


[1]																												
1/30 field	Total job time	ExecutorRuntime	BytesRead		# of fields	without pruning	with pruning																					
w/o pruning #1	32	1513390	11939920307		1	34.8	14.2																					
w/o pruning #2	32	1481642			5	18.4	13.2																					
w/o pruning #3	50	2351618			10	25.6	14.8																					
w/o pruning #4	33	1867898			15	21.2	17.6																					
w/o pruning #5	27	1525897			20	34.6	24.6																					
w/o pruning avg	34.8	1748089			30	22.8	22.6																					
w/ pruning #1	15	798923	1230186185	<div><h3>Job time: with pruning and without pruning</h3><table><thead><tr><th># of fields</th><th>without pruning</th><th>with pruning</th></tr></thead><tbody><tr><td>1</td><td>34.8</td><td>14.2</td></tr><tr><td>5</td><td>18.4</td><td>13.2</td></tr><tr><td>10</td><td>25.6</td><td>14.8</td></tr><tr><td>15</td><td>21.2</td><td>17.6</td></tr><tr><td>20</td><td>34.6</td><td>24.6</td></tr><tr><td>30</td><td>22.8</td><td>22.6</td></tr></tbody></table></div>	# of fields	without pruning	with pruning	1	34.8	14.2	5	18.4	13.2	10	25.6	14.8	15	21.2	17.6	20	34.6	24.6	30	22.8	22.6			
# of fields	without pruning	with pruning																										
1	34.8	14.2																										
5	18.4	13.2																										
10	25.6	14.8																										
15	21.2	17.6																										
20	34.6	24.6																										
30	22.8	22.6																										
w/ pruning #2	17	945437																										
w/ pruning #3	12	555534																										
w/ pruning #4	13	461078																										
w/ pruning #5	14	815851																										
w/ pruning avg	14.2	715364.6																										
w/ pruning #2																												
5/30 field (one array)	Total job time	ExecutorRuntime	BytesRead																									
w/o pruning #1	19	1312890	11939984321																									
w/o pruning #2	20	1070268																										
w/o pruning #3	17	1060790																										
w/o pruning #4	16	921842																										
w/o pruning #5	20	1174904																										
w/o pruning avg	18.4	1108138.8																										
w/ pruning #1	15	947763																										
w/ pruning #2	15	903240																										
w/ pruning #3	8	378407																										
w/ pruning #4	14	604876																										
w/ pruning #5	14	552969																										
w/ pruning avg	13.2	677451																										
10/30 field (one array)	Total job time	ExecutorRuntime	BytesRead																									
w/o pruning #1	20	1067536	11940053129																									
w/o pruning #2	27	980703																										
w/o pruning #3	18	991669																										
w/o pruning #4	35	1854641																										
w/o pruning #5	28	1400340																										
w/o pruning avg	25.6	1258977.8																										
w/ pruning #1	13	529693	4299961615																									
w/ pruning #2	15	638901																										

w/ pruning #3	11	601491							
w/ pruning #4	21	1150256							
w/ pruning #5	14	878139							
<b>w/ pruning avg</b>	<b>14.8</b>	<b>759696</b>							
<b>15/30 field (two arrays)</b>	<b>Total job time</b>	<b>ExecutorRuntime</b>	<b>BytesRead</b>	<b># of fields</b>	<b>without pruning</b>	<b>with pruning</b>			
w/o pruning #1	17	1019042	11940132089	1	<b>1748089</b>	715364			
w/o pruning #2	24	1193403		5	1108139	677451			
w/o pruning #3	21	1323726		10	1258978	759696			
w/o pruning #4	20	1293579		15	1226621	975940			
w/o pruning #5	24	1303359		20	1881825	1329440			
<b>w/o pruning avg</b>	<b>21.2</b>	<b>1226621.8</b>		30	1105217	1204484			
w/ pruning #1	20	1189638							
w/ pruning #2	14	780843							
w/ pruning #3	23	1300785							
w/ pruning #4	13	659361							
w/ pruning #5	18	949076							
<b>w/ pruning avg</b>	<b>17.6</b>	<b>975940.6</b>							
<b>20/30 field (three arrays)</b>	<b>Total job time</b>	<b>ExecutorRuntime</b>	<b>BytesRead</b>						
w/o pruning #1	35	2073196	11940204140						
w/o pruning #2	51	1982131							
w/o pruning #3	34	2382798							
w/o pruning #4	31	1719348							
w/o pruning #5	22	1251653							
<b>w/o pruning avg</b>	<b>34.6</b>	<b>1881825.2</b>							
w/ pruning #1	27	1865219	6998253994						
w/ pruning #2	28	1476422							
w/ pruning #3	23	1111034							
w/ pruning #4	38	1251785							
w/ pruning #5	16	942742							
<b>w/ pruning avg</b>	<b>26.4</b>	<b>1329440.4</b>							
<b>30/30 field</b>	<b>Total job time</b>	<b>ExecutorRuntime</b>	<b>BytesRead</b>						
w/o pruning #1	22	1197862	21958691278						
w/o pruning #2	23	1323532							
w/o pruning #3	20	1026105							
w/o pruning #4	24	851176							
w/o pruning #5	25	1127414							



<b>w/o pruning avg</b>	<b>22.8</b>	<b>1105217.8</b>							
w/ pruning #1	26	1567388							
w/ pruning #2	24	1154639							
w/ pruning #3	20	1085682							
w/ pruning #4	20	1098309							
w/ pruning #5	23	1116402							
<b>w/ pruning avg</b>	<b>22.6</b>	<b>1204484</b>							

[1] Test query:

```
select s.f1, s.f2, ..., s.fn from tbl.
```

where s is a struct of 30 fields of schema:

```
struct<
  string, string, string, bigint, array<string>,
  string, string, string, boolean, string,
  string, string, string, bigint, string,
  string, array<string>, long, string, array<string>,
  string, array<string>, string, long, bigint,
  string, string, bigint, string, bigint
>
```

and f1, f2, ..., fn are fields in the struct.

The tests select 1, 5, 10, 15, 20, 30 fields. Each test run contains 5 iterations and average is taken in the end.

HoS cluster setting:

- locality: any 123, rack local 18
- no dynamic allocation
- # of executors: 30
- exec mem: 5154mb
- exec mem overhead: 1230mb
- driver mem: 2148mb
- driver mem overhead: 400mb

All queries are executed within a single session.