

17 Things Developers should know about Databases


Peter Zaitsev

CEO, Percona
Nizhny Novgorod, Russia

November 9th, 2019



Thank you!



iCluster
Global IT Community

> **it52**

Нижегородское IT сообщество



РОСБАНК

About Percona

Open Source Database Solutions Company

Support, Managed Services, Consulting, Training, Engineering

Focus on MySQL, MariaDB, MongoDB, PostgreSQL

Support Cloud DBaaS Variants on major clouds

Develop Database Software and Tools

Release Everything as 100% Free and Open Source

Widely Deployed Open Source Software



PERCONA
Server for MySQL

5,000,000+ downloads



PERCONA
Monitoring and Management

175,000+ downloads



PERCONA
XtraBackup

4,500,000+ downloads



PERCONA
Server for MongoDB

450,000+ downloads



PERCONA
Toolkit

2,000,000+ downloads



PERCONA
XtraDB Cluster

1,500,000+ downloads

Freshly Released

Percona Distribution for PostgreSQL



Who Are you ?

**More
Developer ?**

More OPS ?

Ops

Focused on Database Only

Generalist

Programming Language

**What Programming
Languages does your
team use ?**

Devs vs Ops

**Is there Separation
between Devs and Ops in
your Organization ?**

Any Tension ?

**Any Tension Between
Devs and Ops ?**

Devs vs Ops Conflict

Devs

- Why is this stupid database always the problem.
- Why can't it just work and work fast

Ops

- Why do not learn schema design
- Why do not you write optimized queries
- Why do not you think about capacity planning

Database are Special Snowflakes

Production Database Changes are Risky and Costly

Often need special handling with Continuous Deployment

Large Organizations

**Ops vs Ops have conflict
too**

Database Responsibility

**Shared Responsibility for
Ultimate Success**

Top Recommendations for Developers

Learn Database Basics

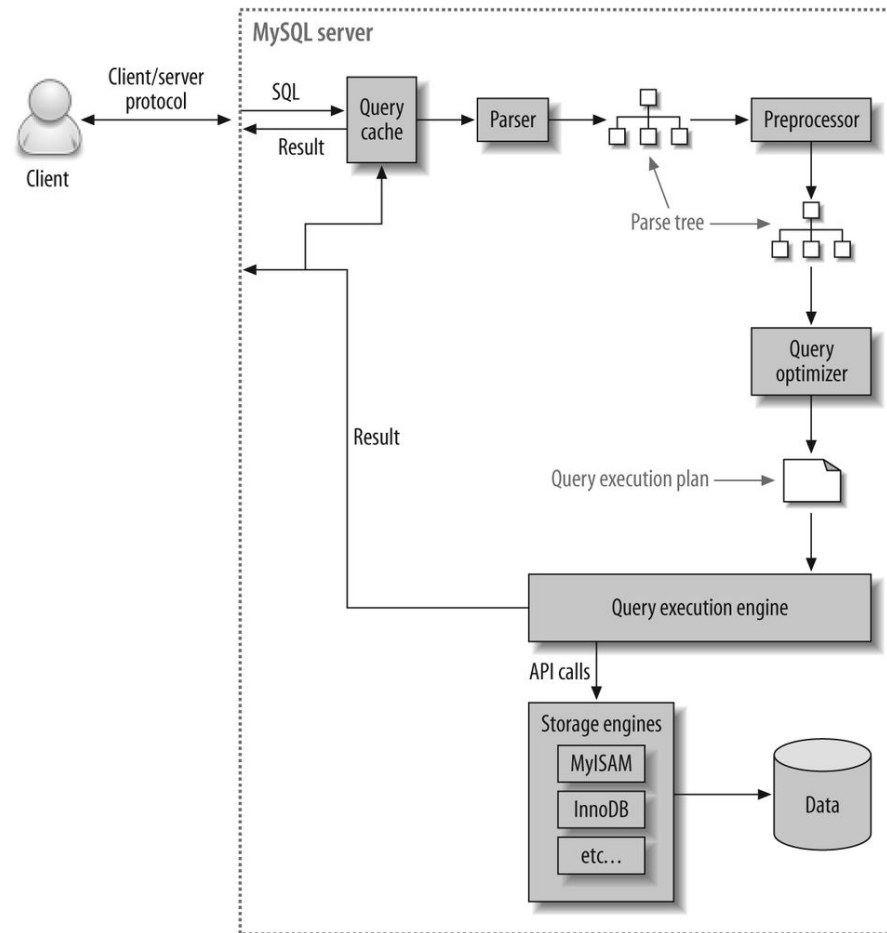
You can't build great database powered applications if you do not understand how databases work

