


Sysv vs mmap

PostgreSQL 9.3 Pgbench Transactions Per Second, 2x Xeon X5650 / 24GB (24 threads), Unix socket												
Clients	1	3	6	9	12	15	18	24	32	48	64	80
DragonFly 3.1 sysv	6752	21345	31234	39761	43654	44498	47779	53583	67211	67415	58364	52799
DragonFly 3.1 mmap	8240	20135	27396	33678	37809	37261	38585	43321	50829	19895	9912	6308
Scientific Linux 6.2 sysv	9275	23827	46349	64578	87794	107274	127029	167956	149705	143835	138290	135267
Scientific Linux 6.2 mmap	9518	24822	47921	65354	88592	107255	128261	167080	149119	143467	138124	135827

PostgreSQL 9.3 - System V shared memory vs mmap



Results

Detailed TPS results excluding connections establishing

Clients (and threads)	1	3	6	9	12	15	18	24	32	48	64	80
DragonFly 3.1 sysv_shm	6,866	21,422	30,665	39,374	43,603	44,420	47,865	53,840	68,382	68,665	57,751	52,190
DragonFly 3.1 sysv_shm	7,995	21,329	29,597	40,538	44,137	44,704	47,947	53,536	76,065	67,516	58,441	53,188
DragonFly 3.1 sysv_shm	5,396	21,283	33,442	39,371	43,223	44,369	47,526	53,374	57,186	66,062	58,901	53,019
DragonFly 3.1 mmap	8,167	20,015	28,133	35,183	37,375	36,557	38,316	42,514	51,566	20,459	9,648	5,780
DragonFly 3.1 mmap	8,177	19,400	27,581	33,361	38,264	36,912	38,449	43,048	51,700	20,756	10,790	7,292
DragonFly 3.1 mmap	8,377	20,990	26,473	32,490	37,789	38,313	38,992	44,400	49,222	18,470	9,298	5,851
Scientific Linux 6.2 sysv_shm	9,979	23,787	45,091	65,151	88,020	106,614	128,796	168,111	149,654	143,889	138,292	134,024
Scientific Linux 6.2 sysv_shm	9,510	23,586	46,439	63,598	88,684	107,172	125,998	167,955	149,710	143,778	138,961	135,882
Scientific Linux 6.2 sysv_shm	8,336	24,109	47,516	64,986	86,678	108,037	126,293	167,801	149,752	143,840	137,615	135,895
Scientific Linux 6.2 mmap	9,883	23,825	49,273	65,392	88,944	106,995	128,606	167,035	149,162	143,405	138,247	135,839
Scientific Linux 6.2 mmap	9,422	26,043	47,051	65,398	89,032	107,528	127,081	167,194	149,207	143,690	137,909	136,024
Scientific Linux 6.2 mmap	9,248	24,597	47,438	65,273	87,800	107,241	129,094	167,011	148,989	143,305	138,218	135,617

DragonFly with mmap performance degradation :

1. as the number of clients increase, so does number of IPI/s
2. After 32 clients, the system starts swapping

Setup

Hardware:

- 2x Xeon X5650 (24 threads total)
- 24 GB RAM

Software

PostgreSQL 9.3-devel, before and after the patch changing setup of the shared memory segment from SYSV shared memory to mmap
Scientific Linux 6.2 or DragonFly 3.1 as of 2012-09-12

Goal:

Test scalability with pgbench, see if the use of mmap() doesn't cause performance problems

postgresql.conf :

```
max_connections = 100
update_process_title = off
autovacuum = off
shared_buffers = 6GB
effective_cache_size = 12GB
```

Initialize database cluster :

```
/usr/local/postgres-9.3.sysv/bin/initdb -D /usr/local/pgdata.93
```

Run Postgres :

```
/usr/local/postgres-9.3.sysv/bin/postgres -D /usr/local/pgdata.93
/usr/local/postgres-9.3.mmap/bin/postgres -D /usr/local/pgdata.93
```

Create test database

```
psql template1
create database bench;
pgbench -i -s 800 bench
```

Scaling factor 800 => ~ = 11GB database

Running tests :

```
tar /usr/local/pgdata > /dev/null # Preload disk cache
pgbench -h 127.0.0.1 -j ${THREADS} -c ${CLIENTS} -T 1800 -S bench # Dummy run to warm up caches for 30 mn
```

```
#!/bin/sh
```

```
for clients in 1 3 6 9 12 15 18 24 32 48 64 80
```

```
do
```

```
  THREADS=${clients}
```

```
  ./pgbench -h 127.0.0.1 -j ${THREADS} -c ${clients} -T 600 -S bench > result_${clients}.txt
```

```
done
```

Repeat three times and average the individual results