Benchmark Data Sheet

Cassandra Throughput Under Service Level Expectations

azul

Benchmark Summary

We used a modified tlp-stress benchmark (https://github.com/AzulSystems/tlp-stress) to determine the levels of throughput different JDK configurations can handle while still maintaining a service level expectation the actually experienced 99.9%'ile in all 10 second periods during the load test remain below 100msec. Read the blog (https://www.azul.com/blog/cassandra-performance-throughout-responsiveness-capacity-and-cost/) for full details. We performed 275 separate cluster runs (5 JVM configurations, 11 load levels between 20K and 120k ops, 5 runs each) lasting 2.5 hours each, treating the first 30 minutes of each run as a warmup period, and considering the experienced service levels during the 120-minute post-warmup period of each run. In the course of collecting the needed P99.9 Max information, we also collected P99 Max, P90 Max, and P50 Max information under the same set of runs.

JVM configurations tested:

- Azul's Prime 11.0.13 JDK
- OpenJDK 11.0.13 with the following garbage collectors:
  - G1
  - CMS
  - Shenandoah
  - ZGC

Cassandra Cluster Information:

- Cassandra 4.0.1
- 3 node cluster, AWS r5d.2xlarge instances
- 40GB heap
- Replication factor = 3
- Consistency level LOCAL_ONE

Load Generator Information:

- One c5.2xlarge AWS instance
- 8 threads with rate limiting
- 50 concurrent queries per thread
- 80% writes/20% reads
- Asynchronous queries (tlp-stress uses asynchronous queries extensively, which can easily overwhelm Cassandra nodes with a limited number of stress threads)

Benchmark Data Sheet

Cassandra Throughput Under Service Level Expectations

azul

Results

Client Experienced P99.9 Max on Queries Under Various Loads

Attempted Throughput (ops/sec)	Prime (C4)	CMS	G1	Shenandoah	ZGC
20K	16.4	207	414	45.2	53.2
30K	28.4	228	432	47.6	58.7
40K	26.9	230	482	58.1	1041760
50K	36.5	227	572	170	988283
60K	28.1	291	589	225	1405092
70K	44.1	339	635	356	2028995
80K	29.8	989	623	3865	1223688
90K	28.5	171049	1174	410518	1262486
100K	1125	414974	118161	522191	811074
110K	11682	537395	284688	638058	1413480
120K	96338	660603	457179	802161	1535115

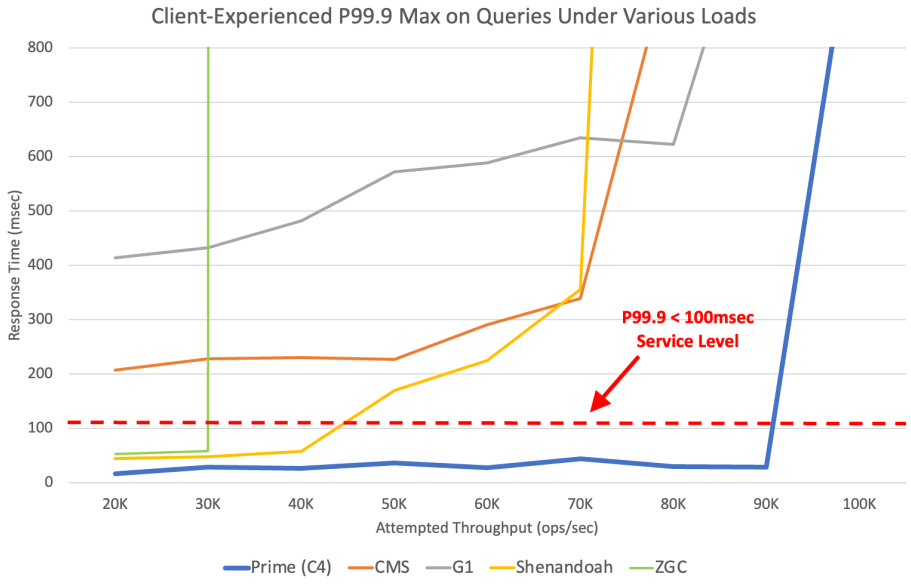

Benchmark Data Sheet

Cassandra Throughput Under Service Level Expectations

azul

Client Experienced P99 Max on Queries Under Various Loads

Attempted Throughput	Prime (C4)	CMS	G1	Shenandoah	ZGC
20K	6.32	153	360	6.92	33.6
30K	6.83	169	388	12.5	35.8
40K	6.99	193	442	23.6	1041760
50K	7.29	200	541	81.4	988283
60K	9.65	264	560	201	1405092
70K	10.3	307	587	293	2027946
80K	9.62	953	605	3811	1223688
90K	14.5	171049	1155	410255	1262486
100K	1092	414974	118161	522191	811074
110K	11682	537395	284688	638058	1413480
120K	96272	660603	457179	801636	1535115