Schema Design

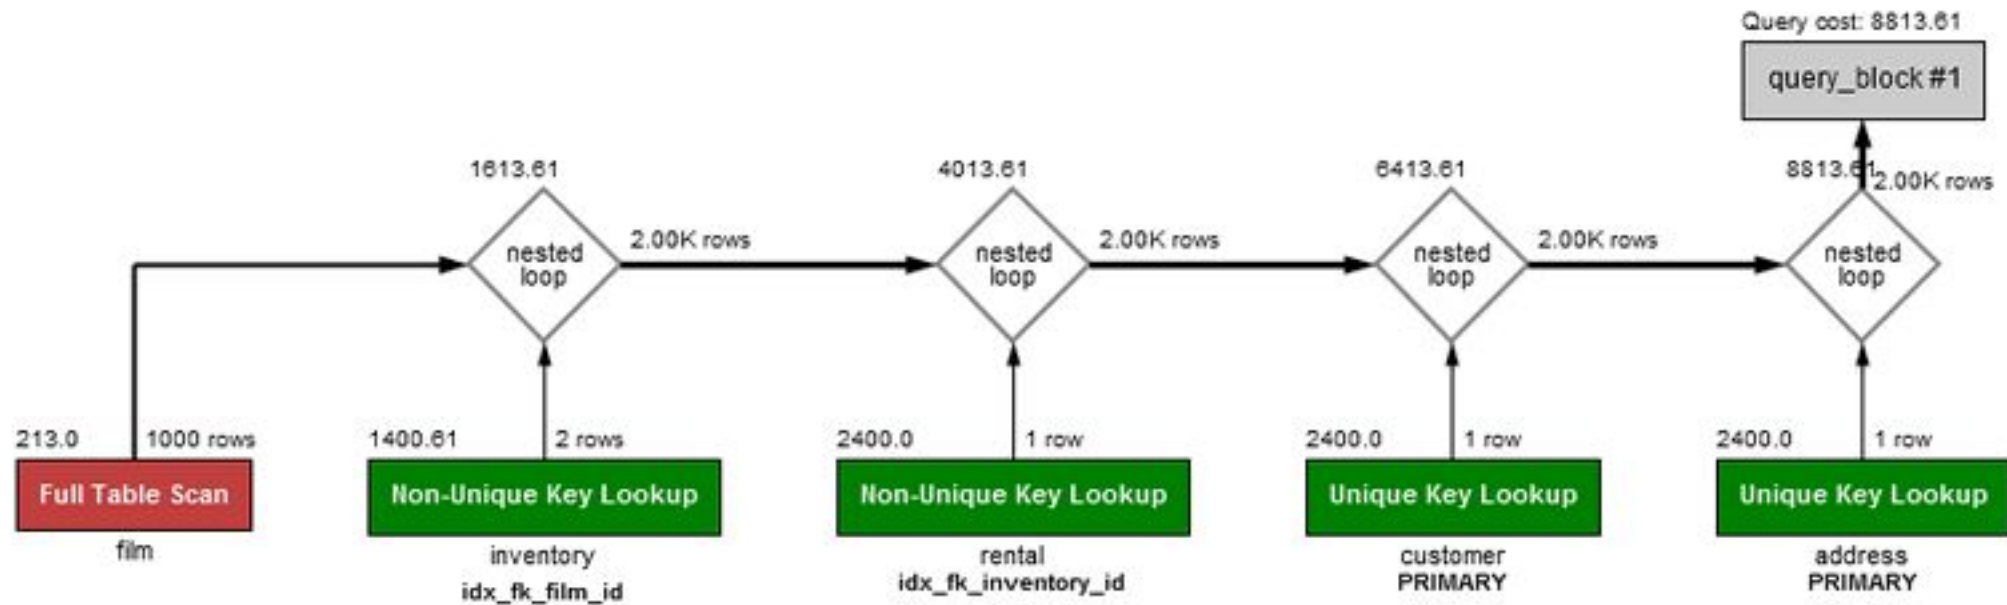
Power of the Database Language

How Database Executes the Query

Query Execution Diagram

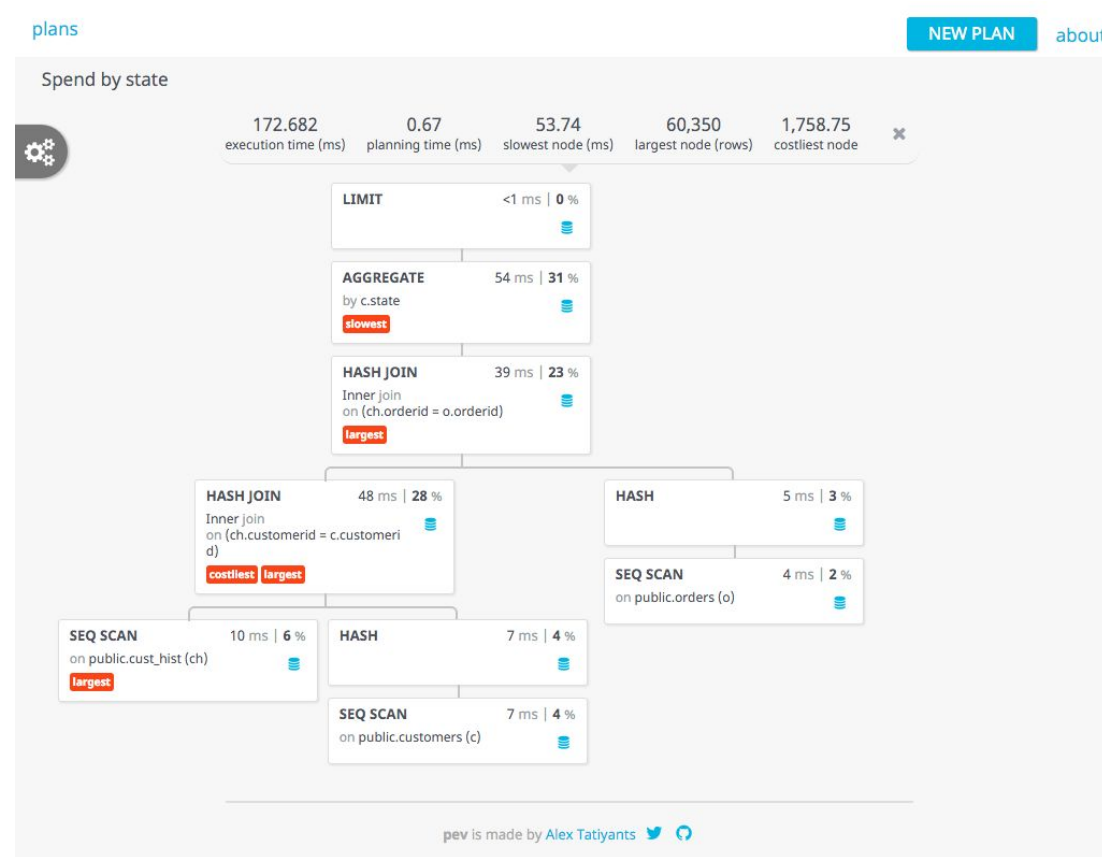


EXPLAIN



<https://dev.mysql.com/doc/refman/8.0/en/execution-plan-information.html>

Visualizing PostgreSQL Plan with PEV

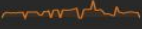





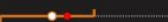




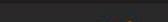
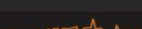

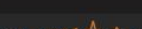
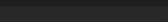

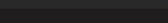
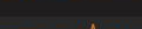
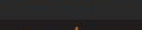
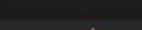


