Vdbench Tips

- Calculate the minimum number of total threads required:

(# of load generators) X (# of disks per load generator) X (# of JVMs per load generator)

- Total # of threads defined depends on where in the job file it is specified:
  - If specified in the sd=default, and only one disk is defined, the value is the total # of threads for the entire test
  - If specified in the sd=default, and multiple disks are defined, it applies to all disks and the total is the sum of all values
  - If specified per disk, the total # of threads is the sum of all values
- Calculate the OIO per disk:

(total threads defined) / ((# of load generators) X (# of disks per load generator))

- Each JVM can service ~100K IOPs, so specifying jvms=1 makes thread calculation easier
- Build a load generator template VM with the SSH keys of the master for no password logins
- Multi-homed master systems must resolve the load generator facing IP to its hostname via /etc/hosts
- Redirect output to a working web server folder, and use .tod for time of day stamps in the folder name

HX Hybrid to All-Flash Performance Comparison

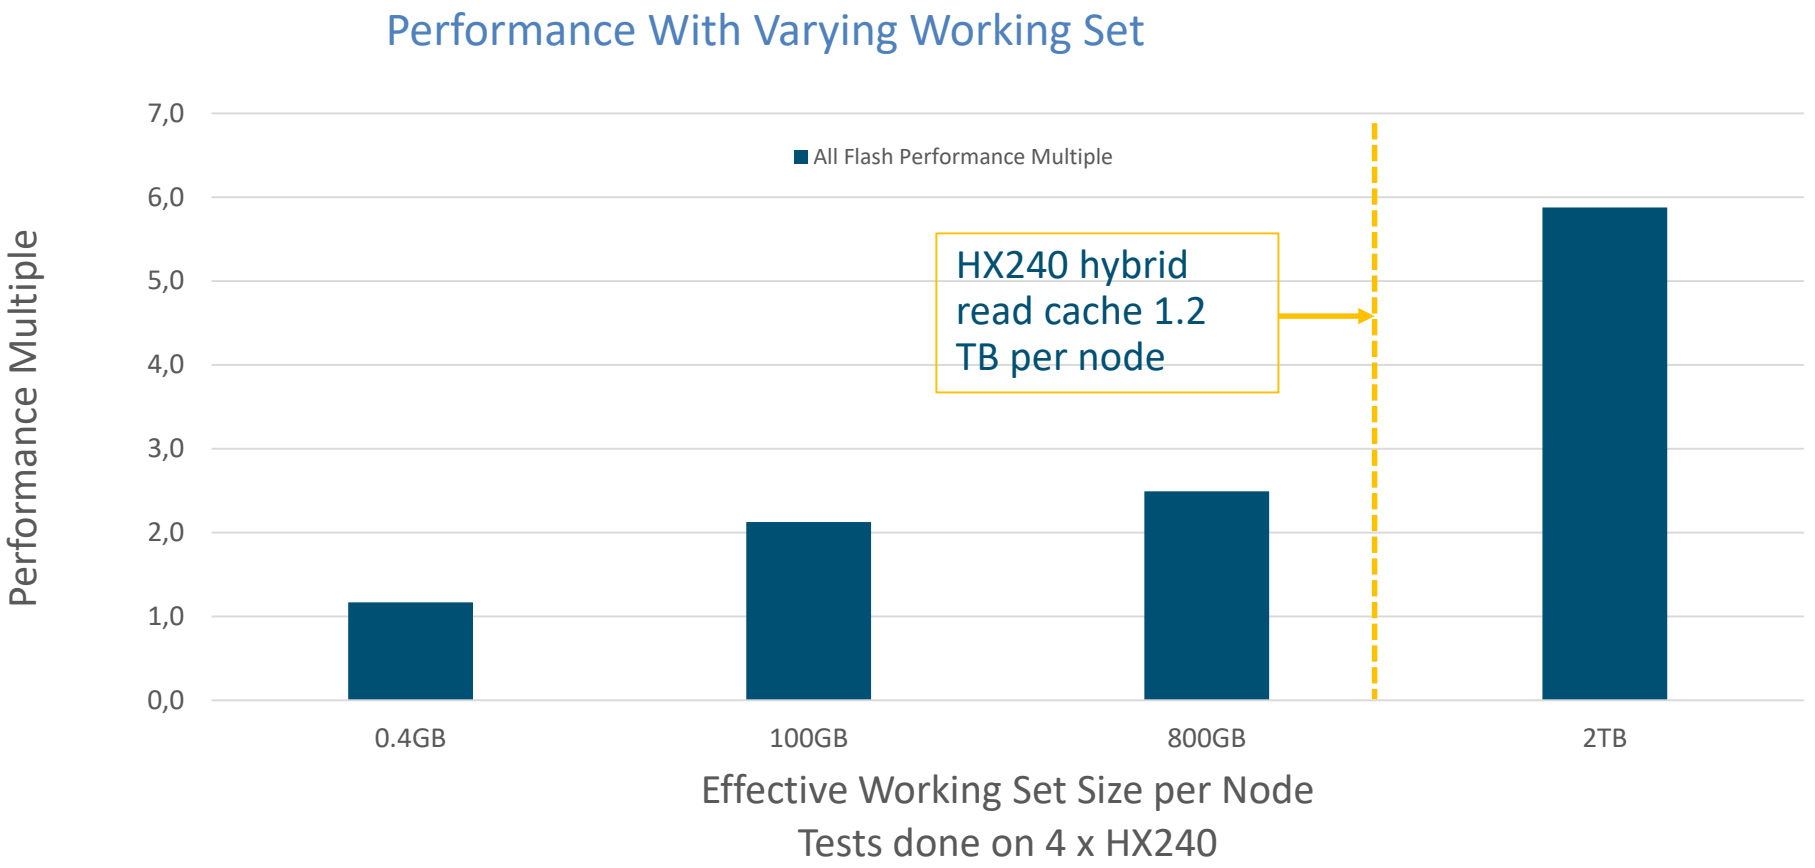

- AF offers performance improvements and lower latency even at small working sets
- AF offers huge performance and latency improvements for working sets that do not fit in the hybrid model read cache

Base Configurations for Vdbench Tests

	HX All-Flash	HX All-Flash + 40GbE	HX All-Flash + 40GbE and NVMe	HX All-Flash M5
Version	HX 2.0.1a	HX 2.0.1a	HX 2.5.1a	HX 2.6.1b
HW Configuration	8 x HXAF240 M4, each with: • WL: 800G SSD • Data: 10 x 960G • VIC 1227 • E5-2690v4 • 512 GB RAM	8 x HXAF240 M4, each with: • WL: 800G SSD • Data: 10 x 960G • VIC 1387 • E5-2690v4 • 512 GB RAM	8 x HXAF240 M4, each with: • WL: 400G NVMe • Data: 10 x 960G • VIC 1387 • E5-2690v4 • 512 GB RAM	8 x HXAF240 M5, each with: • WL: 400G NVMe • Data: 10 x 960G • VIC 1387 • Xeon 6154 • 384 GB RAM
Benchmark Setup	• Vdbench 5.04.06 • 8 Linux VMs per host (64 total) • Each VM has 1 x 200GB virtual disk, HX filesystem ~60% full (12.4 TB working set) • Priming: 64K writes 100% sequential • Tests run 3 times and averaged			

