

	<u># regions</u>	<u># rows</u>	<u>0.18 Branch</u>			<u>0.19 Trunk</u>				<u>Details</u>
			<u>count</u>	<u>msec/row</u>	<u>scan</u>	<u>count</u>	<u>msec/row</u>	<u>% diff</u>	<u>scan</u>	
table1	6	375388	171.657	0.457	22.454	80.803	0.215	52.93%	21.971	Large data size / row
table2	1	343131	94.017	0.274	1.756	29.812	0.087	68.29%	7.511	Small data size / row
table3	1	8146	17.51	2.150	5.392	12.11	1.487	30.84%	0.953	Many cols / row
table4	1	197	0.137	0.695	0.021	0.132	0.670	3.65%	0.012	Few rows

Count is a count 'table' from the shell

Scan is a scan 'table', ['existingfamily:nonexistentcolumn'] from shell

All numbers are averages of three runs. There was much more variability with 0.19 Trunk