



PERCONA
Server for MySQL



40 x Database Response

Sinan Alyuruk
Sr. Database Administrator
[gittigidiyor.com /ebay](http://gittigidiyor.com/ebay)

Agenda



Query Response

Slow queries & Frequent



InnoDB & I/O Tuning

Storage migrations (SSD ->
SSD +MVMem)

Buffer pool

Memory vs CPU tuning



Transactions & Locks

Basket operations



ProxySQL

Multiplexing

Load Balancing

Query routing



Full Table Scan Test: MySQL 5.6 ve Oracle 11g



```
select count(*), ktgr from product
where add_date > DATE('2019-09-01')
and end_date > now()
group by ktgr
order by count(*) desc;
```

Query_time: ~ 120 sec (2 min)




```
select /*+ PARALLEL(16) */count(*), ktgr from
product
where add_date > to_date ('2019-01-01','YYYY-MM-
DD') end_date > CURRENT_DATE
group by ktgr
order by count(*) desc
```

Query_time: 18 sec

Average Active Sessions

Foreground Only Foreground + Background



Performance critical!

SQL Tuning is effective..

```
mysql > create index  
super_multi_idx on urun ( add_date,  
end_date, relist_num) ;
```

Performance:

Query_time: ~ 0.1 sec



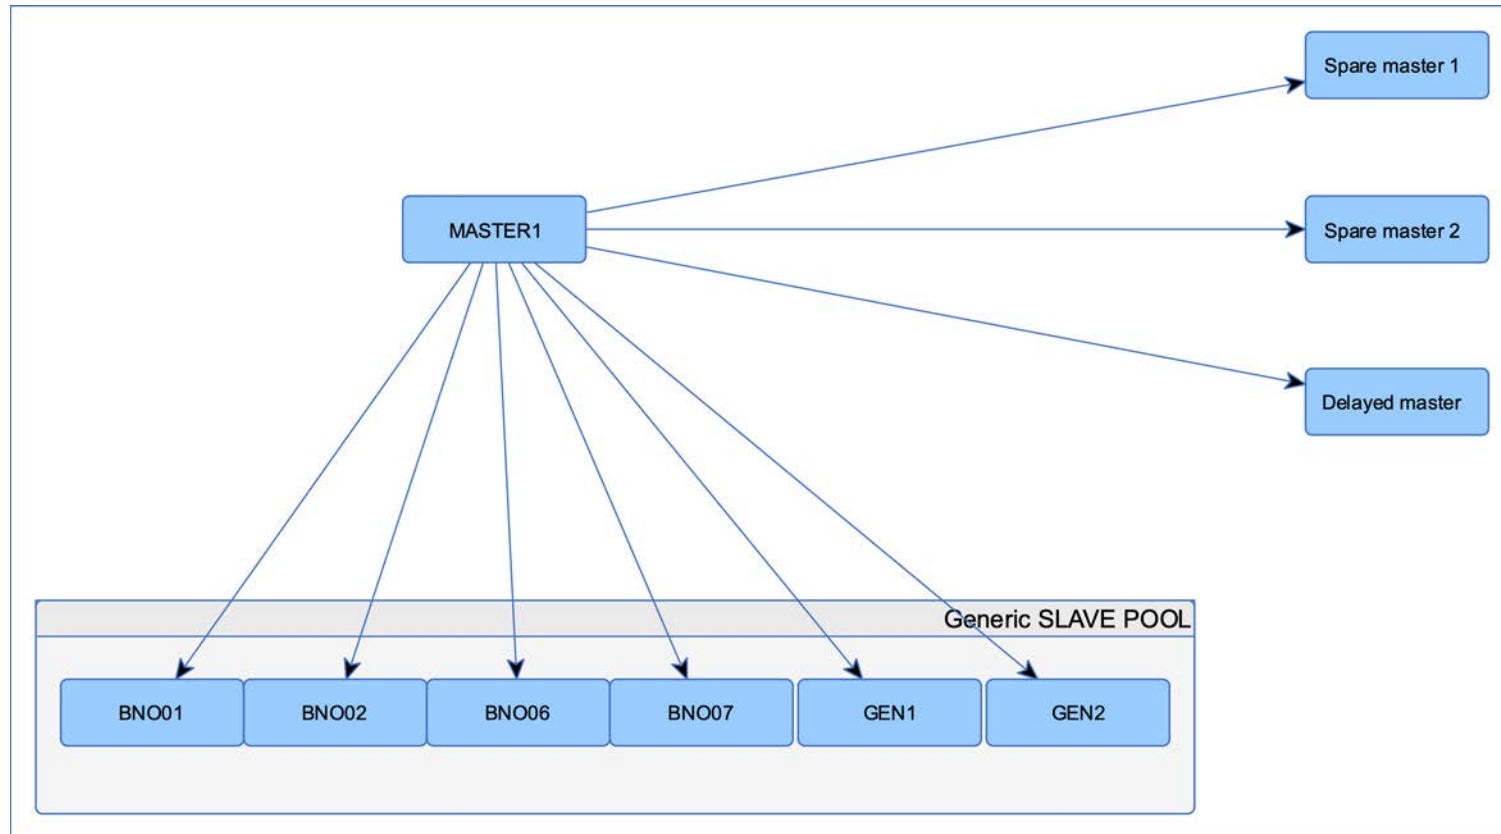
<https://use-the-index-luke.com/>

What About?

Index size > table size?

