_array <<< "$elbencho_args"

    # Parse the options in schedule file for log file naming
    local arg op threads bs random
    for arg in "${elbencho_args_array[@]}"
    do
       case "$arg" in
```

```
            --write)
                op="wr"
            ;;
            --read)
                op="rd"
            ;;
            --threads=*)
                threads="${arg#*=}"
            ;;
            --block=*)
                bs="${arg#*=}"
            ;;
            --rand)  # This is optional
                random=1
            ;;
        esac
    done

    # Sanity check of the mandatory options in schedule file
    if [ "$op" == "invalid" ] || [ "$threads" == "invalid" ] || [ "$iodepth" ==
"invalid" ] || [ "$bs" == "invalid" ]; then
            echo "Please specify --rw/--numjobs/--iodepth/--bs in schedule file"
            exit "$ERROR_CODE_MISSING_ELBENCHO_S3_WORKLOAD_OPTIONS"
    fi

    local ELBENCHO_S3_RESULT_SUBDIR_ON_HOST="elbencho/s3"
    local ELBENCHO_S3_RESULT_DIR_ON_CONTAINER="/vast-scale-testing/s3"
    local result_filename="${op}-${bs}bs-${threads}t-${iodepth}iodepth"
    if [ "$random" == "1" ]; then
        result_filename+="_rnd"
    fi
```

```
    local endpoints
    endpoints=$(eval echo "http://$HOST_SUBNET.{$HOST_SUBNET_FIRST_VIP_LAST_
FIELD..$HOST_SUBNET_LAST_VIP_LAST_FIELD}")

    wl_name="$result_filename"
    if [ "$wl_num" -gt 1 ]; then
        sleep "$INTER_WORKLOAD_WAIT_IN_SECONDS"
    fi
    echo -e "\nBegin #$wl_num workload: $wl_name"
    mkdir -p "$run_dir/$ELBENCHO_S3_RESULT_SUBDIR_ON_HOST"
    docker container kill elbencho-client > /dev/null 2>/dev/null;
    docker container rm elbencho-client > /dev/null 2>/dev/null;

    if [ "$dry_run" == "0" ]; then
        docker run --rm --name elbencho-client --net=host -v "$run_dir/$ELBENCHO_
S3_RESULT_SUBDIR_ON_HOST:$ELBENCHO_S3_RESULT_DIR_ON_CONTAINER" "breuner/
elbencho:$ELBENCHO_TAG" "${elbencho_args_array[@]}" --s3endpoints "$endpoints" --s3key
"$ELBENCHO_S3KEY" --s3secret "$ELBENCHO_S3SECRET" --hosts "$client_nodes" --csvfile
"$ELBENCHO_S3_RESULT_DIR_ON_CONTAINER/${result_filename}.csv"
    fi

    echo -e "End of #$wl_num workload: $wl_name\n"
}

#=================================================================
# Run test with the client count of current step.
# Each step will run all workloads in the schedule file.
#=================================================================
run_step () {

    local client_count=$1
```

```
    before_test

    # Any abort after context is save will do normal shutdown handling
    #graceful_terminate_required=1

    local wl_num=0
    local workloads=()
    local line

    while read -r line
    do
        line=$(echo "$line" | xargs)

        # Skip the comment and empty lines
        if [ "$line" == "" ] || [[ "$line" == "#"* ]]; then continue; fi

        workloads+=("$line ")

    done < "$schedule_name.$schedule_name_ext"  # clush seems to screw up the read sometimes, so read lines to workloads variable first

    for line in  "${workloads[@]}"
    do
        line=$(echo "$line" | xargs)

        # There is still an empty line, so filter again
        if [ "$line" == "" ] || [[ "$line" == "#"* ]]; then continue; fi

        echo "Parsing: $line"
```

```
        local wl_num=$((wl_num + 1))

        case "$test_type" in

          "fio")
            run_test_fio           "$wl_num" "$client_count" "$line"
            ;;
          "elbencho-file")
            run_test_elbencho_file "$wl_num" "$client_count" "$line"
            ;;
          "elbencho-s3")
            run_test_elbencho_s3   "$wl_num" "$client_count" "$line"
            ;;
        esac
    done

    after_test
}

#=====================================================================
# Run test with specified network speed, schedule file and step size
#=====================================================================
run_test () {

    schedule_name="$(basename "$schedule_file" | cut -d'.' -f 1)"
    schedule_name_ext="$(basename "$schedule_file" | cut -d'.' -f 2)"
    hostname=$(hostname -s)

    local count=$STEP_SIZE
    local previous_count=0
```

```
    local step
    for step in $(seq 1 $STEPS)
    do
        if [ "$step" -gt 1 ]; then
            sleep "$INTER_STEP_WAIT_IN_SECONDS"
        fi

        local actual_step_size=$((count - previous_count)) # for incremental type,
previous_count is always 0
        local client_nodes_string
        client_nodes_string="${CLIENT_NODES_ARRAY[*]:$previous_count:$actual_step_
size}"

        IFS=" " read -r -a client_nodes_array <<< "$client_nodes_string"
        client_nodes=${client_nodes_string// /,}

        echo -e "\nBegin the test of $count clients..."
        if [ "$step_type" == "$STEP_TYPE_CONSTANT" ]; then
            result_prefix="${common_prefix}_${actual_step_size}ClientsFrom${previous_
count}"
        else
            result_prefix="${common_prefix}_${count}Clients"
        fi
        run_dir="${RESULTS_DIR}/${result_prefix}_${schedule_name}_${hostname}"
        if [ "$extra_prefix" != "" ]; then
            run_dir="${run_dir}_${extra_prefix}"
        fi
        mkdir -p "$run_dir"

        echo "$run_dir" > "$RUN_DIR_TEMP_FILE"
```

```
        pushd "$run_dir" || exit "$ERROR_CODE_FAIL_TO_GOTO_DIR"

        run_step "$count"

        popd || exit "$ERROR_CODE_FAIL_TO_EXIT_DIR"

        if [ "$step_type" == "$STEP_TYPE_CONSTANT" ]; then
            previous_count="$count"
        fi
        count=$((count + STEP_SIZE))
        if [ $count -gt "$CLIENT_COUNT" ]; then
            count="$CLIENT_COUNT"
        fi
    done
}

