


- The inherent problems of scaling relational databases
- Scaling options
- Sharding And Its Benefits



The inherent problems of scaling relational databases

- Can't scale stateful systems
 - Databases are stateful in nature
- ACID is just too much overhead
 - More transactions, harder to isolate
 - Longer transactions, harder to isolate



Scaling Options

- Tuning
- Hardware upscale
- Partitioning
- New MySQL distribution
- Read/Write Split

- There are many ways to tune your database
- Allot of data online, check out this post
 - <http://forge.mysql.com/wiki/Top10SQLPerformanceTips>

- **innodb_buffer_pool_size**
 - Holds the data and indexes of tables in memory.
 - Bigger buffer results in faster row lookups.
 - The bigger the better.
 - Default – 8M
- **Query Cache**
 - Keeps the result of queries in memory until they are invalidated by writes.
 - query_cache_size
 - total size of memory available to query caching
 - query_cache_limit
 - the maximum number of kilobytes one query may be in order to be cached.
 - query_cache_size = 128MB
 - query_cache_limit = 4MB


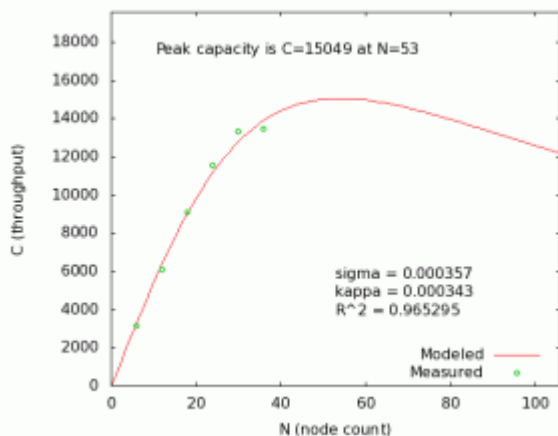
Pros	Cons
May result in major performance improvements	Doesn't scale. No matter how well the tuning is performed, it will reach a limit defined by machine capabilities.
Doesn't require application changes	

- If you write lousy SQL code, you'll get lousy performance
 - Java gurus are not SQL gurus
 - Your ORM code does not know how to write good SQL code
- What will happen when executing
 - `SELECT * FROM`
- Tuning your SQL commands is tedious but very rewarding

- Here are just some examples:
 - Use EXPLAIN to profile the query execution plan
 - Use DISTINCT – not GROUP BY
 - Don't use an indexed column with a function
 - Try not to use a non deterministic functions in where clause

Pros	Cons
May result in major performance improvements	Requires code modifications.
	Doesn't scale. No matter how well the tuning is performed, it will reach a limit defined by machine capabilities.

- Usually DB gets the strongest servers
- However – there is a limit to how much performance gains you can get from increasing hardware
- Some data:



<http://www.mysqlperformanceblog.com/2011/01/26/modeling-innodb-scalability-on-multi-core-servers/>

Pros	Cons
May result in major performance improvements	Scaling is limited
Transparent	Might be expensive
Easy	

- Solid State Drive
 - Better latency and access time than regular HDD
 - Cost more per GB (but prices are dropping)
- Vadim Tkachenko from Percona gave a great lecture on SSD at MySQL Conf 2011
 - (see slides at <http://en.oreilly.com/mysql2011/public/schedule/detail/17117>)
 - Claims you can expect up to X7 performance from SSD

Pros	Cons
May result in major performance improvements	Expensive
Transparent	Still limited scalability

- Partitioning was introduced to MySQL at version 5.1.
- It is a way to split tables across multiple files, a technique that proved very useful for improving database performance.
- Benefits:
 - Helps fit indexes in RAM
 - Faster query/insert
 - Instant delete

- See excellent presentation by Giuseppe Maxia from 2010
 - <http://www.slideshare.net/datacharmer/partitions-performance-with-mysql-51-and-55>

Engine	6 month range query
InnoDB	4min 30s
MyISAM	25.03s
InnoDB partitions	13.19s
MyISAM partiotions	4.45s