MySQL Replication Hosts

Indexes on
Generic Slaves



Proxysql Query Digest:

```
select digest, digest_text, count_star, sum_time from stats_mysql_query_digest  
order by count_star desc limit 6;
```

digest	digest_text	count_star	sum_time
0xB0674A190B03C5EF	SHOW WARNINGS	4058932924	0
0xC51BDBBD6C9BB1D7	SET autocommit=?	3662312142	0
0xC51BDBBD6C9BB1D7	SET autocommit=?	2296354713	0
0xDCC3D6FFA2362D05	select * from promotions	1782754685	579105657208
0x69A3B65862823B76	SELECT USER()	1151605897	343313996682
0xDB3A841EF5443C35	commit	1144095480	2811179005

Cache elsewhere

```
query_cache_size =0
```

Please do not use your favorite OLTP database for caching.

MySQL InnoDB ACID Compliance



A is for **Atomicity**: The database must execute the whole transaction and not just part of it – all or nothing.

C is for **Consistency**: The database must enforce the system rules, for example, you can't put a value in a field, if that value is not allowed.

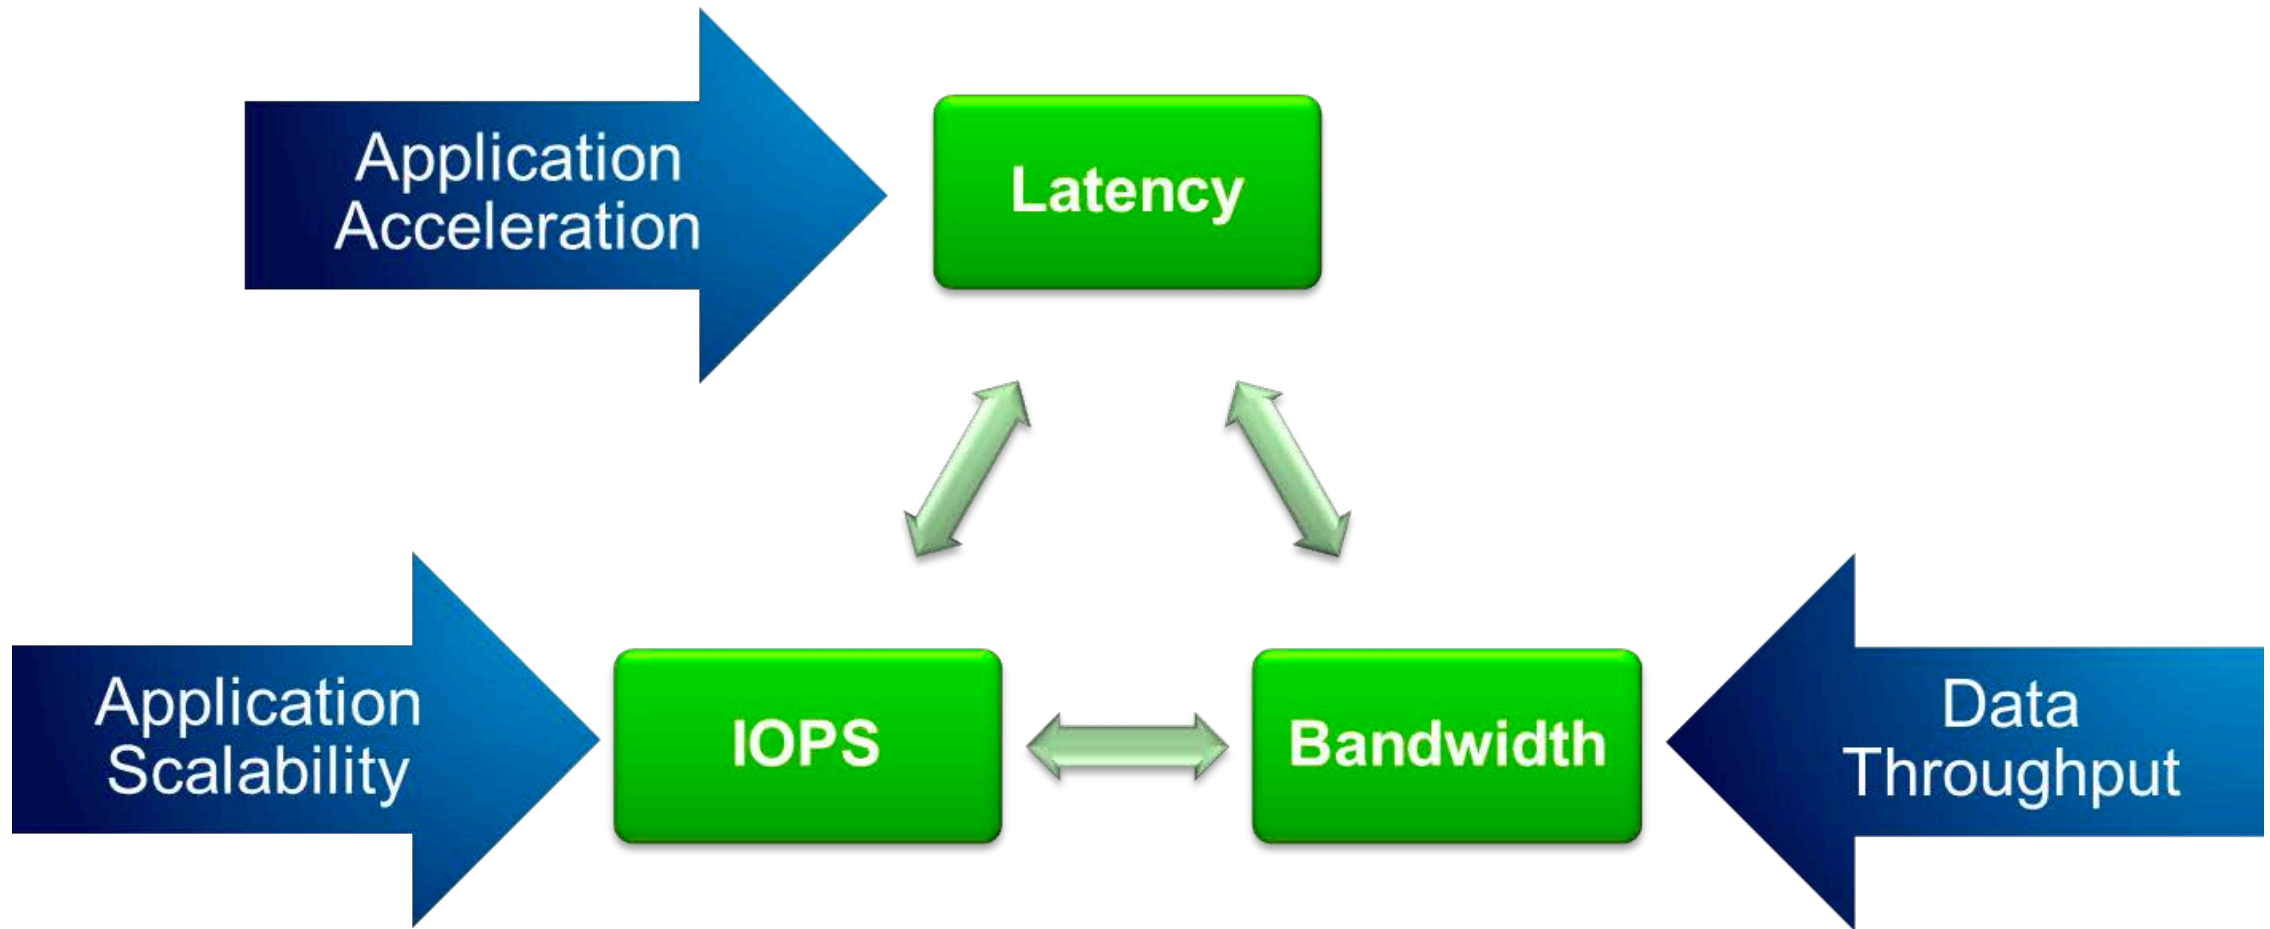
I is for **Isolation**: If multiple transactions are running at the same time they'll run independently and the end result will be as if they were executed one after the other.