#=====================================================================
# Call functions defined above to kick off the test
#=====================================================================
parse_command_line "$@"
parse_env_and_set_globals
check_client_status_and_deploy_files
set_irqb_tb_ht_controls
mount_vips
start_services
run_test
```

5.6.25 SPLIT_JSON.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
```

```
# SPDX-License-Identifier: BSD-3-Clause


IN_FILE="$1"


if [ ! -f "$IN_FILE" ] || [[ -L "$IN_FILE" ]] ; then
    echo "File $IN_FILE not found or is symbolic link!"
    exit 2 # Standard error code ENOENT: No such file or directory
fi


WL=${IN_FILE//_summary.log/}
echo "$WL"
OUT_FILE="${WL}_summary.json"
ALL_CLIENTS_OUT_DIR="All_Clients"


echo -e "Start splitting ${IN_FILE}: $(date)"


sed '/{/,$d' < "$IN_FILE" > "${WL}_summary.inf"
sed '/{/,$!d' < "$IN_FILE" > "${WL}_summary.json"


mkdir -p "${ALL_CLIENTS_OUT_DIR}"


touch "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}"
{
    echo "{"
    echo "  \"fio_version\" : $(jq '."fio version"' < "${OUT_FILE}"),"
    echo "  \"timestamp\" : $(jq '."timestamp"' < "${OUT_FILE}"),"
    echo "  \"time\" : $(jq '."time"' < "${OUT_FILE}"),"
} >> "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.common"


echo '  "global options" : {' > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.options.tmp1"
jq '."global options"' < "${OUT_FILE}" | sed 's/}/},/g' > "${ALL_CLIENTS_OUT_
```

```
DIR}/${OUT_FILE}.options.tmp2"
sed 's/^/    /' "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.options.tmp2" >  "${ALL_CLIENTS_
OUT_DIR}/${OUT_FILE}.options.tmp3"
sed '1d' "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.options.tmp3" > "${ALL_CLIENTS_OUT_
DIR}/${OUT_FILE}.options.tmp4"
cat "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.options.tmp1" "${ALL_CLIENTS_OUT_DIR}/${OUT_
FILE}.options.tmp4" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.options"


# extract summry which is the last item in array (-1 index)
echo '  "client_stats" : [' > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats.tmp1"
jq '."client_stats"[-1]' < "${OUT_FILE}" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats.
tmp2"
sed 's/^/    /' "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats.tmp2" > "${ALL_CLIENTS_OUT_
DIR}/${OUT_FILE}.stats.tmp3"
cat "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats.tmp1" "${ALL_CLIENTS_OUT_DIR}/${OUT_
FILE}.stats.tmp3" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats"
echo '  ],' >> "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats"


echo '  "disk_util" : [' > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util.tmp1"
jq '."disk_util"[-1]' < "${OUT_FILE}" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util.tmp2"
sed 's/^/    /' "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util.tmp2" > "${ALL_CLIENTS_OUT_
DIR}/${OUT_FILE}.util.tmp3"
cat "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util.tmp1" "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.
util.tmp3" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util"
echo '  ]' >> "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.util"



cat "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.common" "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.
options" "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats" "${ALL_CLIENTS_OUT_DIR}/${OUT_
FILE}.util" > "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}"
echo '}' >> "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}"


client_num=$(jq '."client_stats" | length' < "${OUT_FILE}")
```

```
for i in $(seq 0 $((client_num - 2))); do
    jq ".\"client_stats\"[${i}]" < "${OUT_FILE}" > "client-${i}.tmp1"
    sed 's/^/    /' "client-${i}.tmp1" > "client-${i}.tmp2"
    echo '  ]' >> "client-${i}.tmp2"
    cat "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.common" "${ALL_CLIENTS_OUT_DIR}/${OUT_
FILE}.options" "${ALL_CLIENTS_OUT_DIR}/${OUT_FILE}.stats.tmp1" "client-${i}.tmp2" >
"client-${i}.json"
    echo '}' >> "client-${i}.json"
    hostname=$(grep "hostname" "client-${i}.json" | cut -d'"' -f 4)
    mkdir -p "$hostname"
    mv "client-${i}.json" "$hostname/${WL}_summary.json"
done

rm -rf ./*.tmp* ./${ALL_CLIENTS_OUT_DIR}/*.tmp* ./${ALL_CLIENTS_OUT_DIR}/*.common
./${ALL_CLIENTS_OUT_DIR}/*.options ./${ALL_CLIENTS_OUT_DIR}/*.stats ./${ALL_CLIENTS_
OUT_DIR}/*.util

echo -e "Done:  $(date)"
```

5.6.26 VALIDATION

5.6.26.1 RD_RND_QD_128_1024K_16W_SUMMARY.LOG

```
This log is too long to be added as appendix. Clone the repo instead to see the
content.
```

5.6.26.2 README.MD

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# 1. To test how split_json.sh works, run:
#       $../split_json.sh rd_rnd_qd_128_1024k_16w_summary.log
#    rd_rnd_qd_128_1024k_16w_summary.log is a sample FIO summary file in json+ format.
```

```
#
# 2. To test the error conditions, run:
#       $./test_error_code.sh
```

5.6.26.3 TEST_ERROR_CODE.SH

```
#!/bin/bash

# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

validation_dir="$(cd "$(dirname "$0")" || exit; pwd)"
script_dir=${validation_dir/%\/validation/}

source "$script_dir/error_code.conf"

# Todo: not all error codes are tested, as some require more work,
#       i.e. to change the configuration with custom env.conf.override.
#       However, the framework is sufficient to test all.

# Use a normal command line for this, but do not provide valid client names in env.
conf.override for force the error
test_cmd_index_for_ERROR_CODE_CLIENT_NOT_REACHABLE=4

test_cmds=("../run_test.sh -ns 20"
           "../run_test.sh -ns 100 -sf foo.sch"
           "../run_test.sh -ns 100 -st 3"
           "../run_test.sh -ns 100 -crazyoption 1 > /dev/null"
           "../run_test.sh -ns 100 -sf ../../workloads/fio/bw_test_short_100g.sch"
           # This following test assumes env.conf.override doesn't change MULTIPATH_
```

```
ENABLE to 0,
           # otherwise this will be a normal run and the test will fail
           "../run_test.sh -ns 100 -sf ../../workloads/fio/bw_test_short_100g.sch -ei
$ERROR_CODE_MULTIPATH_MOUNT_FAILURE"
           "../nfs_mount.sh /mnt/test fio 1.2.3 2 yes"
           "../nfs_mount.sh / fio 1.2.3 2 yes 1 2 3"
           "../nfs_mount.sh /mnt/test fio subnet 2 yes 1 2 3"
           "../nfs_mount.sh /mnt/test fio 1.2.3 -1 yes 1 2 3"
           "../nfs_mount.sh /mnt/test fio 1.2.3 2 random 1 2 3"
           "../nfs_mount.sh /mnt/test fio 1.2.3 2 yes 1 2 3"
           )

expected_errcode=("$ERROR_CODE_INVALID_NETWORK_SPEED"
                  "$ERROR_CODE_INVALID_SCHEDULE_FILE"
                  "$ERROR_CODE_INVALID_STEPPING_TYPE"
                  "$ERROR_CODE_INVALID_OPTION"
                  "$ERROR_CODE_CLIENT_NOT_REACHABLE"
                  "$ERROR_CODE_MULTIPATH_MOUNT_FAILURE"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NUMBER_OF_INPUTS"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_MOUNT_BASE"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_SUBNET"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_NCONNECT_NUM"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_INVALID_RANDOMNIZE"
                  "$ERROR_CODE_NO_MULTIPATH_MOUNT_FAILURE"
                  )

index=0
passed=0
failed=0
for cmd in "${test_cmds[@]}"
do
```

```
    echo "=================="
    echo "Testing command: $cmd"

    if [ "$index" == "$test_cmd_index_for_ERROR_CODE_CLIENT_NOT_REACHABLE" ]; then
        # Save env.conf.override to backup
        # Assumption is that env.conf only contains the dummy clients and VIPs
        if [[ ! -L $script_dir/env.conf.override ]] && [ -s $script_dir/env.conf.
override ]; then
            mv $script_dir/env.conf.override $script_dir/env.conf.override.bak
        fi
    fi