Benchmark 1

Performance Results Using Vdbench Tests

8K 100% Sequential Write Workload

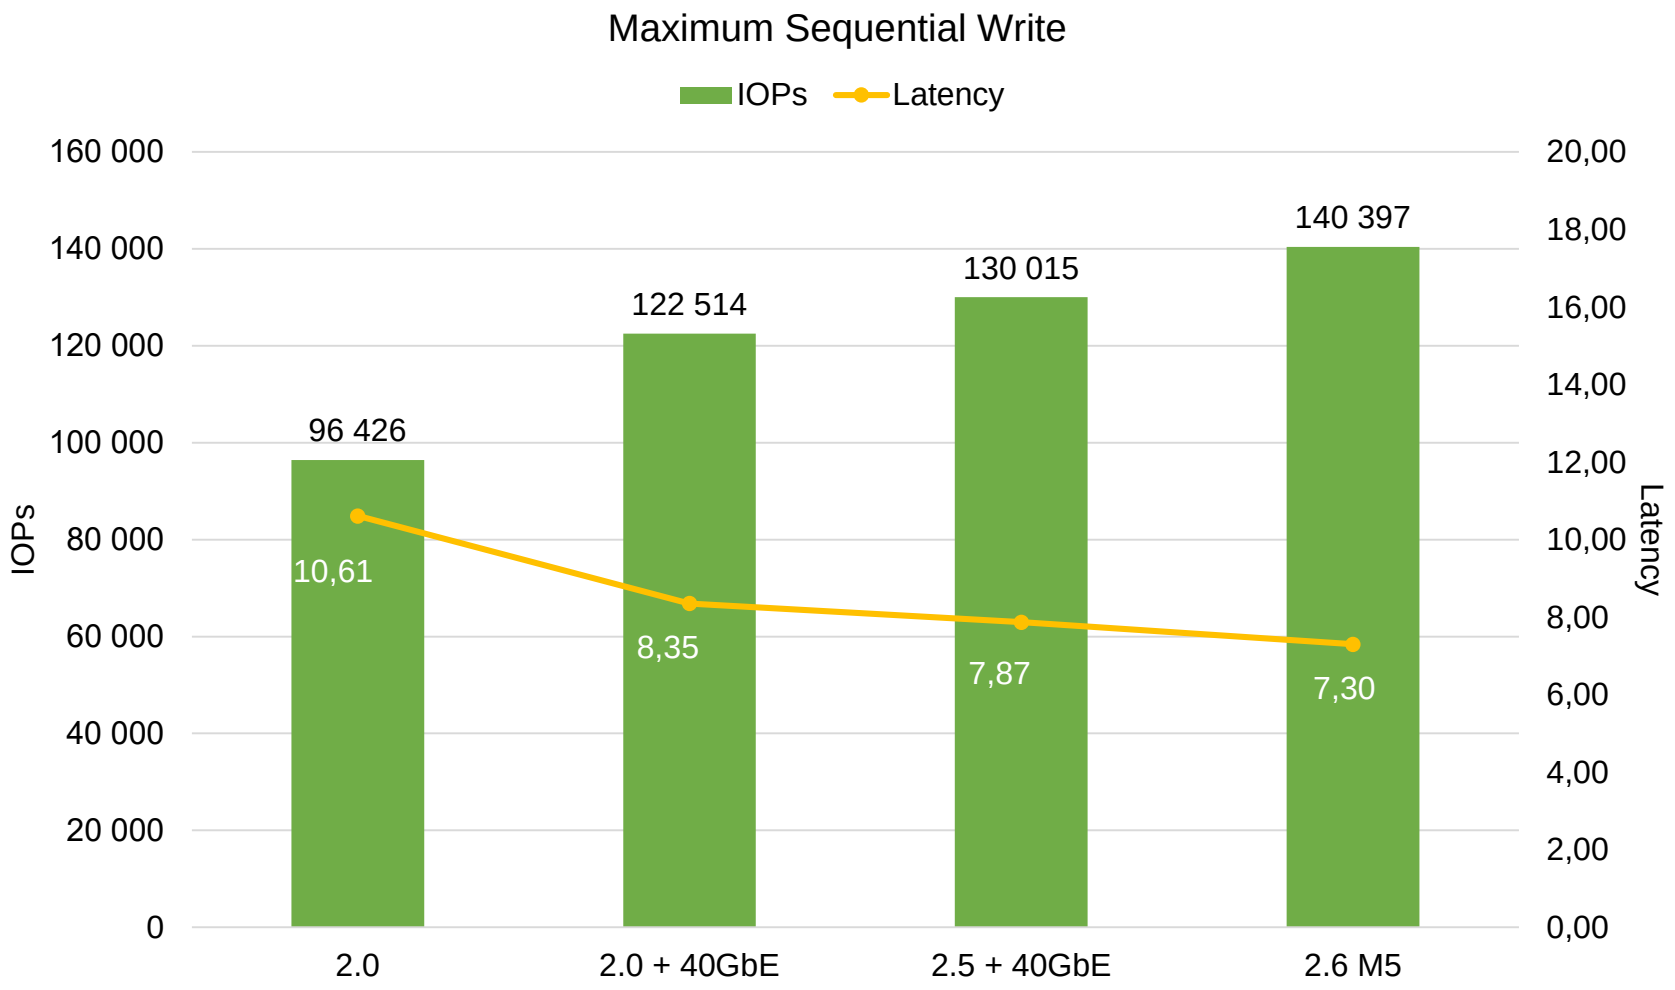

- Target Write Latency < 10ms
- 16 OIO Per Disk
- 27% Latency Reduction and IOPs Gain with 40 GbE
- 6% Latency Reduction and IOPs Gain on 2.5 + NVMe
- 8% Latency Reduction and IOPs Gain with HX 2.6 on M5

