



To SSD or not SSD: That's a POWERful Vertica question.

This project explores the performance of spinning and solid state disks, including CPU energy and heat. I/O kernels (fio, vioperf) and standard disk utility MKFS are examined, as well as loading 500 million rows (125GB of data) into a Vertica database. The following lists the MKFS and kernel results.

	mkfs	mkfs	fio	fio	fio	fio	vioperf	vioperf	vioperf	vioperf
	elapsed	mb/sec	randread	randread	randrw	randrw	writes	rewrites	reads	seeks
			iops	elapseds	iops	elapseds	mb/sec	mb/sec	mb/sec	per sec
ssd2tb-r			11021	23	20209	12	504	219/219	531	56431
ssd2tb-v	106	18.87	5553	47	8919	29	202	214/214	176	27612
esata2tb	353	5.67	177	1473	250	1045	71	32/32	71	182
wdc2tb	398	5.03	171	1525	231	1129	100	39/39	80	377
usb2tb	840	2.38	160	1632	214	1220	32	15/15	32	251
green2tb			140	1870	195	1339	69	29/29	57	289
esata3tbr	1114	1.80	124	2108	76	3432	35	18/18	66	122
vmhome			71	3688	86	3036	33	8/8	35	251

Figure 1 – MKFS, FIO, VIOPERF Results

Two real machines are used in this study, as well as AWS. The desktop runs RHEL 7.0/Vertica 7 native, is an ASUS CM6870 (i7-3770@3.39GHZ, [Geekbench](#) 3745/11744) with 32GB of RAM, 8 logical processors, 4TB of internal disk (wdc2tb, green2tb) and 5TB of esata/ext3 disk (estat2tb, esata3tbr raid).

job	machine	elapsed seconds
create100m	asusrh7	254
	aws8p32g	233
	ssd-r	164
load100m	asusrh7	344
	aws40p160g	230
	ssd-r	117
dbd	asusrh7	245
	aws16p64g	215
	ssd-r	210
load400m	asusrh7	1086
	aws40p160g	680
	ssd-r	799
vsq171q	asusrh7	6.682
	aws40p160g	5.509
	ssd-r	5.390
jdbc50	asusrh7	140
	aws16p64g	99
	ssd-r	84
jdbc50	asusrh7	123
	aws16p64g	96
	ssd-r	83
backup	asusrh7	123
	aws16p64g	140
	ssd-r	24
restore	asusrh7	132
	aws16p64g	132
	ssd-r	18

The laptop is an ASUS G750JM (i7-4700HQ @2.40GHZ, [GBCPU](#) 3701/10757), [intelgpu](#), [nvidiagpu](#), with 32GB of RAM, 8 logical processors, a 1TB internal disk (vmhome), a 2TB external USB drive (usb2tb) and a 2TB Samsung 850 PRO solid state drive (ssd2tb). Debian 8/Vertica 7 runs in a VM on the laptop with 8 processors, 24GB ram, a 40G partition (vmhome), and volumes ssd2tb-v (SCSI), and usb2tb. RHEL70 also runs native on the laptop, and owns (boots off) the entire ssd2tb-r volume.

Figure 2 – Benchmark Jobs Elapsed Times




Figure 3 – CPU Total All Nine Jobs

The graphs are for the native RHEL70/solid state setup on the laptop, and illustrate the CPU/DISK/ENERGY requirements for the nine jobs listed in Figure 2.




Figure 4 – Disk Busy All Nine Jobs




Figure 5 – Disk Writes All Nine Jobs




Figure 6 – Disk Reads All Nine Jobs

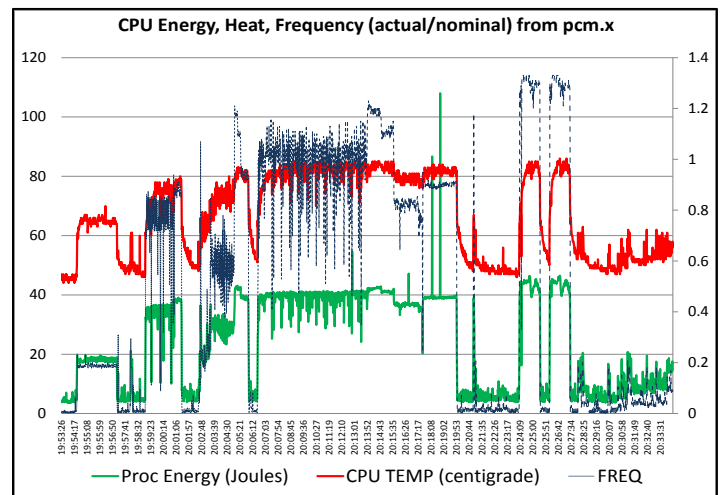


Figure 7 – CPU Energy, Heat, Frequency Ratio All Nine Jobs

We are running out of CPU ([nominal frequency 2.4ghz](#)), redlining for a lot of the time, while there is plenty of capacity left on the SSD. The first draft of the full study can be found at <http://davidyoung.com/cm/pcfverticaonaws.pdf>