D is for **Durability**: Once a transaction is done, it stays done (even if the system crashes thereafter).

[mysqld]

`innodb_flush_log_at_trx_commit=1` (1: durable, 0: Not durable, 2: Depends os/hardware)


Translating Storage Through Database Application



<https://flashdba.com/2013/04/08/the-fundamental-characteristics-of-storage/>


Monitor the Storage:

Read Latency




	max	avg	current
VOL_ASTNOD15_ASLEPPSTL01	2.6 ms	2.6 ms	2.6 ms
VOL_ASTNOD16_SPLUNKINDEXS01	4.9 ms	1.0 ms	1.6 ms
VOL_ASTNOD15_ASLEPYMPMYQ	4.8 ms	1.1 ms	1.1 ms
VOL_ASTNOD15_ASLEMPSTPMYQ03	3.0 ms	1.0 ms	1.1 ms
VOL_ASTNOD15_ASLEMPSTPMYQ01	1.2 ms	0.9 ms	0.9 ms

Read Throughput




	max	avg	current
VOL_ASTNOD16_SPLUNKINDEXS01	797	103	246
VOL_ASTNOD15_SPLUNKINDEXS01	665	58	120
VOL_ASTNOD16_ASLEBNO	41	9	8
VOL_ASTNOD15_K8SROOT01	516	10	2
VOL_ASTNOD15_VMWINROOTS03	111	22	0

Read IOPs




	max	avg	current
VOL_ASTNOD16_SPLUNKINDEXS01	15314	2349	4709
VOL_ASTNOD15_SPLUNKINDEXS01	12411	1231	2103
VOL_ASTNOD16_ASLEBNO	1304	404	436
VOL_ASTNOD15_K8SROOT01	34695	621	18
VOL_ASTNOD15_VMWINROOTS03	1940	489	6

Write Latency




	max	avg	current
VOL_ASTNOD15_ASLEPKPSIE01	10.3 ms	0.9 ms	1.8 ms
VOL_ASTNOD15_ASLENTPRGS01	1.8 ms	1.2 ms	1.3 ms
VOL_ASTNOD16_SPLUNKINDEXS01	2.4 ms	0.8 ms	0.7 ms
VOL_ASTNOD16_MSTVM01	4.0 ms	2.4 ms	0.4 ms
VOL_ASTNOD16_QLIKSENSE_FS	1.3 ms	0.6 ms	0.3 ms

Write Throughput



	max	avg	current
VOL_ASTNOD16_ASLEMGPMYQ05	63	37	47
VOL_ASTNOD15_ASLEMGPMYQ04	44	31	39
VOL_ASTNOD16_SPLUNKINDEXS01	97	33	36
VOL_ASTNOD15_SPLUNKINDEXS01	79	27	30
VOL_ASTNOD16_ASLOSSPMYQ	78	18	25