Benchmark 2

Performance Results Using Vdbench Tests

8K 100% Random Read Workload

- Target Read Latency < 5ms
- 32 OIO Per Disk
- 7% Latency Reduction and IOPs Gain with 40 GbE
- 23% Latency Reduction and IOPs Gain on 2.5 + NVMe
- 3% Latency Reduction and IOPs Gain with HX 2.5 on M5

Benchmark 3

Performance Results Using Vdbench Tests

100% Random 70/30% R/W with Varying Block Sizes

	4K			8K			16K			64K			256K		
	IOPs	R Latency	W Latency	IOPs	R Latency	W Latency	IOPs	R Latency	W Latency	IOPs	R Latency	W Latency	IOPs	R Latency	W Latency
2.0	150,673	2.12	6.36	130,511	2.40	7.47	92,130	3.49	10.35	34,965	10.75	23.70	10,079	45.60	62.80
2.0 + 40 GbE	152,434	2.22	6.00	142,721	2.36	6.47	111,905	3.13	8.03	51,146	6.96	17.16	19,712	11.68	59.28
2.5 + 40 GbE	161,131	2.24	5.37	160,290	2.09	5.75	136,433	2.45	6.78	56,339	6.30	15.74	20,036	9.17	63.61
2.6 M5	180,037	1.72	5.49	173,193	1.91	5.39	140,842	2.36	6.61	61,227	5.45	15.18	20,732	8.96	61.37

- Target Read Latency < 5ms and Write Latency < 10ms for 4-16K blocks
- 8 OIO Per Disk
- ~6% Latency Reduction and IOPs Gain with 40 GbE at 4-8K, 50-100% Gains for Larger Blocks
- 10-22% Latency Reductions and IOPs Gains with HyperFlex 2.5 and NVMe Caching Disks
- 3-11% Latency Reductions and IOPs Gains with HyperFlex 2.6 on M5 Generation Servers

Benchmark 4

Performance Results Using Vdbench Tests

8K 100% Random 70/30% R/W IOPs Curve

	100%			20%			40%			60%			80%		
	IOPs	IOPs STDV	Avg Latency	IOPs	IOPs STDV	Avg Latency	IOPs	IOPs STDV	Avg Latency	IOPs	IOPs STDV	Avg Latency	IOPs	IOPs STDV	Avg Latency
2.0	138,190	19.0%	3.71	27,666	0.5%	1.07	55,299	0.5%	1.57	82,960	0.5%	2.20	110,594	6.6%	3.14
2.0 + 40 GbE	151,568	18.6%	3.38	30,331	0.5%	0.99	60,621	0.5%	1.37	90,957	0.5%	2.08	121,250	2.3%	2.63
2.5 + 40 GbE	169,983	8.6%	3.01	34,031	0.4%	0.86	68,058	0.5%	1.17	102,023	0.5%	1.66	136,018	0.5%	2.28
2.6 M5	173,193	6.2%	2.95	34,699	0.2%	0.80	69,331	0.2%	1.21	103,953	0.2%	1.68	138,580	0.3%	2.38

- Curve test measures IOPs stability by calculating the standard deviation of IOPs as a percentage of the final average result, at various performance levels, starting with a 100% unthrottled test, then again at a percentage of that maximum
- Maximum speed, unthrottled tests will always have more variability due to more aggressive background tasks
- 20% through 60% values are typically very low because the system is not being pushed very hard
- 65% improvement in IOPs stability with 40 GbE at 80% of maximum HX 2.5 shows a further 78% improvement
- HX 2.5 also shows a 54% improvement of IOPs stability in the 100% unthrottled test
- HX 2.6 on M5 generation servers shows improvements in all tests, including 28% at 100%, and 45% at 80%