

RDBMS comparison, read-only benchmark.

Goal: Compare MySQL, Oracle, PostgreSQL, EnterpriseDB in readonly CPU bound benchmarks.

Short conclusions:

1.
MySQL 5.0 / InnoDB shows the best result,
Oracle 10g is slower by 40.7%
PostgreSQL 8 is slower by 23.47%
MySQL 5.0 / MyISAM is slower by 43.18%
EnterpriseDB 8 is slower by 46.02%
2.
Oracle 10g shows the best CPU scalability with factor 3.67 (in comparison 4CPU / HT with 1 CPU),
EnterpriseDB 8 with factor 3.64
MySQL 5.0 / InnoDB – 3.44
PostgreSQL 8 – 3.26
MySQL 5.0 / MyISAM – 2.47
3.
MySQL 5.0-standard is faster than MySQL 4.1-standard on 9.25%
4.
MySQL 5.0 / MyISAM shows the bad result unexpectedly. However, the results with MySQL / MyISAM are very stable even with many threads, in contrast to other RDBMS.

Benchmark description:

System specification:

4x Intel Xeon 2GHz with HT CPU
Hardware RAID 10 with 8 disks
4GB of RAM
RedHat 3 AS, Update 2, kernel 2.4.21-EL.

Benchmark: sysbench (sysbench.sf.net), OLTP readonly benchmark.
Tablesize 1000000 rows (database size ~300MB). Special & Uniform distribution.

Table specification:

```
CREATE TABLE `sbtest` (  
  `id` int(10) unsigned NOT NULL auto_increment,  
  `k` int(10) unsigned NOT NULL default '0',  
  `c` char(120) NOT NULL default '',  
  `pad` char(60) NOT NULL default '',  
  PRIMARY KEY (`id`),  
  KEY `k` (`k`);
```

Each transaction contains:

Point queries (10 times):

```
SELECT c FROM sbtest WHERE id=N
```

Range queries (1 time):

```
SELECT c FROM sbtest WHERE id BETWEEN N AND M
```

Range SUM() queries (1 time):

```
SELECT SUM(c) FROM sbtest WHERE id BETWEEN N and M
```

Range ORDER BY queries (1 time):

```
SELECT c FROM sbtest WHERE id between N and M ORDER BY c
```

Range DISTINCT queries (1 time):

```
SELECT DISTINCT c FROM sbtest WHERE id BETWEEN N and M ORDER BY c
```

Special distribution means that in 75% of cases sysbench generates N & M values from the 1% diapason (in our case with 1000000 rows – in 75% of queries we use rows with ID from 495000 to 505000).

In uniform distribution sysbench uses all rows from table uniformly.

Note 1: This is synthetic OLTP read-only benchmark, accesses done by primary key (this is probably reason why MyISAM is so slow). The results are for given workload and should not be used to draw conclusions about other workloads. Primary key lookups are normally fastest in databases, so this test is to illustrate "peak performance" database can deliver. Also full table is fitted in cache memory, so no I/O activity during benchmark, results are limited only by CPU.

Note 2: For MySQL 4.1 and MySQL 5.0 , InnoDB I created tables separately, so for MySQL 4.1 "old" redundant InnoDB row format was used, for 5.0 – "new" compact.

Note 3: For all RDBMS there was used Prepare Statement API.

Note 4: In all benchmarks (excluding MySQL 4.1 / 5.0 comparison) I used MySQL 5.0.7 custom build.


The results are shown in transaction per second (more is better).

Results:

1. MySQL 4.1 / 5.0 comparison (InnoDB tables):

I used

- mysql-standard-4.1.12-pc-linux-gnu-i686
- mysql-standard-4.1.12-pc-linux-gnu-i686-icc-glibc23
- mysql-standard-5.0.7-beta-linux-i686
- mysql 5.0.7 custom build (configure params see in Appendix)



tps	mysql-standard-4.1.12-pc-linux-gnu-i686		mysql-standard-4.1.12-pc-linux-gnu-i686-icc-glibc23		mysql-standard-5.0.7-beta-linux-i686		MySQL 5.0.7, InnoDB	
threads	special	uniform	special	uniform	special	uniform	special	uniform
1	230.54	228.03	210.53	209.85	230.54	210.53	210.53	230.54
2	396.17	377.80	375.15	374.80	396.17	375.15	375.15	396.17
4	680.83	676.35	651.34	649.19	680.83	651.34	651.34	680.83
8	918.96	900.83	893.05	889.39	918.96	893.05	893.05	918.96
16	832.34	825.45	834.07	826.95	832.34	834.07	834.07	834.07
24	808.91	806.01	817.44	812.38	808.91	817.44	817.44	817.44
32	821.34	817.45	817.08	805.92	821.34	817.08	817.08	821.34
40	810.86	805.57	830.28	818.79	810.86	830.28	830.28	830.28
64	811.45	800.27	836.85	828.89	811.45	836.85	836.85	836.85
88	796.20	791.56	841.24	838.02	796.20	841.24	841.24	841.24
128	779.17	772.84	841.29	842.48	779.17	842.48	842.48	842.48
196	640.33	618.11	786.46	796.23	640.33	796.23	796.23	796.23
256	535.93	546.32	737.49	751.54	546.32	751.54	751.54	751.54