<http://tatiyants.com/postgres-query-plan-visualization/>

Which Queries are Causing the Load

Environment		#	Query	Load	Query Count	
<input type="checkbox"/> mytest	99.99%					
<input type="checkbox"/> n/a	0.01%					
Database						
<input type="checkbox"/> n/a	99.99%					
<input type="checkbox"/> pmm-managed	<0.01					
<input type="checkbox"/> postgres	<0.01					
Schema						
<input type="checkbox"/> tpcc1	20.44%					
<input type="checkbox"/> tpcc5	17.36%					
<input type="checkbox"/> tpcc3	17.32%					
<input type="checkbox"/> tpcc2	17.31%					
<input type="checkbox"/> tpcc4	17.29%					
Node Name						
<input type="checkbox"/> mysql2	41.31%					
<input type="checkbox"/> mysql1	20.47%					
<input type="checkbox"/> mysql3	17.32%					
<input type="checkbox"/> mysql4	17.31%					
<input type="checkbox"/> mysql5	17.29%					
<input type="checkbox"/> mysql6	17.29%					
<input type="checkbox"/> mysql7	17.29%					
<input type="checkbox"/> mysql8	17.29%					
<input type="checkbox"/> mysql9	17.29%					
<input type="checkbox"/> mysql10	17.29%					
<input type="checkbox"/> mysql11	17.29%					
<input type="checkbox"/> mysql12	17.29%					
<input type="checkbox"/> mysql13	17.29%					
<input type="checkbox"/> mysql14	17.29%					
<input type="checkbox"/> mysql15	17.29%					
<input type="checkbox"/> mysql16	17.29%					
<input type="checkbox"/> mysql17	17.29%					
<input type="checkbox"/> mysql18	17.29%					
<input type="checkbox"/> mysql19	17.29%					
<input type="checkbox"/> mysql20	17.29%					
<input type="checkbox"/> mysql21	17.29%					
<input type="checkbox"/> mysql22	17.29%					
<input type="checkbox"/> mysql23	17.29%					
<input type="checkbox"/> mysql24	17.29%					
<input type="checkbox"/> mysql25	17.29%					
<input type="checkbox"/> mysql26	17.29%					
<input type="checkbox"/> mysql27	17.29%					
<input type="checkbox"/> mysql28	17.29%					
<input type="checkbox"/> mysql29	17.29%					
<input type="checkbox"/> mysql30	17.29%					
<input type="checkbox"/> mysql31	17.29%					
<input type="checkbox"/> mysql32	17.29%					
<input type="checkbox"/> mysql33	17.29%					
<input type="checkbox"/> mysql34	17.29%					
<input type="checkbox"/> mysql35	17.29%					
<input type="checkbox"/> mysql36	17.29%					
<input type="checkbox"/> mysql37	17.29%					
<input type="checkbox"/> mysql38	17.29%					
<input type="checkbox"/> mysql39	17.29%					
<input type="checkbox"/> mysql40	17.29%					
<input type="checkbox"/> mysql41	17.29%					
<input type="checkbox"/> mysql42	17.29%					
<input type="checkbox"/> mysql43	17.29%					
<input type="checkbox"/> mysql44	17.29%					
<input type="checkbox"/> mysql45	17.29%					
<input type="checkbox"/> mysql46	17.29%					
<input type="checkbox"/> mysql47	17.29%					
<input type="checkbox"/> mysql48	17.29%					
<input type="checkbox"/> mysql49	17.29%					
<input type="checkbox"/> mysql50	17.29%					
<input type="checkbox"/> mysql51	17.29%					
<input type="checkbox"/> mysql52	17.29%					
<input type="checkbox"/> mysql53	17.29%					
<input type="checkbox"/> mysql54	17.29%					
<input type="checkbox"/> mysql55	17.29%					
<input type="checkbox"/> mysql56	17.29%					
<input type="checkbox"/> mysql57	17.29%					
<input type="checkbox"/> mysql58	17.29%					
<input type="checkbox"/> mysql59	17.29%					
<input type="checkbox"/> mysql60	17.29%					
<input type="checkbox"/> mysql61	17.29%					
<input type="checkbox"/> mysql62	17.29%					
<input type="checkbox"/> mysql63	17.29%					
<input type="checkbox"/> mysql64	17.29%					
<input type="checkbox"/> mysql65	17.29%					
<input type="checkbox"/> mysql66	17.29%					
<input type="checkbox"/> mysql67	17.29%					
<input type="checkbox"/> mysql68	17.29%					
<input type="checkbox"/> mysql69	17.29%					
<input type="checkbox"/> mysql70	17.29%					
<input type="checkbox"/> mysql71	17.29%					
<input type="checkbox"/> mysql72	17.29%					
<input type="checkbox"/> mysql73	17.29%					
<input type="checkbox"/> mysql74	17.29%					
<input type="checkbox"/> mysql75	17.29%					
