Planning Capacity for Vertica on AWS

This research project measures the differences in application execution times and resource requirements for standard Vertica V7 database functions (7.1.2-4 and 7.2.1-0) in RedHat Enterprise Linux 5.9, 6.7, and 7.0 environments. The Setquery Benchmark (Dr. Patrick O'Neil University of Massachusetts, Boston) is investigated, NMON 15e & Vtune 2013/2016 are used for performance data recording and analysis. SQL Server 2016 on Windows 8 is used to investigate Import performance via ado.net. Part II examines the same applications in AWS/EC2 environments.

1) create100M - c++ program creates 100 million rows of input data, eclipse kepler and mars
2) load100M - load 100 million rows into setquery table sq1
3) dbd - run the data base designer, query performance
4) load400M - load 400 million rows into setquery table sq1
5) vsql71 - execute seventy one queries against the 500 million row sq1 table
6) jdbc50 - execute multithreaded jdbc application, with fifty concurrent threads
7) jdbc50 - execute multithreaded jdbc application with fifty concurrent threads a second time
8) backup - backup 500 million rows
9) restore - restore 500 million rows.
10) import200M - import 200 million rows to SQL Server 2016/Windows 8 via ado.net

Figure 1 – Benchmark Tasks

The benchmark tasks are listed preceding in Figure 1. For each measurement run, the server is rebooted, and tasks 2 through 9 are run in succession. For this workload, in general elapsed and cpu times get better for each OS upgrade, as well as the Vertica upgrade from 7.1 to 7.2. Backup/restore has been revamped for 7.2.1-0, notice the significant improvements in elapsed times. DBD processing has also improved for 7.2, the more intense the DBD process, the more the improvements. The measurement results are summarized following in Figures 2, 3, and 4.

		elapsed	elapsed	cpu	cpu	cpu busy	io/sec	MB	MB	vtune	vtune	vtune	vtune #ins
		seconds	ratios	seconds	ratios	avg	avg	read	written	cpi	elapsed	cpu	billions
create100m	r59	463		248		0.079	318	70	27,481	0.545	471.000	248.375	1448
	r67	458	0.99	240	0.97	0.062	146	60	27,587	0.541	471.000	240.445	1492
	r70	254	0.55	190	0.76	0.107	243	50	26,414	0.547	254.000	205.185	1398
	r70e	0	0.00	0	0.00	0	0	0	0	0.000	0.000	0.000	0.000
	aws16p64g	234	0.51	224	0.90	0.057	903	1	27,538				
load100m	r59	321		776		0.293	257	28,751	6,826	0.811	15.062	36.619	152
	r67	333	1.04	722	0.93	0.283	258	27,803	6,759	0.714	15.019	30.615	145
	r70	319	0.99	648	0.84	0.25	306	27,873	9,251	0.709	15.007	27.332	142
	r70e	344	1.07	591	0.76	0.216	290	28,020	6,808	0.712	15.016	26.981	140
	aws16p64g	233	0.73	690	0.89	0.179	1265	27,357	8,716				
dbd	r59	268		807		0.364	104	45	6,525	0.927	15.062	26.916	99
	r67	262	0.98	833	1.03	0.393	115	31	6,402	0.867	15.017	28.320	111
	r70	264	0.99	714	0.88	0.338	153	57	6,716	0.812	15.017	16.040	74
	r70e	245	0.91	694	0.86	0.36	170	76	6,755	1.066	15.011	14.002	50
	aws16p64g	215	0.80	755	0.94								
load400m	r59	1547		4895		0.396	289	122,543	33,618	0.720	15.052	59.413	280
	r67	1229	0.79	4581	0.94	0.474	315	123,479	33,068	0.702	15.019	53.969	261
	r70	1090	0.70	4056	0.83	0.466	383	124,197	33,860	0.672	15.015	51.026	281
	r70e	1086	0.70	3847	0.79	0.442	396	124,412	34,315	0.700	15.014	46.525	245
	aws16p64g	749	0.48	4645	0.95	0.377	1655	119,289	32,973				
vsql71q	r59	8.541		32.507						1.092	15.052	32.507	100
	r67	8.250	0.97	31.17	0.96					1.067	15.022	31.170	99
	r70	7.009	0.82	26.60	0.82					1.043	30.030	26.510	93
	r70e	6.682	0.78	26.39	0.81					1.079	30.030	26.390	89
	aws16p64g	5.764	0.67										