Write IOPs



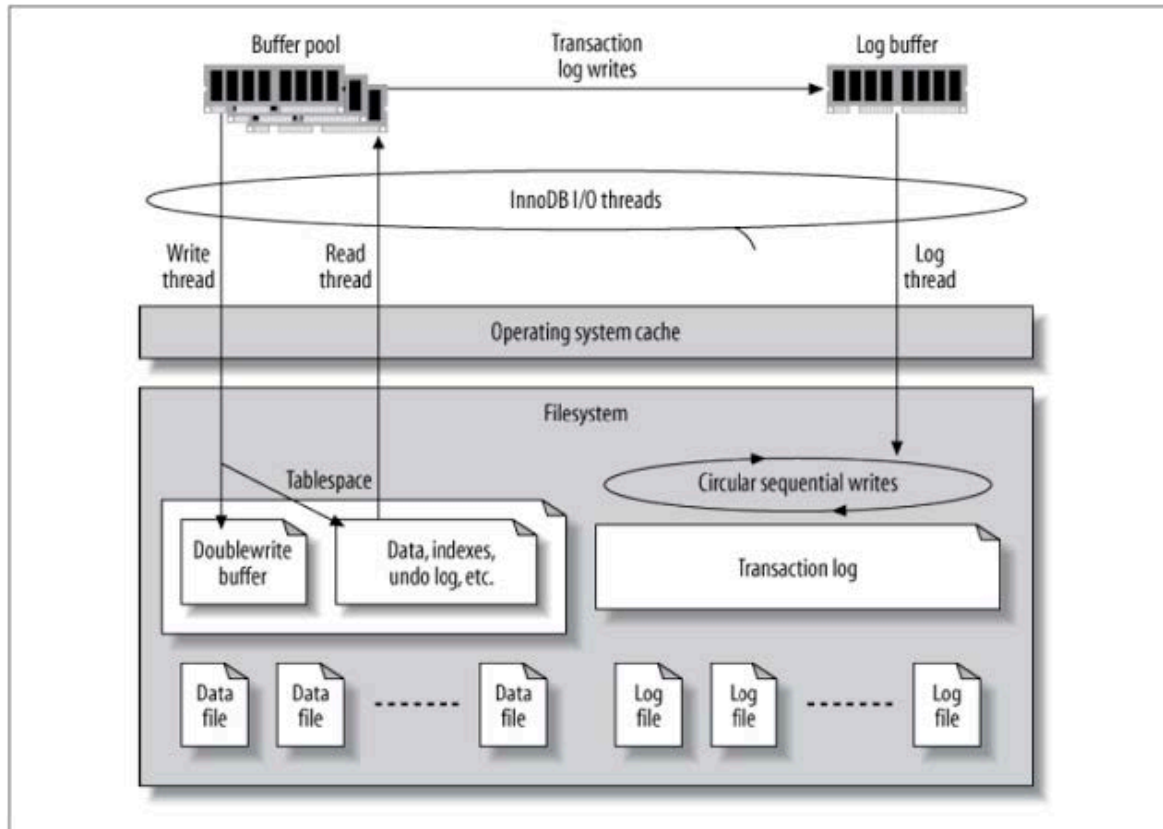
	max	avg	current
VOL_ASTNOD15_ASLEIPMYQ	3416	1480	1997
VOL_ASTNOD16_ASLOSSPMYQ	3086	1327	1782
VOL_ASTNOD16_SPLUNKINDEXS01	2556	919	1053
VOL_ASTNOD15_SPLUNKINDEXS01	1690	699	865
VOL_ASTNOD15_ASLELAC	3118	871	824

Use PMM



INNODB I/O & related settings

- innodb_flush_method O_DIRECT | O_DSYNC
- innodb_old_blocks_time
- innodb_thread_concurrency
- innodb_buffer_pool_instances
- innodb_read_io_threads
- innodb_write_io_thread
- innodb_io_capacity
- innodb_io_capacity_max

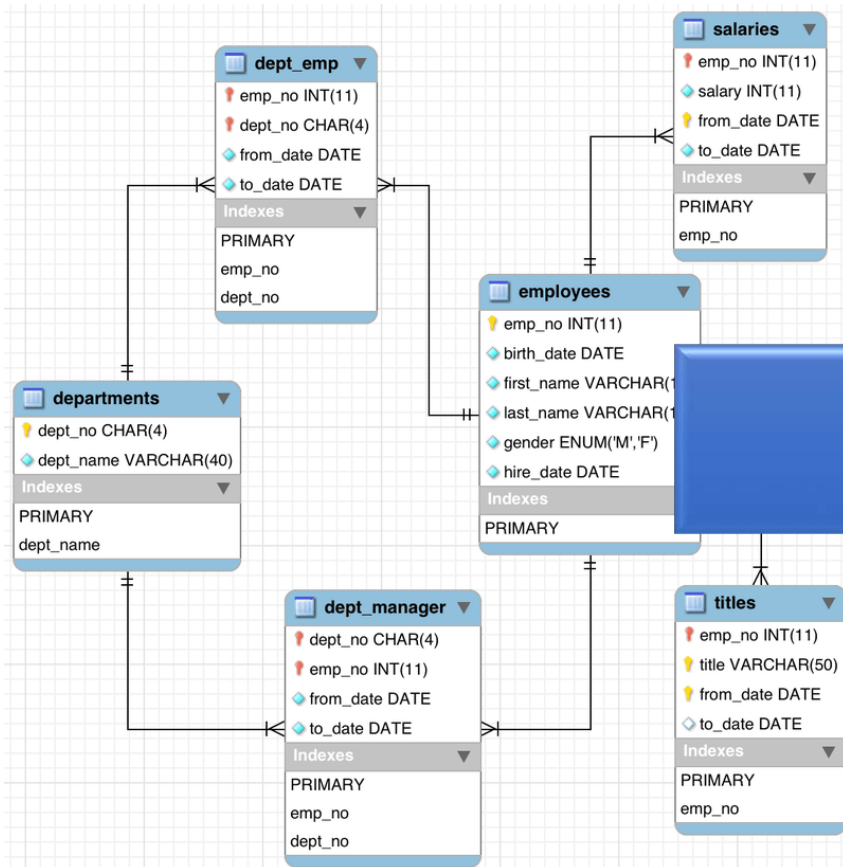


Important in InnoDB Components :