<input type="checkbox"/> mysql76	17.29%					
<input type="checkbox"/> mysql77	17.29%					
<input type="checkbox"/> mysql78	17.29%					
<input type="checkbox"/> mysql79	17.29%					
<input type="checkbox"/> mysql80	17.29%					
<input type="checkbox"/> mysql81	17.29%					
<input type="checkbox"/> mysql82	17.29%					
<input type="checkbox"/> mysql83	17.29%					
<input type="checkbox"/> mysql84	17.29%					
<input type="checkbox"/> mysql85	17.29%					
<input type="checkbox"/> mysql86	17.29%					
<input type="checkbox"/> mysql87	17.29%					
<input type="checkbox"/> mysql88	17.29%					
<input type="checkbox"/> mysql89	17.29%					
<input type="checkbox"/> mysql90	17.29%					
<input type="checkbox"/> mysql91	17.29%					
<input type="checkbox"/> mysql92	17.29%					
<input type="checkbox"/> mysql93	17.29%					
<input type="checkbox"/> mysql94	17.29%					
<input type="checkbox"/> mysql95	17.29%					
<input type="checkbox"/> mysql96	17.29%					
<input type="checkbox"/> mysql97	17.29%					
<input type="checkbox"/> mysql98	17.29%					
<input type="checkbox"/> mysql99	17.29%					
<input type="checkbox"/> mysql100	17.29%					
<input type="checkbox"/> mysql101	17.29%					
<input type="checkbox"/> mysql102	17.29%					
<input type="checkbox"/> mysql103	17.29%					
<input type="checkbox"/> mysql104	17.29%					
<input type="checkbox"/> mysql105	17.29%					
<input type="checkbox"/> mysql106	17.29%					
<input type="checkbox"/> mysql107	17.29%					
<input type="checkbox"/> mysql108	17.29%					
<input type="checkbox"/> mysql109	17.29%					
<input type="checkbox"/> mysql110	17.29%					
<input type="checkbox"/> mysql111	17.29%					
<input type="checkbox"/> mysql112	17.29%					
<input type="checkbox"/> mysql113	17.29%					
<input type="checkbox"/> mysql114	17.29%					
<input type="checkbox"/> mysql115	17.29%					
<input type="checkbox"/> mysql116	17.29%					
<input type="checkbox"/> mysql117	17.29%					
<input type="checkbox"/> mysql118	17.29%					
<input type="checkbox"/> mysql119	17.29%					
<input type="checkbox"/> mysql120	17.29%					
<input type="checkbox"/> mysql121	17.29%					
<input type="checkbox"/> mysql122	17.29%					
<input type="checkbox"/> mysql123	17.29%					
<input type="checkbox"/> mysql124	17.29%					
<input type="checkbox"/> mysql125	17.29%					
<input type="checkbox"/> mysql126	17.29%					
<input type="checkbox"/> mysql127	17.29%					
<input type="checkbox"/> mysql128	17.29%					
<input type="checkbox"/> mysql129	17.29%					
<input type="checkbox"/> mysql130	17.29%					
<input type="checkbox"/> mysql131	17.29%					
<input type="checkbox"/> mysql132	17.29%					
<input type="checkbox"/> mysql133	17.29%					
<input type="checkbox"/> mysql134	17.29%					
<input type="checkbox"/> mysql135	17.29%					
<input type="checkbox"/> mysql136	17.29%					
<input type="checkbox"/> mysql137	17.29%					
<input type="checkbox"/> mysql138	17.29%					
<input type="checkbox"/> mysql139	17.29%					
<input type="checkbox"/> mysql140	17.29%					
<input type="checkbox"/> mysql141	17.29%					
<input type="checkbox"/> mysql142	17.29%					
<input type="checkbox"/> mysql143	17.29%					
<input type="checkbox"/> mysql144	17.29%					
<input type="checkbox"/> mysql145	17.29%					
<input type="checkbox"/> mysql146	17.29%					
<input type="checkbox"/> mysql147	17.29%					
<input type="checkbox"/> mysql148	17.29%					
<input type="checkbox"/> mysql149	17.29%					
<input type="checkbox"/> mysql150	17.29%					
<input type="checkbox"/> mysql151	17.29%					
<input type="checkbox"/> mysql152	17.29%					
<input type="checkbox"/> mysql153	17.29%					
<input type="checkbox"/> mysql154	17.29%					
<input type="checkbox"/> mysql155	17.29%					
<input type="checkbox"/> mysql156	17.29%					
<input type="checkbox"/> mysql157	17.29%					
<input type="checkbox"/> mysql158	17.29%					
<input type="checkbox"/> mysql159	17.29%					
<input type="checkbox"/> mysql160	17.29%					
<input type="checkbox"/> mysql161	17.29%					
<input type="checkbox"/> mysql162	17.29%					