Figure 2 – Measurement Results Part 1

Ratios for Elapsed seconds (r67/r59, r70/r59, r70e/r59) are in the blue column, ratios for CPU seconds are in the green column. CPU seconds are computed using average CPU busy from NMON. Most of the VTUNE sampling runs are 15 seconds in duration (advanced hot spots), elapsed VTUNE time for the import200m step is 60 seconds, and elapsed VTUNE time for the create100m step is the length of the step.

		elapsed	elapsed	cpu	cpu	cpu busy	io/sec	MB	MB	vtune	vtune	vtune	vtune #ins
		seconds	ratios	seconds	ratios	avg	avg	read	written	cpi	elapsed	cpu	billions
jdbc50	r59	147		733		0.607	222	398	353	0.832	15.053	112.571	459
	r67	116	0.79	781	1.07	0.814	216	471	313	0.790	15.018	119.680	514
	r70	128	0.87	542	0.74	0.487	243	544	326	0.788	15.015	93.562	439
	r70e	140	0.95	585	0.80	0.508	311	511	518	0.835	15.014	87.288	386
	aws16p64g	99	0.67	1224	1.67	0.702	847	417	816				
jdbc50	r59	116		777		0.771	240	12	334	0.855	15.059	113.030	448
	r67	99	0.85	770	0.99	0.935	233	17	315	0.792	15.072	119.695	512
	r70	107	0.92	560	0.72	0.614	228	63	385	0.726	15.021	119.970	610
	r70e	123	1.06	572	0.74	0.591	260	82	528	0.816	15.016	107.774	488
	aws16p64g	96	0.83	1191	1.53	0.73	811	76	836				
backup	r59	150		245		0.175	570	2,167	13,074	1.136	15.027	25.756	77
	r67	161	1.07	220	0.90	0.153	335	1,988	12,788	1.032	15.090	23.817	80
	r70	197	1.31	198	0.81	0.125	258	2,196	13,214	1.056	15.015	18.372	66
	r70e	123	0.82	175	0.71	0.166	258	2,144	13,037	0.879	15.015	22.681	100
	aws16p64g	140	0.93	226	0.92	0.094	836	2,036	12,719				
restore	r59	184		236		0.148	345	11	12,307	1.115	15.057	30.177	92
	r67	210	1.14	222	0.94	0.115	202	14	12,779	1.012	15.090	31.941	107
	r70	197	1.07	199	0.84	0.129	207	103	13,292	1.005	15.015	12.428	47
	r70e	132	0.72	174	0.74	0.165	245	86	12,909	0.896	15.015	23.634	101
	aws16p64g	132	0.72	228	0.97	0.099	698	51	12,692				
import200m	r59	2430		1499		0.077	38	4,392	453	0.821	59.630	56.126	227800
	r67	2447	1.01	1884	1.26	0.096	33	2,848	427	0.780	60.005	53.221	228616

Figure 3 – Measurement Results Part 2

The database loads are profiled, the following contains statistics from Vertica's execution_engine_profiles.

	elapsed	response wait	clock time	execution time	temp file	memory allocated
	seconds	seconds	seconds	seconds	bytes	bytes
asusrh7100m	344	231	2,146	373	4,036,458,091	463,392,621
aws8p32g	269	173	1,782	488	4,036,474,057	459,692,876
aws16p64g	233	136	2,916	619	4,036,462,858	632,519,364
aws40p160g	230	145	7,318	637	4,036,527,911	1,150,056,940
asusrh7400m	1086	551	6,542	2,767	23,469,828,553	837,478,006
aws8p32g	1316	600	7,955	3,418	23,492,031,614	830,774,166
aws16p64g	749	162	6,592	4,447	23,191,022,392	983,495,204
aws40p160g	680	117	12,742	5,320	23,324,700,523	1,318,921,250

Figure 4 – Database Loads Elapsed and CPU times

Vertica is housed on an ASUS CM6870 (i7-3770@3.39GHZ, Geekbench 3296/12896) with 32GB of RAM, 8 logical processors, and 5TB of esata/ext3 disk. SQL Server 2016 is on an ASUS gaming laptop G750JM (i7-4700HQ@2.40GHZ, Geekbench 3163/11,813) with 32GB of RAM, 8 logical processors, a 1TB internal disk and a 2TB external USB drive. They are connected to the network via a Netgear GS116E gigabit switch.

The first draft of the full study can be found at http://davidjyoung.com/cmg/pcfverticaonaws.pdf .