    $cmd
    if [ "$?" == "${expected_errcode[index]}" ];
    then
        passed=$((passed + 1))
        echo -e "Result: pass"
    else
        failed=$((failed + 1))
        echo -e "Result: fail"
    fi

    if [ "$index" == "$test_cmd_index_for_ERROR_CODE_CLIENT_NOT_REACHABLE" ]; then
        # Restore env.conf.override
        if [[ ! -L $script_dir/env.conf.override.bak ]] && [ -s $script_dir/env.conf.
override.bak ]; then
            mv $script_dir/env.conf.override.bak $script_dir/env.conf.override
        fi
    fi

    index=$((index + 1))
```

```
done


echo "=================="
echo "Total tests: $index"
echo "Passed:      $passed"
echo "Failed:      $failed"
echo "=================="


if [ "$failed" -ne 0 ];
then
    exit 1 # Catch all shell failure code
fi
```

5.7 FIO Test Schedule Files (scale-testing-for-vastdata/workloads/fio)

5.7.1 README.MD

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


#================================================================================
==================
# FIO options (refer to FIO man page for their definitions) supported by the fio schedule
file:
#
#     Mandatory: rw/numjobs/iodepth/bs
#                They will be used to generate the fio summary log filename, and jobname
when --name is not specified.
#     Optional:  size/runtime/ramp_time/name/fallocate/create_on_open/fill_device/
create_serialize
#                Default values of these are as specified in fio_template.ini, or
```

```
disabled if commented out in fio_template.ini.
#
# Non-FIO options supported by the schedule file:
#     --tag: set the flag to 1 to append workload number to the generated summary log filename,
#            and jobname but only when --name is not specified.
#
#            This can be used for workloads that cannot be differentiated by the mandatory options,
#            to generate unique summary log filename.
#
#            Note that this doesn't affect the jobname if --name is specified.
#     --rwmixreadhost: set this option to run client based mixed workload, eg:
#                           --rw=randrw --rwmixreadhost=70: 70% of clients do random read, 30% of clients do random write
#                           --rw=rw     --rwmixreadhost=70: 70% of clients do sequential read, 30% of clients do sequential write
#                      In comparison, --rwmixread is the native fio option, which specifies the mixed workload per client ratio, eg:
#                           --rw=randrw --rwmixread=70: every client does 70% random read, 30% random write
#                           --rw=rw     --rwmixread=70: every client does 70% sequential read, 30% sequential write
#
# Two ways to support additional options:
# 1) Enhance run_test_fio() function of run_test_no_runlog.sh
# 2) modify fio_template.ini
#
```

5.7.2 FIO_TEMPLATE.INI

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause
```

```
[global]
name=vastfiotest
ioengine=libaio
direct=1
thread=1
buffered=0
randrepeat=0
time_based
norandommap
group_reporting=1
percentile_list=1.0:25.0:50.0:75.0:90.0:99.0:99.9:99.99:99.999:99.9999:99.99999:99.999999:100.0
random_distribution=random
refill_buffers=1
#create_only=0
#fill_device=1
create_serialize=0
#create_on_open=1
#fallocate=none
size=20G
bs=1024k
iodepth=16
numjobs=32
rw=randread
#rwmixread=70
runtime=180
ramp_time=120

[jobname]
directory=/mnt/nfs-multipath-tcp/fio/hostname
```

5.7.3 BW_TEST_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size = 16 jobs x 5g x 20 clients = 1.6T.
--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=5g  --fill_device=1
--name=bw-test-job1-100g --tag=1
--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=5g  --runtime=300
--ramp_time=60 --name=bw-test-job1-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1
--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=5g  --runtime=300
--ramp_time=60 --name=bw-test-job1-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1
--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=5g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job1-100g --tag=1


# Data size = 16 jobs x 30g x 20 clients = 9.6T.
--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --fill_device=1
--name=bw-test-job2-100g --tag=1
--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --runtime=300
--ramp_time=60 --name=bw-test-job2-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1
--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --runtime=300
--ramp_time=60 --name=bw-test-job2-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1
--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job2-100g --tag=1


# Data size = 16 jobs x 75g x 20 clients = 24T.
--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=75g  --fill_device=1
--name=bw-test-job3-100g --tag=1
--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=75g  --runtime=300
--ramp_time=60 --name=bw-test-job3-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1
```

```
--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=75g  --runtime=300
--ramp_time=60 --name=bw-test-job3-100g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=75g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job3-100g --tag=1
```

5.7.4 BW_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size = 16 jobs x 1g x 100 clients = 1.6T.

--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=1g  --fill_device=1
--name=bw-test-job1-10g --tag=1

--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=1g  --runtime=300
--ramp_time=60 --name=bw-test-job1-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=1g  --runtime=300
--ramp_time=60 --name=bw-test-job1-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=1g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job1-10g --tag=1


# Data size = 16 jobs x 6g x 100 clients = 9.6T.

--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --fill_device=1
--name=bw-test-job2-10g --tag=1

--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --runtime=300
--ramp_time=60 --name=bw-test-job2-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --runtime=300
--ramp_time=60 --name=bw-test-job2-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job2-10g --tag=1


# Data size = 16 jobs x 15g x 100 clients = 24T.
```

```
--rw=write      --numjobs=16 --iodepth=8  --bs=1024k --size=15g  --fill_device=1
--name=bw-test-job3-10g --tag=1

--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=15g  --runtime=300
--ramp_time=60 --name=bw-test-job3-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=read       --numjobs=16 --iodepth=8  --bs=1024k --size=15g  --runtime=300
--ramp_time=60 --name=bw-test-job3-10g --fallocate=none --create_on_open=1 --create_
serialize=0 --tag=1

--rw=randwrite  --numjobs=16 --iodepth=8  --bs=1024k --size=15g  --runtime=300
--ramp_time=60 --name=bw-test-rnd-job3-10g --tag=1

5.7.5    bw_test_short_100g.sch

# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size = 16 jobs x 30g x 20 clients = 9.6T.

--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --runtime=300 --ramp_
time=60
```

5.7.6 BW_TEST_SHORT_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size = 16 jobs x 6g x 100 clients = 9.6T.

--rw=randread   --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --runtime=300
--ramp_time=60


5.7.7    iops_test_100g.sch

# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size = 4 jobs x 6400m x 20 clients = 512G.

# 2 x 512g write for randread. Total data for read and write < 4T, both data will be
```

```
hold in Optane

# Create the files first, then read altogether.

--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --create_only --name=iops-test-1-100g --tag=1

--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --name=iops-test-1-100g --tag=1

--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --create_only --name=iops-test-2-100g --tag=1

--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --name=iops-test-2-100g --tag=1

# Another 2 x 512G write. This could be a partial write, as the runtime is too short.

# Only use it for write statistics. For read statistics, need to regenerate the
complete file.

--rw=randwrite  --numjobs=4  --iodepth=64  --bs=4k --size=6400m --runtime=300 --ramp_
time=60

--rw=randwrite  --numjobs=4  --iodepth=16  --bs=4k --size=6400m --runtime=300 --ramp_
time=60

# Write 9.6T (> 2 * 4T) of data to fully destage the previous data to QLC

--rw=write      --numjobs=16 --iodepth=8   --bs=1024k --size=30g --fill_device

# Now re-read the data written before

--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --name=iops-test-1-100g --tag=1

--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=6400m --runtime=300 --ramp_
time=60 --name=iops-test-2-100g --tag=1
```

