



I7-7700 vs I7-4700: Setquery on Vertica, Couchbase, Mongodb

Big Data on Real Hardware, Amazon Web Services, and Google Cloud Platform

The following table summarizes elapsed seconds for creating, loading, and processing 500,000,000 rows (documents) of Vertica/Setquery data (127,279,550,000 bytes, 118.54 GBYTES) on real hardware, AWS, and GCP platforms.

	I7-4700HQ	I7-7700	E5-2676	E5-2676	E5-2676	Xeon	Xeon	
	pc ssd	pc ssd	aws rbs	aws rbs	aws rbs	gcp rbs	gcp ssd	
	8p32g	8p32g	8p32g	16p64g	40p160g	8p30g	8p30g	7700/4700
create100m	164	111	233	234	233	214	206	0.68
load100m	117	78	269	237	230	314	174	0.67
dbd	210	150	241	215	242	265	264	0.71
load400m	799	561	1322	756	680	1542	1026	0.70
vsq171q	5.39	4.53	7.405	5.509	6.244	8.309	7.147	0.84
jdbc50	84	68	112	99	216	129	128	0.81
jdbc50	83	66	113	96	220	127	122	0.80
backup	24	47	163	140	149	180	59	1.96
restore	18	36	137	132	138	178	53	2.00
qph	47,421	56,424	34,517	46,397	40,935	30,762	35,763	1.19
loadmbph	450,563	641,711	272,315	476,190	529,412	233,463	350,877	1.42
jdbcqph	123,614	155,455	90,796	106,875	46,636	80,787	84,098	1.26
backupgbph	2004	1023	295	344	323	267	815	0.51
restoregbph	2672	1336	351	364	349	270	907	0.50

Figure 1 – Vertica 500,000,000 Document Statistics

The following summarizes single user VSQL queries/hour, loaded database Mbytes/hour, and multi-threaded JDBC queries/hour (50 threads) for the measured platforms.

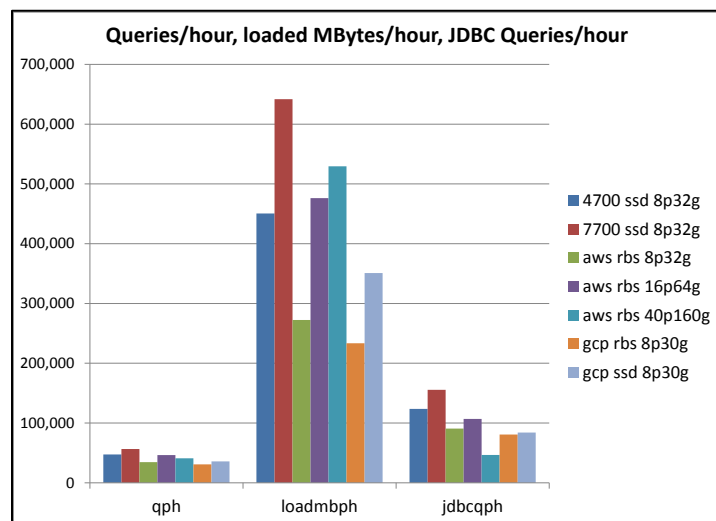


Figure 2 – Vertica VSQL, Database Load, JDBC performance

The I7-7700 loaded 42% more data, processed 19% more VSQL, and 26% more JDBC queries than the I7-4700. For details see <http://davidyoung.com/cmawgcpreal.pdf>.

NoSQL I7-7700 Performance Out Of The Box: Hbase, Couchbase, Mongodb

The following tables summarize one and eight user performance: 100,000 docs for Couchbase, 10,000,000 docs for MongoDB. ETR is the number of transactions per hour.

ITR factors in leftover CPU to project a theoretical maximum transaction rate.

	couchbase	couchbase	mongodb	mongodb
	1 user	8 user	1 user	8 user
	sq1100k	sq1100k	sq110m	sq110m
start	14:01:01	14:05:00	15:10:00	15:24:00
stop	14:03:41	14:18:08	15:21:00	15:39:25
elapsed	0:02:40	0:13:08	0:11:00	0:15:25
elapsed	160	788	660	925
tran	61	488	61	488
etr/hour	1373	2229	333	1899
cpu busy	0.472	0.9	0.113	0.932
itr/hour	2908	2477	2944	2038

Figure 3 – I7-4700 NoSQL Single and Multiple users

	couchbase	couchbase	mongodb	mongodb
	1 user	8 user	1 user	8 user
	sq1100k	sq1100k	sq110m	sq110m
start	7:30:00	7:37:00	8:11:33	8:25:00
stop	7:31:44	7:45:05	8:21:50	8:35:45
elapsed	0:01:44	0:08:05	0:10:17	0:10:45
elapsed	104	485	617	645
tran	61	488	61	488
etr/hour	2112	3622	356	2724
cpu busy	0.444	0.883	0.116	0.93
itr/hour	4756	4102	3068	2929
7700/4700	1.64	1.66	1.04	1.44

Figure 4 – I7-7700 NoSQL Single and Multiple users

In a nutshell, for this NoSQL query workload the I7-7700 performs between 44% and 66% more work, at roughly 2/3 the response time of the I7-4700.

	vcuperf	bogomips	gb single	gb multiple	coremark
4700	9.18	4789	3749	12233	17497
7700	7.5	7197	5091	17224	21953
7700/4700	0.82	1.50	1.36	1.41	1.25
1/(7700/4700)	1.22				

Figure 5 – CPU Kernels I7-4700 vs I7-7700 ([coremark link](#))

For details see <http://davidyoung.com/cmawgcpreal.pdf>

Software/Hardware Configuration

The benchmark machine is an [ASUS G750JM](#) (i7-4700HQ @2.40GHZ, [gbcpu](#) 3620/11722, [iopen](#), [nopen](#), [ncuda](#)) with 32GB of RAM (DDR3@0.799GHZ), a 2TB internal Samsung 850 PRO SSD, and a 1TB usb 3.0 external Samsung Portable SSD T3. An [ASUS G11CD-K](#) I7-7700 @ 3.60GHZ was also measured ([gbcpu](#) 5091/17224, [opencl](#), [cuda](#)), with a 2TB Samsung 850 PRO and 32 GB of RAM (DDR4@2.133GHZ). RHEL70 was installed on the 2TB Internal SSD (Server with GUI), no LVM.

References Jim Gray (Ed.): The Benchmark Handbook for Database and Transaction Systems (2nd Edition). Morgan Kaufmann 1993, ISBN 1-55860-292-5