2. RDBMS comparison

- 4 CPU with HyperThreading

MySQL 5.0.7 custom build, InnoDB tables
MySQL 5.0.7 custom build, MyISAM tables
Oracle 10g
PostgreSQL 8
PostgreSQL 8
EnterpriseDB 8


(The databases startup params see in Appendix)

Note 1. Oracle was tested only with 1 – 196 threads.

Note 2. For case with 4 HT CPU I tested EnterpriseDB only with 1-88 threads.

(The empty results don't mind that benchmark was failed)


RDBMS comparison, read-only special



tps	MySQL 5.0.7, InnoDB			Oracle 10g			PostgreSQL 8			EnterpriseDB 8			MySQL 5.0.7, MySAM		
threads	special	uniform		special	uniform		special	uniform		special	uniform		special	uniform	
1	305.31	302.25		192.55	189.15		248.59	245.76		159.39	157.82		244.13	239.57	
2	501.90	494.01		299.88	303.14		423.30	413.82		273.82	274.81		384.36	379.75	
4	827.42	817.15		490.48	488.00		656.63	642.42		433.11	428.41		542.92	537.99	
8	1048.55	1047.36		606.23	605.26		802.47	776.42		566.04	553.84		595.74	590.47	
16	944.04	949.98		621.63	614.58		734.58	704.12		528.13	511.13		556.22	551.75	
24	946.88	942.02		602.78	595.13		700.36	667.63		491.47	475.01		544.44	544.52	
32	937.68	934.45		575.88	561.66		635.99	606.05		484.96	466.55		545.26	540.02	
40	936.12	925.40		564.64	556.88		623.52	589.07		477.63	457.73		542.97	536.90	
64	957.19	942.29		573.26	562.09		586.23	552.21		465.04	435.94		548.29	541.66	
88	942.47	931.91		560.70	545.74		561.06	513.41		463.49	441.26		548.26	543.70	

- 2 CPU with HT


RDBMS comparison, read-only special



tps	MySQL 5.0.7, InnoDB			Oracle 10g			PostgreSQL 8			EnterpriseDB 8			MySQL 5.0.7, MySAM		
threads	special	uniform		special	uniform		special	uniform		special	uniform		special	uniform	
1	302.62	299.88		186.87	183.89		251.97	247.44		157.29	155.02		241.48	238.65	
2	480.21	474.36		263.13	263.13		380.19	361.64		231.67	225.30		373.59	368.81	
4	661.34	655.71		374.36	373.91		511.56	502.51		312.48	309.25		481.10	474.37	
8	653.10	647.31		374.66	368.02		505.77	494.77		310.26	305.92		472.52	468.37	
16	602.92	597.00		370.42	367.05		465.02	449.86		297.62	291.89		447.74	445.56	
24	615.91	603.49		364.00	356.23		455.29	438.26		298.65	287.90		438.92	435.12	
32	620.06	615.34		360.54	355.25		442.27	429.60		297.49	290.03		443.88	439.37	
40	630.16	621.07		361.63	350.40		463.33	440.73		297.01	288.63		441.63	441.22	
64	606.73	603.76		360.97	349.32		457.75	436.06		291.97	281.10		441.10	437.86	
88	599.25	595.11		352.00	343.07		443.02	417.21		289.42	278.07		438.98	436.97	

