

We've shown the performance improvement of zheap over heap in a few different pgbench scenarios. All of these tests were run with data that fits in *shared_buffers (32GB)*, and *16 transaction slots per zheap page*. *Scenario-1 and Scenario-2 has used synchronous_commit = off and Scenario-3 and Scenario-4 has used synchronous_commit = on*. We've not included the *undo size* since it gets discarded within seconds after the tests complete. Machine : "cthulhu" 8 node numa machine with 128 hyper threads, RAM 503 GB.

Scenario 1: A 15 minutes simple-update pgbench test with scale factor 100 shows 5.13% TPS improvement with 64 clients. The performance improvement increases as we increase the scale factor; at scale factor 1000, it reaches 11.51% with 64 clients.

		HEAP			ZHEAP (tables)			Improvement	
	Scale Factor	Accounts	History	TPS	Accounts	History	TPS	Size	TPS
Before test	100	1281 MB	0 MB	NA	1149 MB	0 MB	NA	-10.30%	NA
	1000	13 GB	0 MB	NA	11 GB	0 MB	NA	-15.38%	NA
After test	100	1455 MB	2722 MB	58393	1149 MB	1950 MB	61388	-26.47%	5.13%
	1000	13 GB	2155 MB	45723	11 GB	1639 MB	51003	-16%	11.51%

Scenario 2: To show the effect of bloat, we've performed another test similar to the previous scenario, but a transaction is kept open for the first 15 minutes of a 30-minute test. This restricts HOT-pruning for the heap and undo-discarding for zheap for the first half of the test. Scale factor 1000 - 75.86% TPS improvement for zheap at 64 client count. Scale factor 3000 - 98.18% TPS improvement for zheap at 64 client count.

		HEAP			ZHEAP (tables)			Improvement	
	Scale Factor	Accounts	History	TPS	Accounts	History	TPS	Size	TPS
Before test	1000	13 GB	0 MB	NA	11 GB	0 MB	NA	-15.38%	NA
	3000	38 GB	0 MB	NA	34 GB	0 MB	NA	-10.52%	NA
After test	1000	17 GB	2639 MB	29423	11 GB	3263 MB	51743	-26.30%	75.86%
	3000	43 GB	2499 MB	27893	34 GB	3537 MB	55280	-17.70%	98.18%

Scenario 3: A 15 minutes simple-update pgbench test with scale factor 100 shows 6% TPS improvement with 64 clients. The performance improvement increases as we increase the scale factor to 1000 achieving 11.8% with 64 clients.

		HEAP			ZHEAP (tables)			Improvement	
	Scale Factor	Accounts	History	TPS	Accounts	History	TPS	Size	TPS
Before test	100	1281 MB	0 MB	NA	1149 MB	0 MB	NA	-10.30%	NA
	1000	13 GB	0 MB	NA	11 GB	0 MB	NA	-15.38%	NA
After test	100	1388 MB	1566 MB	33822	1149 MB	1111 MB	35851	-23.61%	6.00%
	1000	13 GB	926 MB	20623	11 GB	744 MB	23069	-15.80%	11.80%

Scenario 4: To amplify the effect of bloats in scenario 3, we've performed another test similar to scenario, but a transaction is kept open for the first 15 minutes of a 30 minute test. This restricts HOT-pruning for heap and undo-discarding for zheap for the first half of the test.

		HEAP			ZHEAP (tables)			Improvement	
--	--	------	--	--	----------------	--	--	-------------	--

	Scale Factor	Accounts	History	TPS	Accounts	History	TPS	Size	TPS
Before test	1000	13 GB	0 MB	NA	11 GB	0 MB	NA	-15.38%	NA
	3000	38 GB	0 MB	NA	34 GB	0 MB	NA	-10.52%	NA
After test	1000	14 GB	1554 MB	17422	11 GB	1439 MB	23213	-20%	33%
	3000	39 GB	1259 MB	13693	34 GB	1169 MB	17428	-12.90%	27.27%