5.7.8 IOPS_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size = 4 jobs x 1280m x 100 clients = 512G.

# 2 x 512g write for randread iops. Total data for read and write < 4T, both data will
be hold in Optane
```

```
# Create the files first, then read altogether.
--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --create_only --name=iops-test-1-10g --tag=1
--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --name=iops-test-1-10g --tag=1
--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --create_only --name=iops-test-2-10g --tag=1
--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --name=iops-test-2-10g --tag=1
# Another 2 x 512G write. This could be a partial write, as the runtime is too short.
# Only use it for write statistics. For read statistics, need to regenerate the
complete file.
--rw=randwrite  --numjobs=4  --iodepth=64  --bs=4k --size=1280m --runtime=300 --ramp_
time=60
--rw=randwrite  --numjobs=4  --iodepth=16  --bs=4k --size=1280m --runtime=300 --ramp_
time=60
# Write 9.6T (> 2 * 4T) of data to fully destage the previous data to QLC
--rw=write      --numjobs=16 --iodepth=8   --bs=1024k --size=6g --fill_device
# Now re-read the data written before
--rw=randread   --numjobs=4  --iodepth=64  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --name=iops-test-1-10g --tag=1
--rw=randread   --numjobs=4  --iodepth=16  --bs=4k --size=1280m --runtime=300 --ramp_
time=60 --name=iops-test-2-10g --tag=1
```

5.7.9 SWEEP_WORKERS_AND_IODEPTH_RANDOM_1024K_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 15g x 20 clients = 9.6T.
--rw=randread --numjobs=32 --iodepth=32 --bs=1024k --size=15g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=16 --iodepth=32 --bs=1024k --size=30g  --name=sweep-workers-16
--tag=1
```

```
--rw=randread --numjobs=8  --iodepth=32 --bs=1024k --size=60g  --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=4  --iodepth=32 --bs=1024k --size=120g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=1  --iodepth=32 --bs=1024k --size=480g --name=sweep-workers-1
--tag=1


--rw=randread --numjobs=32 --iodepth=16 --bs=1024k --size=15g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=16 --iodepth=16 --bs=1024k --size=30g  --name=sweep-workers-16
--tag=1
--rw=randread --numjobs=8  --iodepth=16 --bs=1024k --size=60g  --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=4  --iodepth=16 --bs=1024k --size=120g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=1  --iodepth=16 --bs=1024k --size=480g --name=sweep-workers-1
--tag=1


--rw=randread --numjobs=32 --iodepth=8  --bs=1024k --size=15g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=16 --iodepth=8  --bs=1024k --size=30g  --name=sweep-workers-16
--tag=1
--rw=randread --numjobs=8  --iodepth=8  --bs=1024k --size=60g  --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=4  --iodepth=8  --bs=1024k --size=120g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=1  --iodepth=8  --bs=1024k --size=480g --name=sweep-workers-1
--tag=1
```

5.7.10 SWEEP_WORKERS_AND_IODEPTH_RANDOM_1024K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 3g x 100 clients = 9.6T.
```

```
--rw=randread --numjobs=32 --iodepth=32 --bs=1024k --size=3g  --name=sweep-workers-32
--create_only --tag=1
--rw=randread --numjobs=32 --iodepth=32 --bs=1024k --size=3g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=32 --iodepth=16 --bs=1024k --size=3g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=32 --iodepth=8  --bs=1024k --size=3g  --name=sweep-workers-32
--tag=1
--rw=randread --numjobs=16 --iodepth=32 --bs=1024k --size=6g  --name=sweep-workers-16
--create_only --tag=1
--rw=randread --numjobs=16 --iodepth=32 --bs=1024k --size=6g  --name=sweep-workers-16
--tag=1
--rw=randread --numjobs=16 --iodepth=16 --bs=1024k --size=6g  --name=sweep-workers-16
--tag=1
--rw=randread --numjobs=16 --iodepth=8  --bs=1024k --size=6g  --name=sweep-workers-16
--tag=1
--rw=randread --numjobs=8  --iodepth=32 --bs=1024k --size=12g --name=sweep-workers-8
--create_only --tag=1
--rw=randread --numjobs=8  --iodepth=32 --bs=1024k --size=12g --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=8  --iodepth=16 --bs=1024k --size=12g --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=8  --iodepth=8  --bs=1024k --size=12g --name=sweep-workers-8
--tag=1
--rw=randread --numjobs=4  --iodepth=32 --bs=1024k --size=24g --name=sweep-workers-4
--create_only --tag=1
--rw=randread --numjobs=4  --iodepth=32 --bs=1024k --size=24g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=4  --iodepth=16 --bs=1024k --size=24g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=4  --iodepth=8  --bs=1024k --size=24g --name=sweep-workers-4
--tag=1
--rw=randread --numjobs=1  --iodepth=32 --bs=1024k --size=96g --name=sweep-workers-1
--create_only --tag=1
```

```
--rw=randread --numjobs=1  --iodepth=32 --bs=1024k --size=96g --name=sweep-workers-1
--tag=1
--rw=randread --numjobs=1  --iodepth=16 --bs=1024k --size=96g --name=sweep-workers-1
--tag=1
--rw=randread --numjobs=1  --iodepth=8  --bs=1024k --size=96g --name=sweep-workers-1
--tag=1
```

5.7.11 SWEEP_WORKERS_AND_IODEPTH_SEQUENTIAL_1024K_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 15g x 20 clients = 9.6T.
--rw=read --numjobs=32 --iodepth=16 --bs=1024k --size=15g
--rw=read --numjobs=16 --iodepth=16 --bs=1024k --size=30g
--rw=read --numjobs=8  --iodepth=16 --bs=1024k --size=60g
--rw=read --numjobs=4  --iodepth=16 --bs=1024k --size=120g
--rw=read --numjobs=1  --iodepth=16 --bs=1024k --size=480g

--rw=read --numjobs=32 --iodepth=8  --bs=1024k --size=15g
--rw=read --numjobs=16 --iodepth=8  --bs=1024k --size=30g
--rw=read --numjobs=8  --iodepth=8  --bs=1024k --size=60g
--rw=read --numjobs=4  --iodepth=8  --bs=1024k --size=120g
--rw=read --numjobs=1  --iodepth=8  --bs=1024k --size=480g
```

5.7.12 SWEEP_WORKERS_AND_IODEPTH_SEQUENTIAL_1024K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 3g x 100 clients = 9.6T.
--rw=read --numjobs=32 --iodepth=16 --bs=1024k --size=3g
--rw=read --numjobs=16 --iodepth=16 --bs=1024k --size=6g
--rw=read --numjobs=8  --iodepth=16 --bs=1024k --size=12g
```

```
--rw=read --numjobs=4  --iodepth=16 --bs=1024k --size=24g
--rw=read --numjobs=1  --iodepth=16 --bs=1024k --size=96g

--rw=read --numjobs=32 --iodepth=8  --bs=1024k --size=3g
--rw=read --numjobs=16 --iodepth=8  --bs=1024k --size=6g
--rw=read --numjobs=8  --iodepth=8  --bs=1024k --size=12g
--rw=read --numjobs=4  --iodepth=8  --bs=1024k --size=24g
--rw=read --numjobs=1  --iodepth=8  --bs=1024k --size=96g
```

5.7.13 SWEEP_WORKERS_AND_IODEPTH_RANDOM_4K_100G.SCH

```
 # Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 15g x 20 clients = 9.6T.
--rw=randread --numjobs=32 --iodepth=16 --bs=4k --size=15g
--rw=randread --numjobs=16 --iodepth=16 --bs=4k --size=30g
--rw=randread --numjobs=8  --iodepth=16 --bs=4k --size=60g
--rw=randread --numjobs=4  --iodepth=16 --bs=4k --size=120g
--rw=randread --numjobs=1  --iodepth=16 --bs=4k --size=480g

--rw=randread --numjobs=32 --iodepth=8  --bs=4k --size=15g
--rw=randread --numjobs=16 --iodepth=8  --bs=4k --size=30g
--rw=randread --numjobs=8  --iodepth=8  --bs=4k --size=60g
--rw=randread --numjobs=4  --iodepth=8  --bs=4k --size=120g
--rw=randread --numjobs=1  --iodepth=8  --bs=4k --size=480g
```

5.7.1.4 SWEEP_WORKERS_AND_IODEPTH_RANDOM_4K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause
```

```
# Data size per workload = 32 jobs x 3g x 100 clients = 9.6T.
--rw=randread --numjobs=32 --iodepth=16 --bs=4k --size=3g
--rw=randread --numjobs=16 --iodepth=16 --bs=4k --size=6g
--rw=randread --numjobs=8  --iodepth=16 --bs=4k --size=12g
--rw=randread --numjobs=4  --iodepth=16 --bs=4k --size=24g
--rw=randread --numjobs=1  --iodepth=16 --bs=4k --size=96g

--rw=randread --numjobs=32 --iodepth=8  --bs=4k --size=3g
--rw=randread --numjobs=16 --iodepth=8  --bs=4k --size=6g
--rw=randread --numjobs=8  --iodepth=8  --bs=4k --size=12g
--rw=randread --numjobs=4  --iodepth=8  --bs=4k --size=24g
--rw=randread --numjobs=1  --iodepth=8  --bs=4k --size=96g
```