<input type="checkbox"/> mysql163	17.29%					
<input type="checkbox"/> mysql164	17.29%					
<input type="checkbox"/> mysql165	17.29%					
<input type="checkbox"/> mysql166	17.29%					
<input type="checkbox"/> mysql167	17.29%					
<input type="checkbox"/> mysql168	17.29%					
<input type="checkbox"/> mysql169	17.29%					
<input type="checkbox"/> mysql170	17.29%					
<input type="checkbox"/> mysql171	17.29%					
<input type="checkbox"/> mysql172	17.29%					
<input type="checkbox"/> mysql173	17.29%					
<input type="checkbox"/> mysql174	17.29%					
<input type="checkbox"/> mysql175	17.29%					
<input type="checkbox"/> mysql176	17.29%					
<input type="checkbox"/> mysql177	17.29%					
<input type="checkbox"/> mysql178	17.29%					
<input type="checkbox"/> mysql179	17.29%					
<input type="checkbox"/> mysql180	17.29%					
<input type="checkbox"/> mysql181	17.29%					
<input type="checkbox"/> mysql182	17.29%					
<input type="checkbox"/> mysql183	17.29%					
<input type="checkbox"/> mysql184	17.29%					
<input type="checkbox"/> mysql185	17.29%					
<input type="checkbox"/> mysql186	17.29%					
<input type="checkbox"/> mysql187	17.29%					
<input type="checkbox"/> mysql188	17.29%					
<input type="checkbox"/> mysql189	17.29%					
<input type="checkbox"/> mysql190	17.29%					
<input type="checkbox"/> mysql191	17.29%					
<input type="checkbox"/> mysql192	17.29%					
<input type="checkbox"/> mysql193	17.29%					
<input type="checkbox"/> mysql194	17.29%					
<input type="checkbox"/> mysql195	17.29%					
<input type="checkbox"/> mysql196	17.29%					
<input type="checkbox"/> mysql197	17.29%					
<input type="checkbox"/> mysql198	17.29%					
<input type="checkbox"/> mysql199	17.29%					
<input type="checkbox"/> mysql200	17.29%					
<input type="checkbox"/> mysql201	17.29%					
<input type="checkbox"/> mysql202	17.29%					
<input type="checkbox"/> mysql203	17.29%					
<input type="checkbox"/> mysql204	17.29%					
<input type="checkbox"/> mysql205	17.29%					
<input type="checkbox"/> mysql206	17.29%					
<input type="checkbox"/> mysql207	17.29%					
<input type="checkbox"/> mysql208	17.29%					
<input type="checkbox"/> mysql209	17.29%					
<input type="checkbox"/> mysql210	17.29%					
<input type="checkbox"/> mysql211	17.29%					
<input type="checkbox"/> mysql212	17.29%					
<input type="checkbox"/> mysql213	17.29%					
<input type="checkbox"/> mysql214	17.29%					
<input type="checkbox"/> mysql215	17.29%					
<input type="checkbox"/> mysql216	17.29%					
<input type="checkbox"/> mysql217	17.29%					
<input type="checkbox"/> mysql218	17.29%					
<input type="checkbox"/> mysql219	17.29%					
<input type="checkbox"/> mysql220	17.29%					
<input type="checkbox"/> mysql221	17.29%					
<input type="checkbox"/> mysql222	17.29%					
<input type="checkbox"/> mysql223	17.29%					
<input type="checkbox"/> mysql224	17.29%					
<input type="checkbox"/> mysql225	17.29%					
<input type="checkbox"/> mysql226	17.29%					
<input type="checkbox"/> mysql227	17.29%					
<input type="checkbox"/> mysql228	17.29%					
<input type="checkbox"/> mysql229	17.29%					
<input type="checkbox"/> mysql230	17.29%					
<input type="checkbox"/> mysql231	17.29%					
<input type="checkbox"/> mysql232	17.29%					
<input type="checkbox"/> mysql233	17.29%					
<input type="checkbox"/> mysql234	17.29%					
<input type="checkbox"/> mysql235	17.29%					
<input type="checkbox"/> mysql236	17.29%					
<input type="checkbox"/> mysql237	17.29%					
<input type="checkbox"/> mysql238	17.29%					
<input type="checkbox"/> mysql239	17.29%					
<input type="checkbox"/> mysql240	17.29%					
<input type="checkbox"/> mysql241	17.29%					
<input type="checkbox"/> mysql242	17.29%					
<input type="checkbox"/> mysql243	17.29%					
<input type="checkbox"/> mysql244	17.29%					
<input type="checkbox"/> mysql245	17.29%					
<input type="checkbox"/> mysql246	17.29%					
<input type="checkbox"/> mysql247	17.29%					
<input type="checkbox"/> mysql248	17.29%					