- 1) InnoDB buffer pool size
- 2) Transaction log buffer
- 3) InnoDB IO threads
- 4) Transaction Log files
- 5) Table-space files
- 6) Datafiles

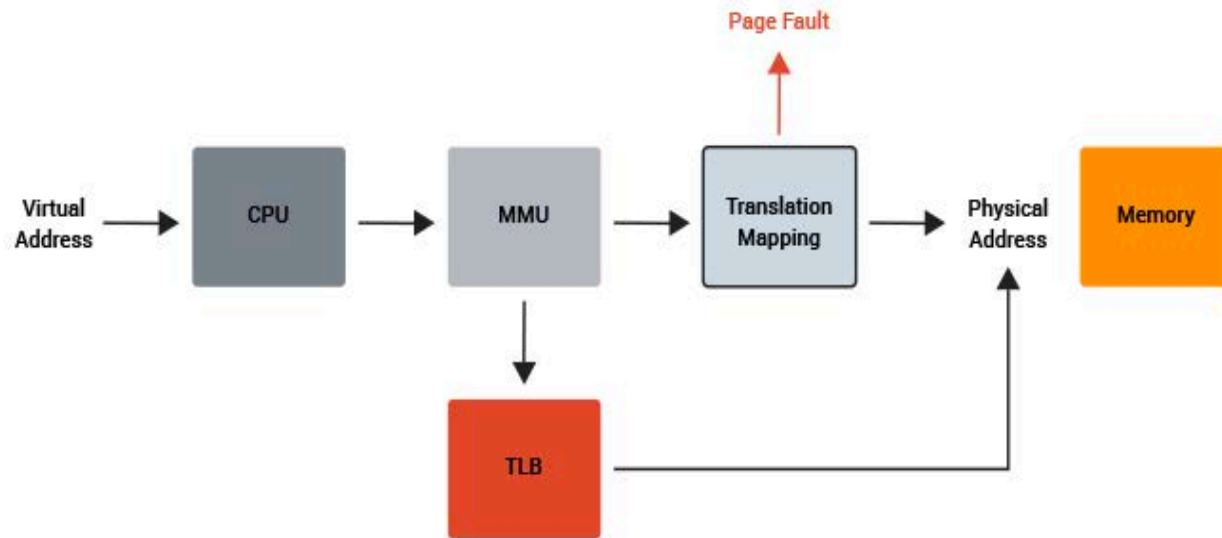
InnoDB Buffer Pool Size

```
SELECT CEILING(Total_InnoDB_Bytes*1.6/POWER(1024,3)) RIBPS  
FROM (SELECT SUM(data_length+index_length) Total_InnoDB_Bytes  
FROM information_schema.tables WHERE engine='InnoDB') A;
```



- innodb_buffer_pool_size=128G

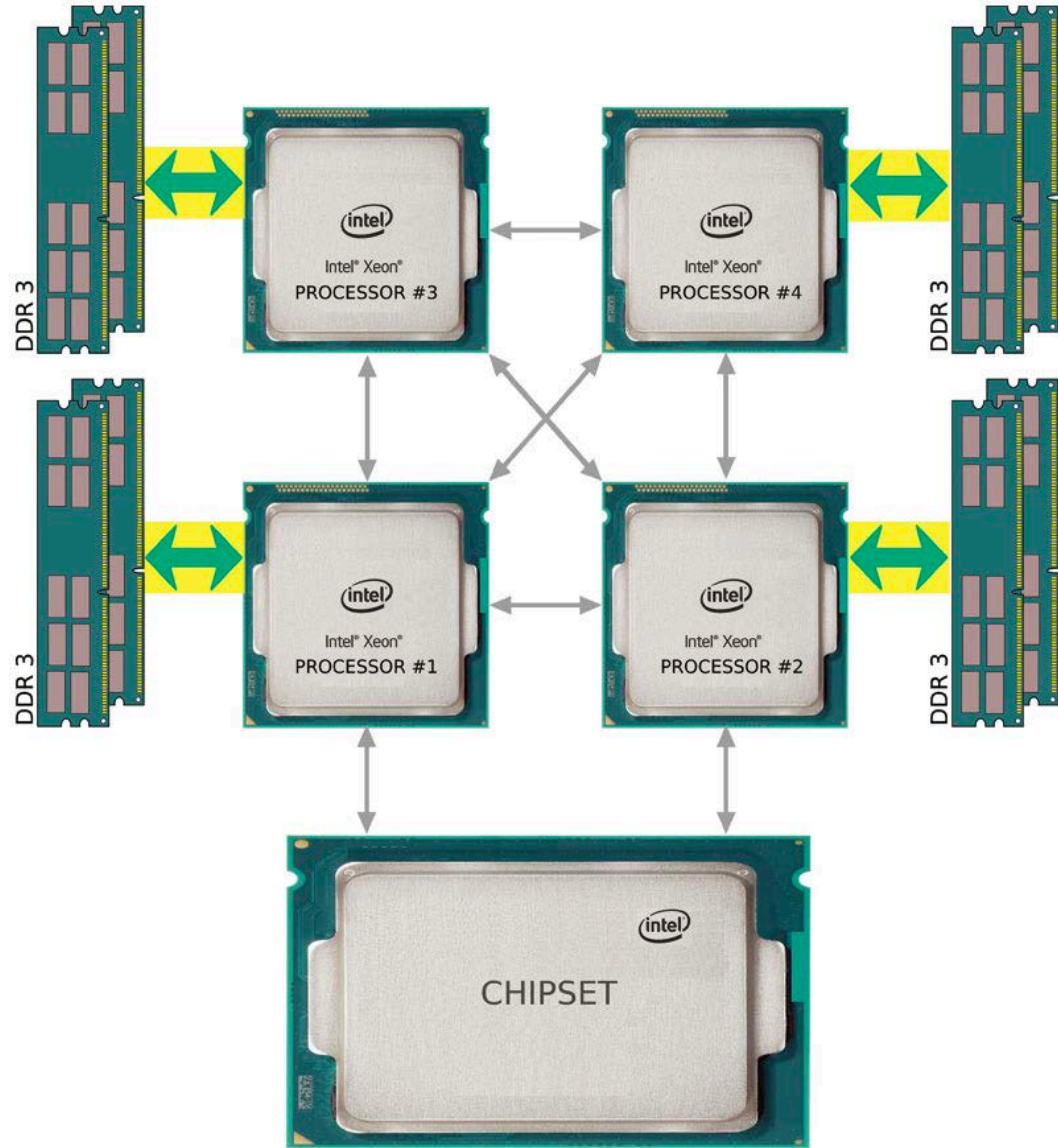
Page Faults!