5.7.15 SWEEP_WORKERS_AND_IODEPTH_SEQUENTIAL_4K_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 15g x 20 clients = 9.6T.
--rw=read --numjobs=32 --iodepth=16 --bs=4k --size=15g
--rw=read --numjobs=16 --iodepth=16 --bs=4k --size=30g
--rw=read --numjobs=8  --iodepth=16 --bs=4k --size=60g
--rw=read --numjobs=4  --iodepth=16 --bs=4k --size=120g
--rw=read --numjobs=1  --iodepth=16 --bs=4k --size=480g

--rw=read --numjobs=32 --iodepth=8  --bs=4k --size=15g
--rw=read --numjobs=16 --iodepth=8  --bs=4k --size=30g
--rw=read --numjobs=8  --iodepth=8  --bs=4k --size=60g
--rw=read --numjobs=4  --iodepth=8  --bs=4k --size=120g
--rw=read --numjobs=1  --iodepth=8  --bs=4k --size=480g
```

5.7.16 SWEEP_WORKERS_AND_IODEPTH_SEQUENTIAL_4K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 3g x 100 clients = 9.6T.
--rw=read --numjobs=32 --iodepth=16 --bs=4k --size=3g
--rw=read --numjobs=16 --iodepth=16 --bs=4k --size=6g
--rw=read --numjobs=8  --iodepth=16 --bs=4k --size=12g
--rw=read --numjobs=4  --iodepth=16 --bs=4k --size=24g
--rw=read --numjobs=1  --iodepth=16 --bs=4k --size=96g

--rw=read --numjobs=32 --iodepth=8  --bs=4k --size=3g
--rw=read --numjobs=16 --iodepth=8  --bs=4k --size=6g
--rw=read --numjobs=8  --iodepth=8  --bs=4k --size=12g
--rw=read --numjobs=4  --iodepth=8  --bs=4k --size=24g
--rw=read --numjobs=1  --iodepth=8  --bs=4k --size=96g
```

5.7.17 SWEEP_IOSIZES_100RD_RANDOM_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=randread --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=randread --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.18 SWEEP_IOSIZES_100RD_RANDOM_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
--rw=randread --numjobs=16 --iodepth=8 --bs=1024k --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=512k  --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=256k  --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=128k  --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=64k   --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=32k   --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=16k   --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=8k    --size=6g
--rw=randread --numjobs=16 --iodepth=8 --bs=4k    --size=6g
```

5.7.19 SWEEP_IOSIZES_100RD_SEQUENTIAL_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=read --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=read --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.20 SWEEP_IOSIZES_100RD_SEQUENTIAL_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
--rw=read --numjobs=16 --iodepth=8 --bs=1024k --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=512k  --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=256k  --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=128k  --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=64k   --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=32k   --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=16k   --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=8k    --size=6g
--rw=read --numjobs=16 --iodepth=8 --bs=4k    --size=6g
```

5.7.21 SWEEP_IOSIZES_100WR_RANDOM_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=randwrite --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.22 SWEEP_IOSIZES_100WR_RANDOM_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
```

```
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=randwrite --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=randwrite --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.23 SWEEP_IOSIZES_100WR_SEQUENTIAL_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=write --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.24 SWEEP_IOSIZES_100WR_SEQUENTIAL_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause
```

```
# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
--rw=write --numjobs=16 --iodepth=8 --bs=1024k --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=512k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=256k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=128k  --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=64k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=32k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=16k   --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=8k    --size=30g
--rw=write --numjobs=16 --iodepth=8 --bs=4k    --size=30g
```

5.7.25 SWEEP_IOSIZES_70RD_RANDOM_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
# Workload on each client is 70 read, 30 write
--rw=randrw --numjobs=16 --iodepth=8 --bs=1024k --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=512k  --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=256k  --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=128k  --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=64k   --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=32k   --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=16k   --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=8k    --size=30g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=4k    --size=30g --rwmixread=70
```

5.7.26 SWEEP_IOSIZES_70RD_RANDOM_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
```

```
# Workload on each client is 70 read, 30 write
--rw=randrw --numjobs=16 --iodepth=8 --bs=1024k --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=512k  --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=256k  --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=128k  --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=64k   --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=32k   --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=16k   --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=8k    --size=6g --rwmixread=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=4k    --size=6g --rwmixread=70
```

5.7.27 SWEEP_IOSIZES_70RD_SEQUENTIAL_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
# Workload on each client is 70 read, 30 write
--rw=rw --numjobs=16 --iodepth=8 --bs=1024k --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=512k  --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=256k  --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=128k  --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=64k   --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=32k   --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=16k   --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=8k    --size=30g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=4k    --size=30g --rwmixread=70
```

5.7.28 SWEEP_IOSIZES_70RD_SEQUENTIAL_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
```

```
# Workload on each client is 70 read, 30 write
--rw=rw --numjobs=16 --iodepth=8 --bs=1024k --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=512k  --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=256k  --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=128k  --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=64k   --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=32k   --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=16k   --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=8k    --size=6g --rwmixread=70
--rw=rw --numjobs=16 --iodepth=8 --bs=4k    --size=6g --rwmixread=70
```

5.7.29 SWEEP_IOSIZES_70HOSTRD_RANDOM_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
# 70% of the clients: random read; 30% of the clients: random write
--rw=randrw --numjobs=16 --iodepth=8 --bs=1024k --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=512k  --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=256k  --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=128k  --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=64k   --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=32k   --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=16k   --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=8k    --size=30g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=4k    --size=30g --rwmixreadhost=70
```

5.7.30 SWEEP_IOSIZES_70HOSTRD_RANDOM_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
```

```
# 70% of the clients: random read; 30% of the clients: random write
--rw=randrw --numjobs=16 --iodepth=8 --bs=1024k --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=512k  --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=256k  --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=128k  --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=64k   --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=32k   --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=16k   --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=8k    --size=6g --rwmixreadhost=70
--rw=randrw --numjobs=16 --iodepth=8 --bs=4k    --size=6g --rwmixreadhost=70
```

5.7.31 SWEEP_IOSIZES_70HOSTRD_SEQUENTIAL_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 30g x 20 clients = 9.6T.
# 70% of the clients: om read; 30% of the clients: om write
--rw=rw --numjobs=16 --iodepth=8 --bs=1024k --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=512k  --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=256k  --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=128k  --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=64k   --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=32k   --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=16k   --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=8k    --size=30g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=4k    --size=30g --rwmixreadhost=70
```

5.7.32 SWEEP_IOSIZES_70HOSTRD_SEQUENTIAL_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 16 jobs x 6g x 100 clients = 9.6T.
```

```
# 70% of the clients: om read; 30% of the clients: om write
--rw=rw --numjobs=16 --iodepth=8 --bs=1024k --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=512k  --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=256k  --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=128k  --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=64k   --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=32k   --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=16k   --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=8k    --size=6g --rwmixreadhost=70
--rw=rw --numjobs=16 --iodepth=8 --bs=4k    --size=6g --rwmixreadhost=70
```

5.8 Elbencho File Test Schedule Files (scale-testing-for-vastdata/workloads/elbencho/file)

5.8.1 README.MD

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


#================================================================================
==================
# Elbencho options supported by the Elbencho File testing schedule file:
#
#     Mandatory: rw/numjobs/iodepth/bs
#                They will be used to generate the result filename
#     Optional:  all valid Elbencho options for file testing are supported in schedule
file. Common ones to be used are:
#                rand/blockvarpct/direct/size/nosvcshare/dirs/files/mkdirs/timelimit/
infloop
#
#
# Non-Elbencho options supported by the schedule file:
#     --name: This is added for workload to create the files under the specified folder,
```

```
instead of the default.
#             The difference is whether the file will be reused or new files will be
generated.
#             For example, following 2 lines will write to the same file
#                 --write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k
--direct --size=1g --nosvcshare --dirs=1 --files=1 --mkdirs
#                 --write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k
--direct --size=1g --nosvcshare --dirs=1 --files=1 --mkdirs
#             While the lines below will write to different files:
#                 --write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k
--direct --size=1g --nosvcshare --dirs=1 --files=1 --mkdirs --name=bw-test-100g-job1
#                 --write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k
--direct --size=1g --nosvcshare --dirs=1 --files=1 --mkdirs --name=bw-test-100g-job2
#     --tag: set the flag to 1 to append workload number to the generated .csv filename,
#
#            This can be used for workloads that cannot be differentiated by the
mandatory options,
#            to generate a unique summary .csv filename.
#
#            Note that:
#                 - this doesn't affect the jobname,
#                 - elbencho will append the subsequent results to .csv even for
workloads that have the same mandatory options.
#                   So for the 2 examples above, it's perfectly fine not to use --tag at
all, although using tag would make it easier to identify each workload's result .csv.
#
```