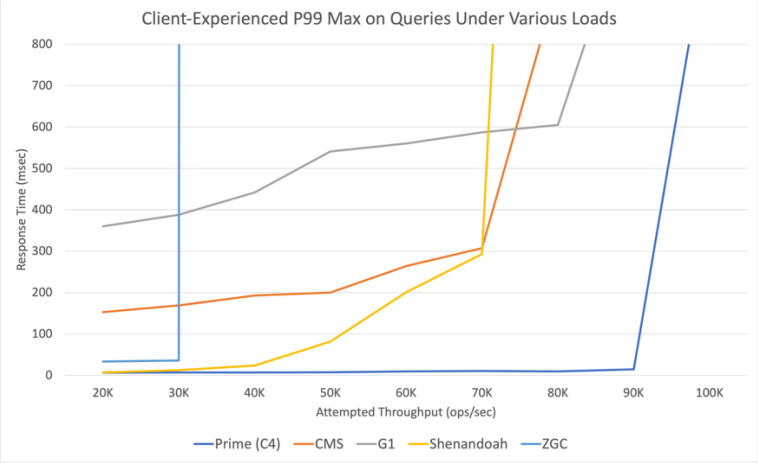

Benchmark Data Sheet

Cassandra Throughput Under Service Level Expectations

azul

Client Experienced P90 Max on Queries Under Various Loads

Attempted Throughput	Prime (C4)	CMS	G1	Shenandoah	ZGC
20K	1.21	24.2	176	1.34	1.64
30K	1.7	78.4	203	2.55	2.75
40K	2.29	90.5	243	2.32	1041236
50K	2.31	128	272	4.31	988283
60K	2.63	169	412	9.33	1405092
70K	2.5	229	471	14.7	2027946
80K	2.46	889	482	3372	1222640
90K	2.62	170656	1074	409993	1261437
100K	1001	414712	117965	522191	810549
110K	11657	536871	284164	637534	1412432
120K	95879	660079	456917	801112	1535115

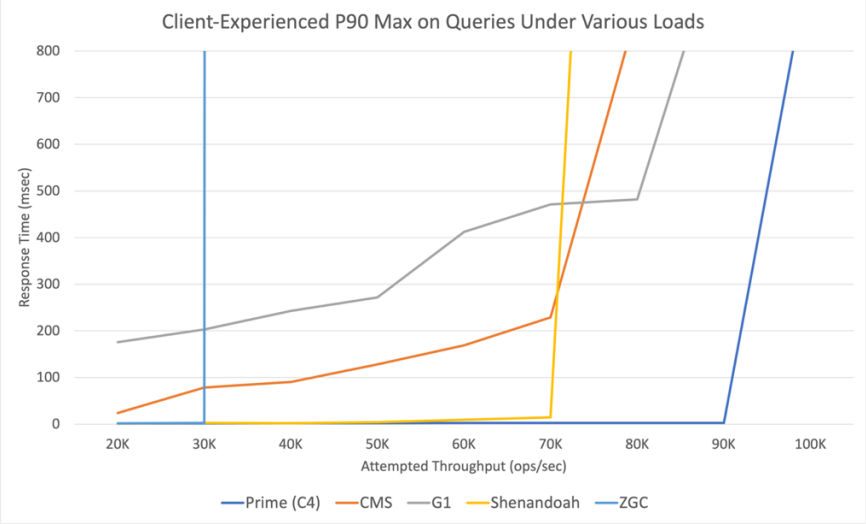

Benchmark Data Sheet

Cassandra Throughput Under Service Level Expectations

azul

Client Experienced P50 Max on Queries Under Various Loads

Attempted Throughput	Prime (C4)	CMS	G1	Shenandoah	ZGC
20K	0.637	0.649	0.632	0.644	0.591
30K	0.692	0.701	0.689	0.681	0.678
40K	0.725	0.736	0.814	0.749	1039663
50K	0.739	11.8	2.77	0.774	985661
60K	0.832	55.9	47.5	0.839	1404043
70K	0.815	112	141	0.892	2026897
80K	0.831	804	155	1656	1221591
90K	0.839	169869	884	401867	1261437
100K	745	414188	116654	519307	809501
110K	4174	535560	281543	635437	1411383
120K	86901	659030	456131	799015	1534067