`ps -eo min_flt,maj_flt,cmd`

```
gary@gary-1304-VB: ~  
221 0 upstart-udev-bridge --daemon  
1483 0 /lib/systemd/systemd-udev --daemon  
1778 0 dbus-daemon --system --fork  
0 0 [kpsmoused]  
0 0 [kworker/0:2]  
0 0 [iprt]  
404 2 /usr/sbin/bluetoothd  
844 1 /lib/systemd/systemd-logind  
0 0 [krfcommd]  
379 0 rsyslogd  
488 0 avahi-daemon: running [gary-1304-VB.local]  
87 0 avahi-daemon: chroot helper  
211 0 upstart-file-bridge --daemon  
6530 35 smbd -F  
164 0 upstart-socket-bridge --daemon  
942 1 smbd -F  
1304 19 /usr/sbin/ModemManager  
1891 3 NetworkManager  
2958 1 /usr/lib/policykit-1/polkitd --no-debug  
539 4 /usr/sbin/dnsmasq --no-resolv --keep-in-foreground --no-hosts --bi  
1636 1 /usr/sbin/winbindd -F  
545 2 /usr/sbin/winbindd -F  
413 0 nmbd -D  
374 0 /sbin/getty -8 38400 tty4
```

NUMA ARCHITECTURE



```
[mysqld]  
# Percona support case (#174734) 12.04.2017  
innodb_numa_interleave=1
```

```
[mysqld_safe]  
flush_caches
```

Even vNuma

The image shows a screenshot of the VMware vSphere configuration interface for a virtual machine's CPU settings. The interface has four tabs: "Virtual Hardware" (selected), "VM Options", "SDRS Rules", and "vApp Options". A red box highlights the "CPU" section, which includes the following settings:

- CPU:** 6 (with an information icon)
- Cores per Socket:** 3 (with "Sockets: 2" displayed next to it)
- CPU Hot Plug:** Enable CPU Hot Add

Below the highlighted section, other CPU-related settings are visible:

- Reservation:** 0 MHz
- Limit:** Unlimited MHz
- Shares:** Normal, 6000
- CPUID Mask:** Expose the NX/XD flag to guest (with an "Advanced..." link)
- Hardware virtualization:** Expose hardware assisted virtualization to the guest OS (with an information icon)
- Performance counters:** Enable virtualized CPU performance counters

Large Pages

[mysqld]

large-pages

Applications that perform a lot of memory accesses may obtain performance improvements by using large pages due to reduced Translation Lookaside Buffer (TLB) misses.


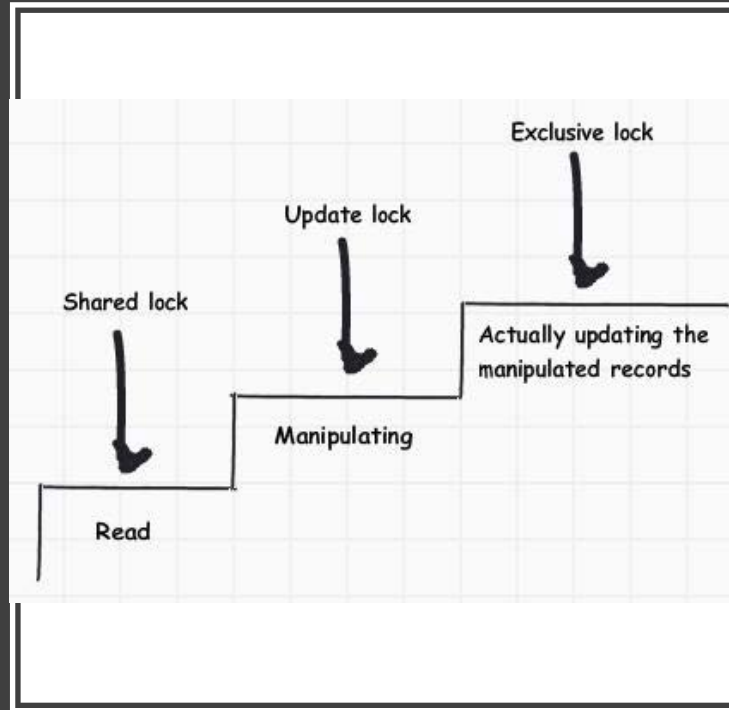

In MySQL, large pages can be used by InnoDB, to allocate memory for its buffer pool and additional memory pool. Standard use of large pages in MySQL attempts to use the largest size supported, up to 4MB

Before large pages can be used on Linux, the kernel must be enabled to support them and it is necessary to configure the HugeTLB memory pool.