5.8.2 BW_TEST_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 1g x 20 clients x 1 dirs x 5 files = 1.6T.
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
```

```
--nosvcshare --dirs=1 --files=5 --mkdirs --name=bw-test-100g-job1

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=bw-test-100g-job1

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop --name=bw-test-100g-job1

# Short random write for statistics

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=5 --mkdirs --rand --timelimit=300 --name=bw-test-100g-job2


# Data size per workload = 16 jobs x 6g x 20 clients x 1 dirs x 5 files = 9.6T.

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs --name=bw-test-100g-job3

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=bw-test-100g-job3

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop --name=bw-test-100g-job3

# Short random write for statistics

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs --rand --timelimit=300 --name=bw-test-100g-job4


# Data size per workload = 16 jobs x 15g x 20 clients x 1 dirs x 5 files = 24T.

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=5 --mkdirs --name=bw-test-100g-job5

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=bw-test-100g-job5

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop --name=bw-test-100g-job5

# Short random write for statistics

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=5 --mkdirs --rand --timelimit=300 --name=bw-test-100g-job6
```

5.8.3 BW_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause
```

```
# Data size per workload = 16 jobs x 1g x 100 clients x 1 dirs x 1 file = 1.6T.
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=1 --mkdirs --name=bw-test-10g-job1
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=bw-test-10g-job1
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop --name=bw-test-10g-job1
# Short random write for statistics
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=1g
--nosvcshare --dirs=1 --files=1 --mkdirs --rand --timelimit=300 --name=bw-test-10g-job2

# Data size per workload = 16 jobs x 6g x 100 clients x 1 dirs x 1 file = 9.6T.
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs --name=bw-test-10g-job3
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=bw-test-10g-job3
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop --name=bw-test-10g-job3
# Short random write for statistics
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs --rand --timelimit=300 --name=bw-test-10g-job4

# Data size per workload = 16 jobs x 15g x 100 clients x 1 dirs x 1 file = 24T.
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=1 --mkdirs --name=bw-test-10g-job5
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=bw-test-10g-job5
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop --name=bw-test-10g-job5
# Short random write for statistics
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=15g
--nosvcshare --dirs=1 --files=1 --mkdirs --rand --timelimit=300 --name=bw-test-10g-job6
```

5.8.4 IOPS_TEST_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 4 threads x 1280m x 20 clients x 1 dirs x 5 files = 512G.
# Total data for read and write < 4T, both data will be held in Optane
--write --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --mkdirs --name=iops-test-100g-job1
--read  --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=iops-test-100g-
job1


--write --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --mkdirs --name=iops-test-100g-job2
--read  --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=iops-test-100g-
job2
# For random write statistics only.
--write --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --mkdirs --rand --timelimit=300 --name=iops-test-100g-
job3
--write --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --mkdirs --rand --timelimit=300 --name=iops-test-100g-
job4
# Write 9.6T (> 2 * 4T) of data to fully destage the previous data to QLC
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs --name=iops-test-100g-job5
# Now re-read the data written before
--read  --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=iops-test-100g-
job1
--read  --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop --name=iops-test-100g-
job2
```

5.8.5 IOPS_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 4 threads x 1280m x 100 clients x 1 dirs x 1 files = 512G.
# Total data for read and write < 4T, both data will be held in Optane
# Create the files first, then read altogether.
--write --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --mkdirs --name=iops-test-10g-job1
--read  --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=iops-test-10g-
job1


--write --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --mkdirs --name=iops-test-10g-job2
--read  --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=iops-test-10g-
job2
# For random write statistics only.
--write --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --mkdirs --rand --timelimit=300 --name=iops-test-10g-
job3
--write --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --mkdirs --rand --timelimit=300 --name=iops-test-10g-
job4


# Write 9.6T (> 2 * 4T) of data to fully destage the previous data to QLC
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs --name=iops-test-10g-job5
# Now re-read the data written before
--read  --blockvarpct=100 --iodepth=64 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=iops-test-10g-
job1
```

```
--read  --blockvarpct=100 --iodepth=16 --threads=4 --block=4k --direct --size=1280m
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop --name=iops-test-10g-
job2
```

5.8.6 SWEEP_RANDOM_IOSIZES_100G.SCH

```
# Data size per workload = 16 jobs x 6g x 20 clients x 1 dirs x 5 files = 9.6T.
--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs
# Random read
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=512k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=256k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=128k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=64k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=32k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=16k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=8k    --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k    --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
```

5.8.7 SWEEP_RANDOM_IOSIZES_10G.SCH

```
# Data size per workload = 16 jobs x 6g x 100 clients x 1 dirs x 1 files = 9.6T.
--write --blockvarpct=100 --iodepth=16 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs
# Random read
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=512k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=256k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=128k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=64k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=32k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=16k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=8k    --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k    --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

5.8.8 SWEEP_RANDOM_THREADS_AND_IODEPTH_1024K_100G.SCH

```
# Data size per workload = 32 jobs x 3g x 20 clients x 1 dirs x 5 files = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --mkdirs

# Random read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

# Random read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --rand --timelimit=300 --infloop
```

5.8.9 SWEEP_RANDOM_THREADS_AND_IODEPTH_1024K_10G.SCH

```
# Data size per workload = 32 jobs x 3g x 100 clients x 1 dirs x 1 file = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --mkdirs

# Random read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

# Random read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

5.8.10 SWEEP_RANDOM_THREADS_AND_IODEPTH_4K_100G.SCH

```
# Data size per workload = 32 jobs x 3g x 100 clients x 1 dirs x 1 files = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --mkdirs

# Random read, iodepth 8
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

# Random read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

# Random read, iodepth 64

--read  --blockvarpct=100 --iodepth=64 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

5.8.11 SWEEP_RANDOM_THREADS_AND_IODEPTH_4K_10G.SCH

```
# Data size per workload = 32 jobs x 3g x 100 clients x 1 dirs x 1 file = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --mkdirs

# Random read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

# Random read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

# Random read, iodepth 64

--read  --blockvarpct=100 --iodepth=64 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=64 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --rand --timelimit=300 --infloop
```

5.8.12 SWEEP_SEQUENTIAL_THREADS_AND_IODEPTH_1024K_100G.SCH

```
# Data size per workload = 32 jobs x 3g x 20 clients x 1 dirs x 5 files = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --mkdirs

# Sequential read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

# Sequential read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop
```

5.8.13 SWEEP_SEQUENTIAL_THREADS_AND_IODEPTH_1024K_10G.SCH

```
# Data size per workload = 32 jobs x 3g x 100 clients x 1 dirs x 1 file = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --file=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --file=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --file=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --file=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --file=1 --mkdirs

# Sequential read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

# Sequential read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=1024k --direct --size=3g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=1024k --direct --size=12g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=1024k --direct --size=24g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=1024k --direct --size=96g
--nosvcshare --dirs=1 --file=1 --timelimit=300 --infloop
```

5.8.14 SWEEP_SEQUENTIAL_THREADS_AND_IODEPTH_4K_100G.SCH

```
# Data size per workload = 32 jobs x 3g x 20 clients x 1 dirs x 5 files = 9.6T.