- **1 CPU with HT**


RDBMS comparison, read-only special



tps	MySQL 5.0.7, InnoDB		Oracle 10g		PostgreSQL 8		EnterpriseDB 8		MySQL 5.0.7, MyISAM	
threads	special	uniform	special	uniform	special	uniform	special	uniform	special	uniform
1	299.56	285.64	170.29	168.67	243.63	238.67	154.38	152.62	235.19	234.37
2	347.49	344.56	192.98	192.15	251.65	248.44	152.08	154.46	283.00	280.78
4	362.12	358.35	193.31	191.28	274.42	268.96	162.40	160.37	294.65	291.82
8	354.10	346.48	193.94	190.60	266.99	260.99	158.84	156.33	280.35	278.83
16	345.17	340.27	192.02	187.89	267.02	260.43	159.07	155.44	284.85	283.95
24	346.27	341.56	191.90	187.17	266.92	256.06	158.46	153.34	284.30	278.51
32	348.04	342.87	191.86	186.73	259.42	249.12	155.64	152.07	275.64	273.04
40	342.91	337.69	191.00	188.93	255.38	254.77	154.68	150.65	273.04	281.22
64	335.27	333.65	187.23	185.30	260.07	232.34	153.68	145.59	271.87	268.46
88	330.59	326.97	188.96	180.43	237.54	227.56	147.18	142.89	268.39	269.68

3. 1 CPU without HT

RDBMS comparison, read-only special




tps	MySQL 5.0.7, InnoDB		Oracle 10g		PostgreSQL 8		EnterpriseDB 8		MySQL 5.0.7, MyISAM	
	special	uniform	special	uniform	special	uniform	special	uniform	special	uniform
1	273.30	271.24	169.18	167.54	245.58	241.15	155.46	153.54	225.11	217.82
2	298.06	297.78	156.72	154.89	246.44	241.69	155.23	153.08	240.85	237.82
4	302.67	299.78	161.22	159.51	245.66	239.71	154.77	152.14	241.10	238.99
8	304.62	299.55	163.30	160.59	243.06	236.37	153.89	150.66	235.82	230.26
16	294.48	295.04	163.77	159.60	242.16	233.50	152.51	147.97	236.45	234.05
24	298.63	295.45	163.74	159.71	239.88	227.90	151.73	147.69	237.35	234.29
32	297.10	293.04	164.83	162.64	239.31	226.64	150.85	145.56	237.11	233.01
40	290.63	291.31	163.05	156.91	237.27	226.08	147.58	143.79	236.75	235.31
64	291.46	292.06	151.46	158.54	227.64	222.07	145.87	143.52	234.28	232.60
88	287.78	284.24	155.86	151.05	220.57	190.86	140.50	129.04	225.42	224.85

CPU scalability:

Peak results:

	MySQL 5.0.7, InnoDB	Oracle 10g	PostgreSQL 8	EnterpriseDB 8	MySQL 5.0.7, MyISAM
1CPU	304.62	169.18	246.44	155.46	241.1
1CPU HT	362.12	193.94	274.42	162.4	294.65
2CPU HT	661.34	374.66	511.56	312.48	481.1
4CPU HT	1048.55	621.63	802.47	566.04	595.74

Peak results



Scalability:

peak result with N CPU / peak result with N/2 CPU

	MySQL 5.0.7, InnoDB	Oracle 10g	PostgreSQL 8	EnterpriseDB 8	MySQL 5.0.7, MyISAM
1CPU	X	X	X	X	X
1CPU HT	1.19	1.15	1.11	1.04	1.22
2CPU HT	1.83	1.93	1.86	1.92	1.63
4CPU HT	1.59	1.66	1.57	1.81	1.24

peak result with N CPU / peak result with 1 CPU

	MySQL 5.0.7, InnoDB	Oracle 10g	PostgreSQL 8	EnterpriseDB 8	MySQL 5.0.7, MyISAM
1CPU	X	X	X	X	X
1CPU HT	1.19	1.15	1.11	1.04	1.22
2CPU HT	2.17	2.21	2.08	2.01	2
4CPU HT	3.44	3.67	3.26	3.64	2.47

Threads scalability:

I'm using data for 2 CPU with HT, 256 threads.

tps		MySQL 5.0.7, InnoDB		Oracle 10g		PostgreSQL 8		EnterpriseDB 8		MySQL 5.0.7, MyISAM	
threads		special	uniform	special	uniform	special	uniform	special	uniform	special	uniform
peak		661.34	655.71	374.66	373.91	511.56	502.51	312.48	309.25	481.10	474.37
256		548.84	544.08	308.07	290.80	322.53	305.35	233.93	221.84	433.74	434.45
sag, %		17.01	17.02	17.77	22.23	36.95	39.24	25.14	28.27	9.84	8.42

Appendix

Startup params.