Why Are they Causing this Load

SELECT sbtest				737F39F04B198EF6	
Metrics				Query first seen: ☉ Aug 3, 2017 1:55 PM ... Last seen: ☉ Today at 9:46 AM	
Metrics	Rate/Sec		Sum	Per Query Stats	
Query Count	104.05 (per sec)		374.58 k 4.27% of total		
Query Time	19.00 load		18:59:56 29.73% of total	183.66 ms avg	
Lock Time	0.11 (avg load)		0:06:42 1.35% of total 0.61% of query time	1.13 ms avg	
Innodb IO Read Wait	0.61 (avg load)		0:36:44 9.10% of total 3.38% of query time	6.20 ms avg	
Innodb Read Ops	52.35 (per sec)		188.45 k 7.62% of total	0.00 avg	
Innodb Read Bytes	857.64 KB (per sec)		3.09 GB 7.62% of total 16.38 KB avg io size	8.22 KB avg	
Innodb Distinct Pages			-	4.69 avg	
Rows Sent	10.41 k (per sec)		37.46 m 30.52% of total	100.00 avg	
Bytes Sent	1.30 MB (per sec)		4.67 GB 30.78% of total 124.71 Bytes bytes/row	12.47 KB avg	
Rows Examined	1.14 m (per sec)		4.11 b 39.17% of total 109.79 per row sent	10.47 k avg	
External Sorts (Filesort)	104.05 (per sec)		374.58 k 49.93% of total 100.00% of queries	-	
Full Table Scans	0.01 (per sec)		40.00 0.17% of total 0.01% of queries	-	
Queries Requiring Tmp Table In Memory	104.05 (per sec)		374.58 k 95.17% of total 100.00% of queries	-	

How to Improve their Performance

▼Example

```
SELECT DISTINCT c
FROM sbtest1
WHERE id
      BETWEEN 5559
      AND 5658
ORDER BY c
```

▼CREATE

```
CREATE TABLE `sbtest1` (
  `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_1` (`k`)
) ENGINE=InnoDB AUTO_INCREMENT=100000001 DEFAULT
```

▼JSON

Expand All

```
-{
  "query_block": -{
    "select_id": 1,
    "cost_info": +{...},
    "ordering_operation": -{
      "using_filesort": false,
      "duplicates_removal": -{
        "using_temporary_table": true,
        "using_filesort": true,
        "cost_info": +{...},
        "table": +{...}
      }
    }
  }
}
```

Check out PMM

<http://pmmdemo.percona.com>

PMM v 2 is now GA

How are Queries Executed ?