Or jemalloc

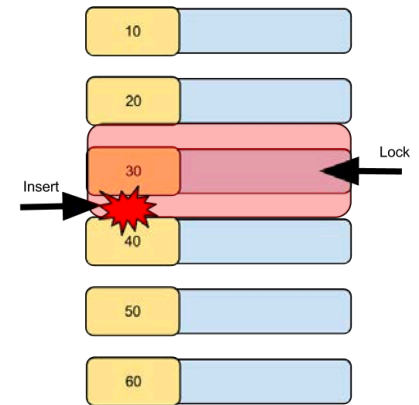

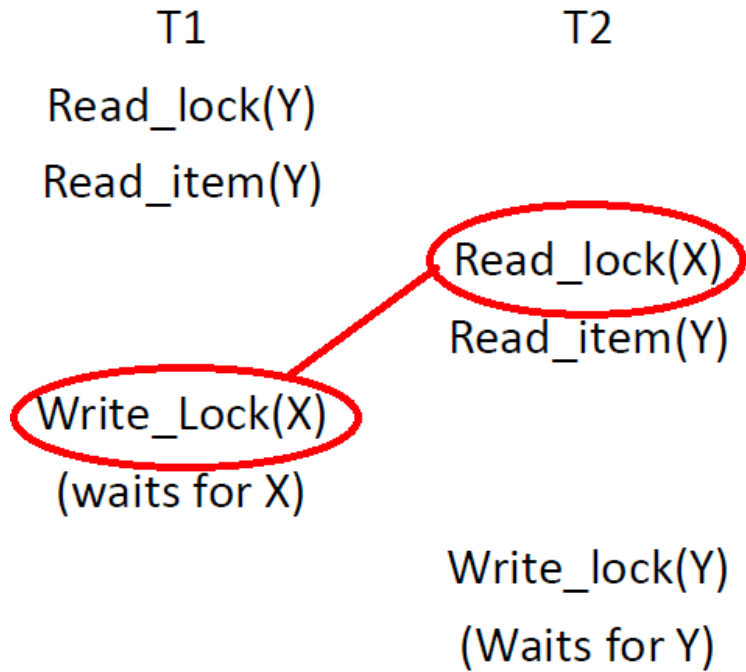
The benefits of [jemalloc versus glibc](#) for use with MySQL have been widely discussed. With jemalloc (along with [Transparent Huge Pages disabled](#)) you have less memory fragmentation, and thus more efficient resource management of the available server memory.

<https://www.percona.com/blog/2017/01/03/enabling-and-disabling-jemalloc-on-percona-server/>