Oracle 10g:

```
*.background_dump_dest='/home/oracle/product/10.1.0/db_1/admin/orcl/bdump'
*.compatible='10.1.0.2.0'
*.control_files='/home/oracle/product/10.1.0/oradata/orcl/control01.ctl','/home/oracle/product/10
.1.0/oradata/orcl/control02.ctl','/
home/oracle/product/10.1.0/oradata/orcl/control03.ctl'
*.core_dump_dest='/home/oracle/product/10.1.0/db_1/admin/orcl/cdump'
*.db_block_checking=FALSE
*.db_block_checksum=FALSE
*.db_block_size=8192
*.db_cache_size=400M
*.db_domain=''
*.db_file_multiblock_read_count=16
*.db_name='orcl'
*.db_recovery_file_dest='/home/oracle/product/10.1.0/db_1/flash_recovery_area'
*.db_recovery_file_dest_size=2147483648
*.dispatchers=''
*.java_pool_size=50M
*.job_queue_processes=10
*.large_pool_size=50M
*.open_cursors=2000
*.pga_aggregate_target=1048576000
*.processes=1100
*.remote_login_passwordfile='EXCLUSIVE'
*.session_cached_cursors=100
*.sessions=1500
*.sga_max_size=1G
*.sga_target=1G
*.shared_pool_size=209715200
*.sort_area_size=65536
*.undo_management='AUTO'
*.undo_tablespace='UNDOTBS1'
*.user_dump_dest='/home/oracle/product/10.1.0/db_1/admin/orcl/udump'
*.workarea_size_policy='AUTO'
```

PostgreSQL 8 / EnterpriseDB 8:

```
max_connections = 300
shared_buffers = 50000          # min 16, at least max_connections*2, 8KB each
effective_cache_size = 48000    # typically 8KB each
lc_messages = 'en_US'          # locale for system error message strings
lc_monetary = 'en_US'          # locale for monetary formatting
lc_numeric = 'en_US'           # locale for number formatting
lc_time = 'en_US'              # locale for time formatting
```

MySQL InnoDB:

```
--port=3306 \
--socket=/tmp/mysql.sock \
--user=root \
--log-error=/benchmarks/work/err4.log \
--datadir=$FSPATH \
--basedir=$BASEPATH \
--max_connections=1600 \
--max_connect_errors=10 \
--table_cache=2048 \
--max_allowed_packet=1048576 \
--binlog_cache_size=1048576 \
--max_heap_table_size=67108864 \
--sort_buffer_size=65536 \
--join_buffer_size=1048576 \
--thread_cache=16 \
--thread_concurrency=16 \
--thread_stack=64K \
--query_cache_size=0 \
--ft_min_word_len=4 \
--default_table_type=MYISAM \
--transaction_isolation=REPEATABLE-READ \
```

```

--tmp_table_size=64M \
--skip-locking \
--server-id=1 \
--innodb_status_file=0 \
--innodb_data_home_dir=$FSPATH \
--innodb_data_file_path=ibdata1:100M:autoextend \
--innodb_log_group_home_dir=$FSPATH \
--innodb_buffer_pool_size=1800M \
--innodb_additional_mem_pool_size=20M \
--innodb_log_file_size=900M \
--innodb_log_files_in_group=2 \
--innodb_log_buffer_size=8M \
--innodb_flush_log_at_trx_commit=2 \
--innodb_lock_wait_timeout=300 \
--innodb_locks_unsafe_for_binlog=1 \
--innodb_thread_concurrency=1000

```

MySQL MyISAM:

```

--port=3306 \
--socket=/tmp/mysql.sock \
--user=root \
--log-error=/benchmarks/work/err \
--datadir=$FSPATH \
--basedir=$BASEPATH \
--max_connections=1600 \
--max_connect_errors=10 \
--table_cache=2048 \
--max_allowed_packet=1048576 \
--binlog_cache_size=1048576 \
--max_heap_table_size=67108864 \
--sort_buffer_size=512K \
--join_buffer_size=1048576 \
--key_buffer=1800M \
--thread_cache=16 \
--thread_concurrency=16 \
--thread_stack=64K \
--query_cache_size=0 \
--ft_min_word_len=4 \
--default_table_type=MYISAM \
--transaction_isolation=REPEATABLE-READ \
--tmp_table_size=64M \
--skip-locking \
--server-id=1 \
--skip-innodb

```

MySQL custom build configure line:

```

CFLAGS="-g -O3 -fno-omit-frame-pointer" CXX="gcc" CXXFLAGS="-g -felide-constructors
-fno-exceptions -fno-rtti -O3 -fno-omit-frame-pointer" CXXLDFLAGS="" \
./configure --prefix=/usr/local/mysql \
    --enable-asm \
    --with-extra-charsets=complex \
    --enable-thread-safe-client \
    --with-readline \
    --with-big-tables \
    --enable-local-infile

```