

NUMA Performance Results

Introduction

We have measured the NUMA patch performance improvement with the Big Bench by using the Hive-On-MR as execution engine. We ran different workloads (100G, 1TB) and below is the overall performance benefit for each workload.

- 100 GB Workload – 6.18%
- 1TB Workload – 15.14%

Query level execution time and improvements can be found in the below pages.

System Configurations

- 5 nodes (0-4)
- 0th node : NN, RM, DN, NM
- (1-4) nodes : DN, NM
- Each Node Configuration
 - OS - CentOS release 6.8
 - Main Memory – 250 GB
 - Processors – 88

BB – Hive On MR – 100G

QUERY No	QUERY Exec Time W/O NUMA(hh:mm:ss.msec)	QUERY Exec Time With NUMA (hh:mm:ss.msec)	Improvement (hh:mm:ss.msec)	Improvement % (Impr/Exec Time W/O NUMA)*100
1	0:09:36.330	0:08:09.121	0:01:27.209	15.13%
2	0:35:29.288	0:31:48.597	0:03:40.691	10.36%
3	0:21:03.414	0:19:56.293	0:01:07.121	5.31%
4	1:04:18.695	1:03:00.818	0:01:17.877	2.02%
5	0:37:22.037	0:35:28.722	0:01:53.315	5.05%
6	0:31:02.191	0:26:27.745	0:04:34.446	14.74%
7	0:28:47.012	0:27:13.727	0:01:33.285	5.40%
8	0:26:53.918	0:23:05.275	0:03:48.643	14.17%
9	0:22:58.277	0:23:18.300	-0:00:20.023	-1.45%
10	0:49:53.587	0:50:20.769	-0:00:27.182	-0.91%
11	0:08:26.100	0:08:16.272	0:00:09.828	1.94%
12	0:13:57.730	0:09:09.259	0:04:48.471	34.43%
13	0:17:03.804	0:17:29.660	-0:00:25.856	-2.53%
14	0:03:31.321	0:05:36.848	-0:02:05.527	-59.40%
15	0:07:59.925	0:08:04.127	-0:00:04.202	-0.88%

BB – Hive On MR – 100G (Cont)

QUERY No	QUERY Exec Time W/O NUMA (hh:mm:ss.msec)	QUERY Exec Time With NUMA (hh:mm:ss.msec)	Improvement (hh:mm:ss.msec)	(Imprmt/Exec Time W/O NUMA)*100
16	0:39:15.111	0:37:40.848	0:01:34.263	4.00%
17	0:06:48.830	0:06:44.047	0:00:04.783	1.17%
18	0:41:33.907	0:29:54.300	0:11:39.607	28.05%
19	0:15:38.666	0:14:55.559	0:00:43.107	4.59%
20	0:12:41.417	0:12:24.869	0:00:16.548	2.17%
21	0:14:37.699	0:13:35.918	0:01:01.781	7.04%
22	0:03:51.879	0:03:56.710	-0:00:04.831	-2.08%
23	0:17:53.199	0:17:27.170	0:00:26.03	2.42%
24	0:11:56.497	0:11:50.056	0:00:06.44	0.84%
25	0:13:54.834	0:15:54.471	-0:01:59.637	-14.33%
26	0:09:07.838	0:07:45.547	0:01:22.291	15.02%
27	0:04:08.422	0:04:06.256	0:00:02.166	0.87%
28	0:11:42.285	0:08:02.729	0:03:39.556	31.26%
29	0:06:52.139	0:05:39.661	0:01:12.478	17.59%
30	1:14:20.208	1:14:23.560	-0:00:03.352	-0.08%
Total	11:02:46.560	10:21:47.234	0:46:29.936	6.18%

BB – Hive On MR – 1T

QUERY No	QUERY Exec Time W/O NUMA(hh:mm:ss.msec)	QUERY Exec Time With NUMA (hh:mm:ss.msec)	Improvement (hh:mm:ss.msec)	Improvement % (Impr/Exec Time W/O NUMA)*100
1	0:26:57.424	0:25:19.143	0:01:38.281	6.08%
2	5:55:38.790	5:09:31.906	0:46:06.884	12.97%
3	1:29:06.039	1:23:32.765	0:05:33.274	6.23%
4	3:59:20.774	3:05:23.636	0:53:57.138	22.54%
5	5:58:09.045	4:10:19.197	1:47:49.848	30.11%
6	2:38:15.845	2:00:20.909	0:37:54.936	23.96%
7	0:36:51.733	0:32:43.695	0:04:08.038	11.21%
8	3:00:04.111	2:28:41.655	0:31:22.456	17.42%
9	0:44:09.806	0:39:31.227	0:04:38.579	10.51%
10	0:51:52.269	0:51:15.096	0:00:37.173	1.19%
11	0:15:40.548	0:13:33.767	0:02:06.781	13.48%
12	0:30:18.956	0:27:47.187	0:02:31.769	8.34%
13	0:37:54.178	0:36:33.568	0:01:20.610	3.54%
14	0:06:51.893	0:06:37.560	0:00:14.333	3.48%
15	0:11:18.726	0:10:17.290	0:01:01.436	9.05%

BB – Hive On MR – 1T (Cont)

QUERY No	QUERY Exec Time W/O NUMA(hh:mm:ss.msec)	QUERY Exec Time With NUMA (hh:mm:ss.msec)	Improvement (hh:mm:ss.msec)	Improvement % (Impr/Exec Time W/O NUMA)*100
16	2:13:28.889	1:50:55.890	0:22:32.999	16.89%
17	0:15:27.248	0:14:04.685	0:01:22.563	8.90%
18	1:50:51.138	1:41:42.206	0:09:08.932	8.25%
19	0:28:36.825	0:28:07.362	0:00:29.463	1.72%
20	0:39:57.797	0:33:52.184	0:06:05.613	15.25%
21	0:35:47.749	0:35:21.329	0:00:26.420	1.23%
22	0:07:51.338	0:07:53.269	-0:00:01.931	-0.41%
23	0:29:03.146	0:26:00.382	0:03:02.764	10.48%
24	0:20:25.184	0:20:14.264	0:00:10.920	0.89%
25	1:06:26.779	1:01:20.752	0:05:06.027	7.68%
26	0:34:06.367	0:26:21.788	0:07:44.579	22.70%
27	0:04:20.948	0:04:15.219	0:00:05.729	2.20%
28	0:42:58.541	0:47:53.223	-0:04:54.682	-11.43%
29	0:22:52.555	0:20:04.695	0:02:47.860	12.23%
30	5:05:44.011	4:05:49.539	0:59:54.472	19.59%
Total	41.44.28.652	35.25.25.388	6:19:03.264	15.14%

References

- Big Bench - <https://github.com/intel-hadoop/Big-Data-Benchmark-for-Big-Bench>
- numactl - <https://linux.die.net/man/8/numactl>
- numatop - <https://01.org/numatop>