InnoDB Transactions

Avoid deadlocks: Show engine innodb status



Connections

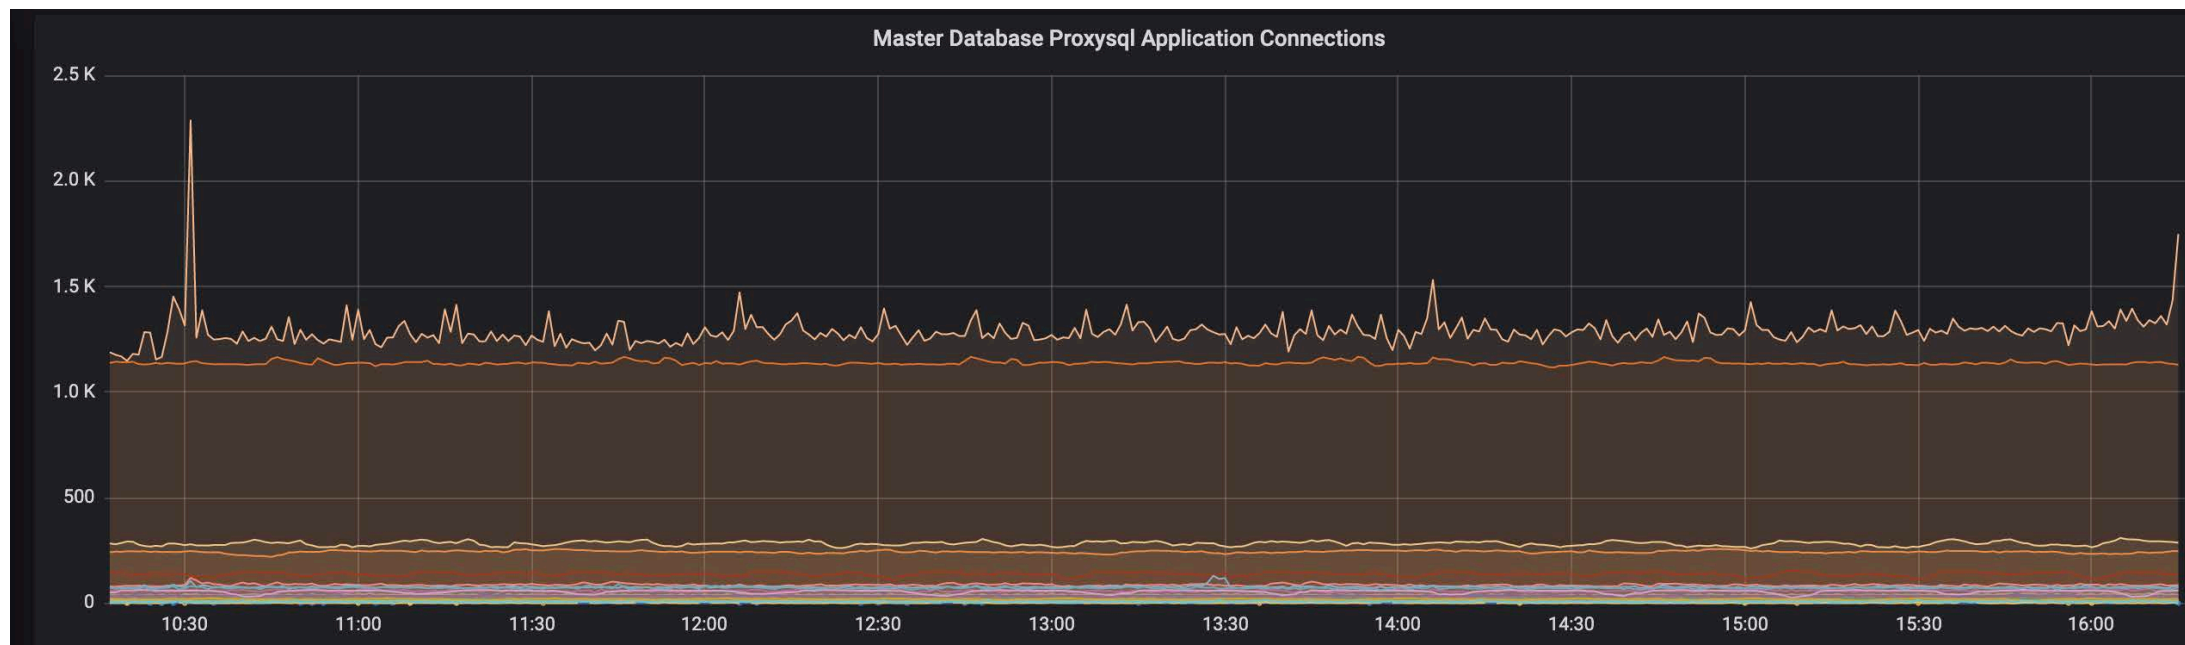
光 HikariCP

- BoneCP
- C3p0
- DBCP2
- ..
- .

* mysql (php)




Kubernetes autoscale !






ProxySQL

-
- ✓ Multiplexing
 - ✓ Load balancing
 - ✓ Query Routing



Multiplexing

Application servers



ProxySQL

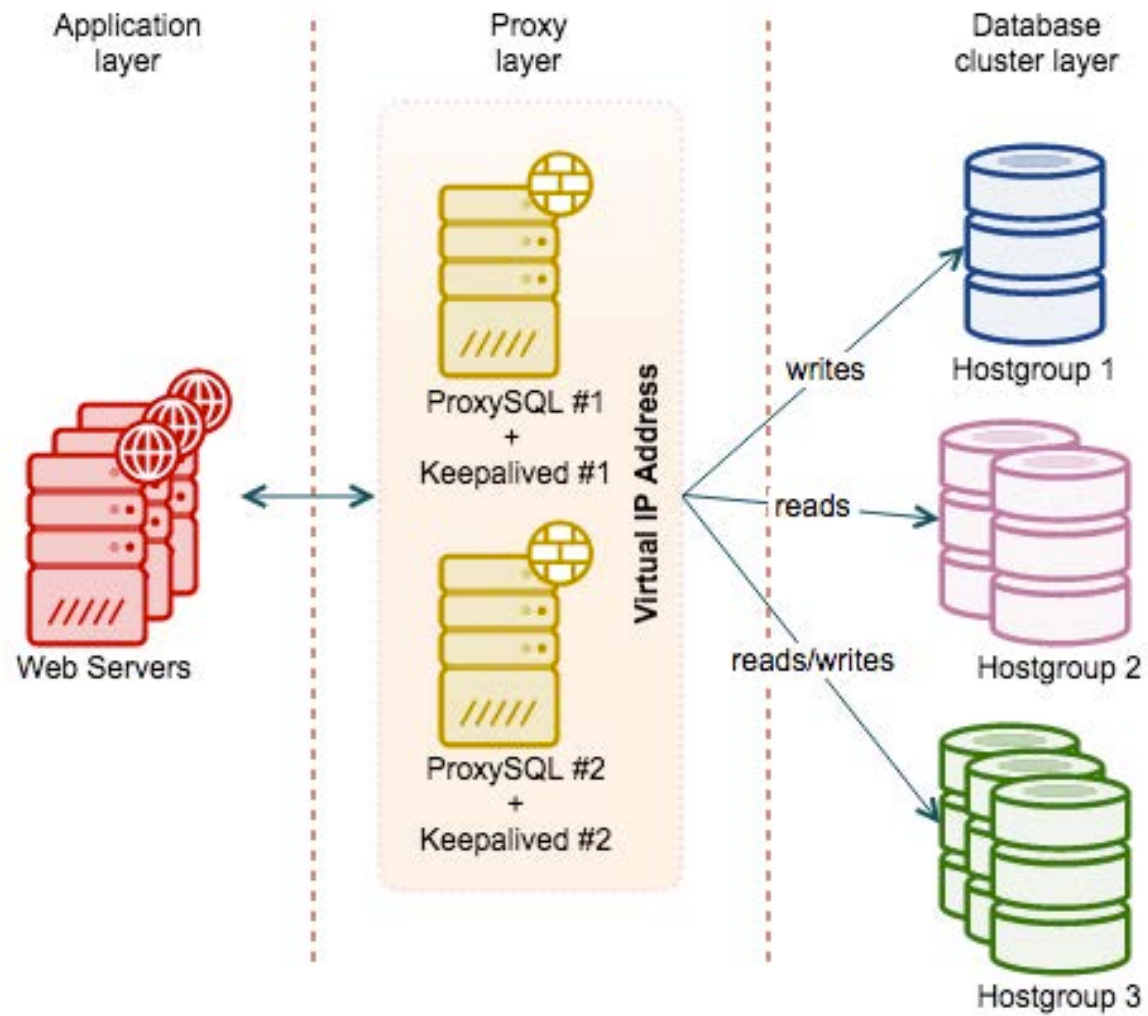


Multiplexing

MySQL



Load Balancing



Query Routing

```
$query = "SELECT /* hostgroup=1 */
```

Q&A