# Generate all data first

--write --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --mkdirs

# Sequential read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop
```

```
--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

# Sequential read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

# Sequential read, iodepth 32

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop
```

5.8.15 SWEEP_SEQUENTIAL_THREADS_AND_IODEPTH_4K_10G.SCH

```
# Data size per workload = 32 jobs x 3g x 100 clients x 1 dirs x 1 files = 9.6T.

# Generate all data first
```

```
--write --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --mkdirs

--write --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --mkdirs

# Sequential read, iodepth 8

--read  --blockvarpct=100 --iodepth=8 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

# Sequential read, iodepth 16

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

# Sequential read, iodepth 32

--read  --blockvarpct=100 --iodepth=16 --threads=32 --block=4k --direct --size=3g
```

```
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=16 --block=4k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=8  --block=4k --direct --size=12g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=4  --block=4k --direct --size=24g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=16 --threads=1  --block=4k --direct --size=96g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop
```

5.8.16 SWEEP_SEQUENTIAL_IOSIZES_100G.SCH

```
# Data size per workload = 16 jobs x 6g x 20 clients x 1 dirs x 5 files = 9.6T.

--write --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --mkdirs

# Sequential read

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=512k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=256k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=128k  --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=64k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=32k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=16k   --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=8k    --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k    --direct --size=6g
--nosvcshare --dirs=1 --files=5 --timelimit=300 --infloop
```

5.8.17 SWEEP_SEQUENTIAL_IOSIZES_10G.SCH

```
# Data size per workload = 16 jobs x 6g x 100 clients x 1 dirs x 1 files = 9.6T.

--write --blockvarpct=100 --iodepth=16 --threads=32 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --mkdirs

# Sequential read

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=1024k --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=512k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=256k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=128k  --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=64k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=32k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=16k   --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=8k    --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop

--read  --blockvarpct=100 --iodepth=8 --threads=16 --block=4k    --direct --size=6g
--nosvcshare --dirs=1 --files=1 --timelimit=300 --infloop
```

5.9 Elbencho S3 Test Schedule Files (scale-testing-for-vastdata/workloads/elbencho/s3)

5.9.1 README.MD

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


#================================================================================
==================
```

```
# Elbencho options supported by the Elbencho Object testing schedule file:
#
#     Mandatory: rw/numjobs/iodepth/bs
#                They will be used to generate the result filename
#     Optional:  all valid Elbencho options for object testing are supported in
schedule file. Common ones to be used are:
#                s3randobj/blockvarpct/size/dirs/files/mkdirs/timelimit/infloop
#
#
# Non-Elbencho options supported by the schedule file: none
```

5.9.2 BW_TEST_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 20 clients x 1 dirs x 32 files = 1.6T.
--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=32
--mkdirs elbencho-s3-bucket1-100g
--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=32
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-100g
--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=32
--timelimit=300 --infloop elbencho-s3-bucket1-100g


# Data size per workload = 16 jobs x 160m x 20 clients x 1 dirs x 200 files = 10T.
--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=200
--mkdirs elbencho-s3-bucket1-100g
--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-100g
--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket1-100g


# Data size per workload = 16 jobs x 160m x 20 clients x 1 dirs x 400 files = 24T.
```

```
--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=480
--mkdirs elbencho-s3-bucket1-100g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=480
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-100g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=480
--timelimit=300 --infloop elbencho-s3-bucket1-100g
```

5.9.3 BW_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 100 clients x 1 dirs x 6 files = 1.5T.

--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=6
--mkdirs elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=6
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=6
--timelimit=300 --infloop elbencho-s3-bucket1-10g


# Data size per workload = 16 jobs x 160m x 100 clients x 1 dirs x 40 files = 10T.

--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=40
--mkdirs elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=40
--timelimit=300 --infloop elbencho-s3-bucket1-10g


# Data size per workload = 16 jobs x 160m x 100 clients x 1 dirs x 96 files = 24T.

--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=96
--mkdirs elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=96
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket1-10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=96
--timelimit=300 --infloop elbencho-s3-bucket1-10g
```

5.9.4 IOPS_TEST_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 4 jobs x 1280m x 20 clients x 1 dirs x 5 files = 512G.
# Total data for read and write < 4T, both data will be held in Optane
# Create the files first, then read altogether.
#
# s3 protocol defines that:
#    (1) an upload object cannot consist of more than 10000 parts
#    (2) an upload object cannot be less than 5MiB
# VAST doesn't have limitation of (2), but does also have (1).
# Workaround: use 5MiB for write, but read with 4KiB.
--write --blockvarpct=100 --threads=4 --block=5m --size=1280m --dirs=1 --files=5
--mkdirs elbencho-s3-bucket1-100g
--read  --blockvarpct=100 --threads=4 --block=4k --size=1280m --dirs=1 --files=5
--s3randobj --timelimit=300 --infloop elbencho-s3-bucket1-100g
# Write 10T (> 2 * 4T) of data to fully destage the previous data to QLC
--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=1 --files=200
--mkdirs elbencho-s3-bucket2-100g
# Now re-read the data written before
--read --blockvarpct=100 --threads=4 --block=4k --size=1280m --dirs=1 --files=5
--s3randobj --timelimit=300 --infloop elbencho-s3-bucket1-100g
```

5.9.5 IOPS_TEST_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 4 jobs x 1280m x 100 clients x 1 dirs x 1 file = 512G.
# Total data for read and write < 4T, both data will be held in Optane
# Create the files first, then read altogether.
```

```
#
# s3 protocol defines that:
#    (1) an upload object cannot consist of more than 10000 parts
#    (2) an upload object cannot be less than 5MiB
# VAST doesn't have limitation of (2), but does also have (1).
# Workaround: use 5MiB for write, but read with 4KiB.
--write --blockvarpct=100 --threads=4 --block=5m --size=1280m --dirs=1 --files=1
--mkdirs elbencho-s3-bucket1-10g
--read  --blockvarpct=100 --threads=4 --block=4k --size=1280m --dirs=1 --files=1
--s3randobj --timelimit=300 --infloop elbencho-s3-bucket1-10g
# Write 10T (> 2 * 4T) of data to fully destage the previous data to QLC
--write --blockvarpct=100 --threads=16 --block=32m --size=160m --dirs=40 --files=1
--mkdirs elbencho-s3-bucket2-10g
# Now re-read the data written before
--read  --blockvarpct=100 --threads=4 --block=4k --size=1280m --dirs=1 --files=1
--s3randobj --timelimit=300 --infloop elbencho-s3-bucket1-10g
```

5.9.6 SWEEP_IOSIZES_RANDOM_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 20 clients x 1 dirs x 200 files = 10T.
--write --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=200
--mkdirs elbencho-s3-bucket-sweep-iosizes-100g
# Random read
--read  --blockvarpct=100 --threads=16  --block=32m  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g
--read  --blockvarpct=100 --threads=16  --block=16m  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g
--read  --blockvarpct=100 --threads=16  --block=8m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g
--read  --blockvarpct=100 --threads=16  --block=5m   --size=160m  --dirs=1 --files=200
```

```
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=4m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=2m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=1m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=512k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=256k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=128k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=64k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=32k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=16k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=8k   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=4k   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-100g
```

5.9.7 SWEEP_IOSIZES_RANDOM_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 100 clients x 1 dirs x 40 files = 10T.

--write --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=40
--mkdirs elbencho-s3-bucket-sweep-iosizes-10g

# Random read

--read  --blockvarpct=100 --threads=16  --block=32m  --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g
```

```
--read  --blockvarpct=100 --threads=16  --block=16m  --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=8m   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=5m   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=4m   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=2m   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=1m   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=512k --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=256k --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=128k --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=64k  --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=32k  --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=16k  --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=8k   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=4k   --size=160m  --dirs=1 --files=40
--timelimit=300 --infloop --s3randobj elbencho-s3-bucket-sweep-iosizes-10g
```