Single Threaded

Single Node

Distributed

Indexes

**Indexes are
Must**

**Indexes are
Expensive**

Capacity Planning

No Database can handle “unlimited scale”


Scalability is very application dependent

Trust Measurements more than Promises


Can be done or can be done Efficiently ?

Vertical and Horizontal Scaling

Vertical Scaling



Horizontal Scaling



Scalable != Efficient

The Systems promoted as scalable can be less efficient

Hadoop, Cassandra, TiDB are great examples

By only the wrong thing you can get in trouble

Throughput != Latency

If I tell you system can do
100.000 queries/sec
would you say it is fast ?

Speed of Light Limitations

High Availability Design Choices

You want instant durable replication over wide geography or Performance ?

Understanding Difference between High Availability and Disaster Recovery protocols

Network Bandwidth is not the same as Latency

Wide Area Networks

Tend not to be as stable as Local Area Networks (LAN)

Expect increased Jitter

Expect Short term unavailability

Also Understand

Connections to the database are expensive

Especially if doing TLS Handshake

Query Latency Tends to Add Up

Especially on real network and not your laptop

ORM (Object-Relational-Mapping)


Allows Developers to query the database without need to understand SQL

Can create SQL which is very inefficient

Learn SQL Generation “Hints”, Learn JPQL/HQL advanced features

Be ready to manually write SQL if there is no other choice

Understanding Optimal Concurrency



Embrace Queueing

Request Queueing is Normal

With requests coming at “Random Arrivals” some queueing will happen with any system scale

Should not happen too often or for very long

Queueing is “Cheaper” Close to the User

Benefits of Connection Pooling

**Avoiding
Connection
Overhead,
especially TLS/SSL**

**Avoiding using
Excessive Number
of Database
Connections**

**Multiplexing/Load
Management**

Configuring Connection Pool

Default and Maximum Connection Pool Size

Scaling Parameters

Combined Connection Pool Max Size should be smaller than number of connections database can support

Waiting for free connection to become available is OK

Scale Matters

Developing and Testing with Toy Database is risky

Queries Do not slow down linearly

The slowest query may slow down most rapidly

Memory or Disk

Data Accessed in memory is much faster than on disk

It is true even with modern SSDs

SSD accesses data in large blocks, memory does not

Fitting data in Working Set

Newer is not Always Faster

Upgrading to the
new
Software/Hardware
is not always faster

Test it out

Defaults Change are
often to blame

Upgrades are needed but not seamless

**Major Database Upgrades often
require application changes**

**Having Conversation on Application
Lifecycle is a key**

Character Sets


Performance Impact

Pain to Change

Wrong Character Set can cause Data Loss

Character Sets


MySQL 5.7 utf8mb4_general_ci (default) and latin1



<https://per.co.na/MySQLCharsetImpact>

Less impact In MySQL 8

MySQL 8.0 utf8mb4_0900_ai_ci and latin1



Operational Overhead

Operations Take Time, Cost Money, Cause Overhead

10TB Database Backup ?

Adding The Index to Large Table ?

Distributed Systems

10x+ More Complicated

Better High Availability

Many Failure Scenarios

Test how application performs

Risks of Automation

**Automation is
Must**

**Mistakes can
destroy
database at
scale**

Security

Database is where the most sensitive data tends to live

Shared Devs and Ops Responsibility

Law of Gravity

**Shitty Application at
scale will bring down any
Database**

What Else

What Would you Add ?

Check Out <http://per.co.na/careers>



**BECOME A
PERCONA SUPERHERO**

WE'RE HIRING

**CONTACT
careers@percona.com**

Thank You!

Twitter: @percona @peterzaitsev