5.9.8 SWEEP_IOSIZES_SEQUENTIAL_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 20 clients x 1 dirs x 200 files = 10T.
```

```
--write --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=200
--mkdirs elbencho-s3-bucket-sweep-iosizes-100g

# Sequential read

--read  --blockvarpct=100 --threads=16  --block=32m  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=16m  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=8m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=5m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=4m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=2m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=1m   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=512k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=256k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=128k --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=64k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=32k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=16k  --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=8k   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g

--read  --blockvarpct=100 --threads=16  --block=4k   --size=160m  --dirs=1 --files=200
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-100g
```

5.9.9 SWEEP_IOSIZES_SEQUENTIAL_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 16 jobs x 160m x 100 clients x 1 dirs x 40 files = 10T.
--write --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --threads=40
--mkdirs elbencho-s3-bucket-sweep-iosizes-10g
# Sequential read
--read  --blockvarpct=100 --threads=16  --block=32m  --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=16m  --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=8m   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=5m   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=4m   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=2m   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=1m   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=512k --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=256k --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=128k --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=64k  --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=32k  --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
--read  --blockvarpct=100 --threads=16  --block=16k  --size=160m  --dirs=1 --threads=40
```

```
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=8k   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g

--read  --blockvarpct=100 --threads=16  --block=4k   --size=160m  --dirs=1 --threads=40
--timelimit=300 --infloop elbencho-s3-bucket-sweep-iosizes-10g
```

5.9.10 SWEEP_THREADS_RANDOM_32M_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 20 clients x 1 dirs x 100 files = 10T.

--write --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=100
--mkdirs elbencho-s3-bucket-sweep-threads-100g

# Random read

--read  --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g

--read  --blockvarpct=100 --threads=8  --block=32m --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g

--read  --blockvarpct=100 --threads=4  --block=32m --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g

--read  --blockvarpct=100 --threads=1  --block=32m --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
```

5.9.11 SWEEP_THREADS_RANDOM_32M_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 100 clients x 1 dirs x 20 files = 10T.
```

```
--write --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=20
--mkdirs elbencho-s3-bucket-sweep-threads-10g

# Random read

--read  --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g

--read  --blockvarpct=100 --threads=8  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g

--read  --blockvarpct=100 --threads=4  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g

--read  --blockvarpct=100 --threads=1  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g
```

5.9.12 SWEEP_THREADS_RANDOM_4K_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 20 clients x 1 dirs x 100 files = 10T.

# s3 protocol defines that:

#    (1) an upload object cannot consist of more than 10000 parts

#    (2) an upload object cannot be less than 5MiB

# VAST doesn't have limitation of (2), but does also have (1).

# Workaround: use 5MiB for write, but read with 4KiB.

--write --blockvarpct=100 --threads=32 --block=5m --size=160m  --dirs=1 --files=100
--mkdirs elbencho-s3-bucket-sweep-threads-100g

# Random read

--read  --blockvarpct=100 --threads=32 --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
```

```
100g
--read  --blockvarpct=100 --threads=16 --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
--read  --blockvarpct=100 --threads=8  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
--read  --blockvarpct=100 --threads=4  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
--read  --blockvarpct=100 --threads=2  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
--read  --blockvarpct=100 --threads=1  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
100g
```

5.9.13 SWEEP_THREADS_RANDOM_4K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 100 clients x 1 dirs x 20 files = 10T.
# s3 protocol defines that:
#    (1) an upload object cannot consist of more than 10000 parts
#    (2) an upload object cannot be less than 5MiB
# VAST doesn't have limitation of (2), but does also have (1).
# Workaround: use 5MiB for write, but read with 4KiB.
--write --blockvarpct=100 --threads=32 --block=5m --size=160m  --dirs=1 --files=20
--mkdirs elbencho-s3-bucket-sweep-threads-10g
# Random read
--read  --blockvarpct=100 --threads=32 --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-
10g
--read  --blockvarpct=100 --threads=16 --block=4k --size=160m  --dirs=1 --files=20
```

```
--timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=8  --block=4k --size=160m  --dirs=1 --files=20 --timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=4  --block=4k --size=160m  --dirs=1 --files=20 --timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=2  --block=4k --size=160m  --dirs=1 --files=20 --timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=1  --block=4k --size=160m  --dirs=1 --files=20 --timelimit=300 --infloop --randamount=1p --s3randobj elbencho-s3-bucket-sweep-threads-10g
```

5.9.14 SWEEP_THREADS_SEQUENTIAL_32M_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause

# Data size per workload = 32 jobs x 160m x 20 clients x 1 dirs x 100 files = 10T.

--write --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=100 --mkdirs elbencho-s3-bucket-sweep-threads-100g

# Sequential read

--read  --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=8  --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=4  --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=2  --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=1  --block=32m --size=160m  --dirs=1 --files=100 --timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g
```

5.9.15 SWEEP_THREADS_SEQUENTIAL_32M_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 100 clients x 1 dirs x 20 files = 10T.
--write --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=20
--mkdirs elbencho-s3-bucket-sweep-threads-10g
# Sequential read
--read  --blockvarpct=100 --threads=32 --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
--read  --blockvarpct=100 --threads=16 --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
--read  --blockvarpct=100 --threads=8  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
--read  --blockvarpct=100 --threads=4  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
--read  --blockvarpct=100 --threads=2  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
--read  --blockvarpct=100 --threads=1  --block=32m --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
```

5.9.16 SWEEP_THREADS_SEQUENTIAL_4K_100G.SCH

```
# Intel Copyright © 2022, Intel Corporation.
# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 20 clients x 1 dirs x 100 files = 10T.
# s3 protocol defines that:
#    (1) an upload object cannot consist of more than 10000 parts
#    (2) an upload object cannot be less than 5MiB
# VAST doesn't have limitation of (2), but does also have (1).
# Workaround: use 5MiB for write, but read with 4KiB.
--write --blockvarpct=100 --threads=32 --block=5m --size=160m  --dirs=1 --files=100
```

```
--mkdirs elbencho-s3-bucket-sweep-threads-100g

# Sequential read

--read  --blockvarpct=100 --threads=32 --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=16 --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=8  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=4  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=2  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g

--read  --blockvarpct=100 --threads=1  --block=4k --size=160m  --dirs=1 --files=100
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-100g
```

5.9.17 SWEEP_THREADS_SEQUENTIAL_4K_10G.SCH

```
# Intel Copyright © 2022, Intel Corporation.

# SPDX-License-Identifier: BSD-3-Clause


# Data size per workload = 32 jobs x 160m x 100 clients x 1 dirs x 20 files = 10T.

# s3 protocol defines that:

#    (1) an upload object cannot consist of more than 10000 parts

#    (2) an upload object cannot be less than 5MiB

# VAST doesn't have limitation of (2), but does also have (1).

# Workaround: use 5MiB for write, but read with 4KiB.

--write --blockvarpct=100 --threads=32 --block=5m --size=160m  --dirs=1 --files=20
--mkdirs elbencho-s3-bucket-sweep-threads-10g

# Sequential read

--read  --blockvarpct=100 --threads=32 --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=16 --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=8  --block=4k --size=160m  --dirs=1 --files=20
```

```
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=4  --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=2  --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g

--read  --blockvarpct=100 --threads=1  --block=4k --size=160m  --dirs=1 --files=20
--timelimit=300 --infloop elbencho-s3-bucket-sweep-threads-10g
```