Pros	Cons
May result in major performance improvements	MySQL server itself introduces limits. User concurrency, transaction chains, isolation, are still bottlenecked by the single MySQL that owns all partitions
Mostly transparent to the application	

- There are many MySQL drop-in replacements
- Are MySQL, but tuned differently, different extensions
- Leading examples
 - PerconaServer
 - MariaDB

Pros	Cons
Provide performance improvements	Still limited scalability
Transparent	

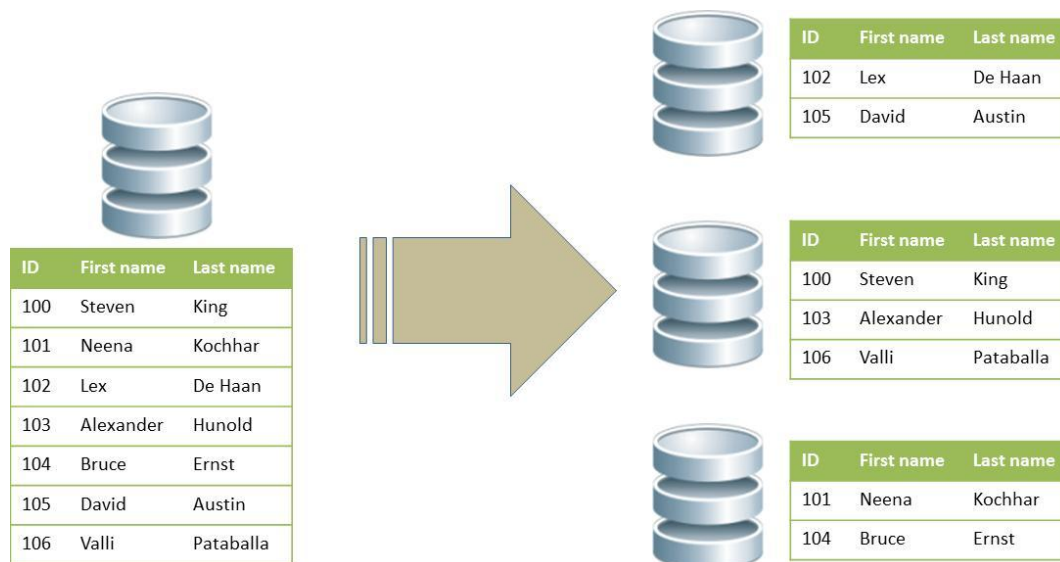
- Write to MySQL master, read from 1 (or more) slaves
- Excellent read scaling
- Many issues:
 - Since replication is a-synchronous – read might not be up to date
 - Transactions create stickiness
 - Code changes

Pros	Cons
Provides performance improvements	Requires code changes
Scale out the database	Good for scaling reads, not writes
	Since replication is asynchronous, some reads might get data that is not up to date.



Sharding And Its Benefits


- With Sharding, a database's data is split into multiple databases, each storing a subset of the data. Consider the following diagram:



- A single database table is **split into multiple databases**, each storing some of the original table rows.
- The application needs to be aware of the shards, and keep track of which data subset is stored in which database.
- **Sharding is the next step of partitioning**, where partitioning over different files in the same server, evolves into harnessing the power of several different servers for our partitions.
- **Sharding has proved itself as the ultimate scaling solution for MySQL**, and is used in most popular MySQL-backed websites. As a matter of fact, it's the only way to scale reads as well as writes. However, the development effort for building shards supporting code is huge, which is why most companies use sharding only as a "Last Resort" (see a list of problems here - <http://www.scalebase.com/don%E2%80%99t-ever-ever-write-your-own-sharding-code/>)



ID	First name	Last name
100	Steven	King
101	Neena	Kochhar
102	Lex	De Haan
103	Alexander	Hunold
104	Bruce	Ernst
105	David	Austin
106	Valli	Pataballa



ID	First name	Last name
102	Lex	De Haan
105	David	Austin



ID	First name	Last name
100	Steven	King
103	Alexander	Hunold
106	Valli	Pataballa



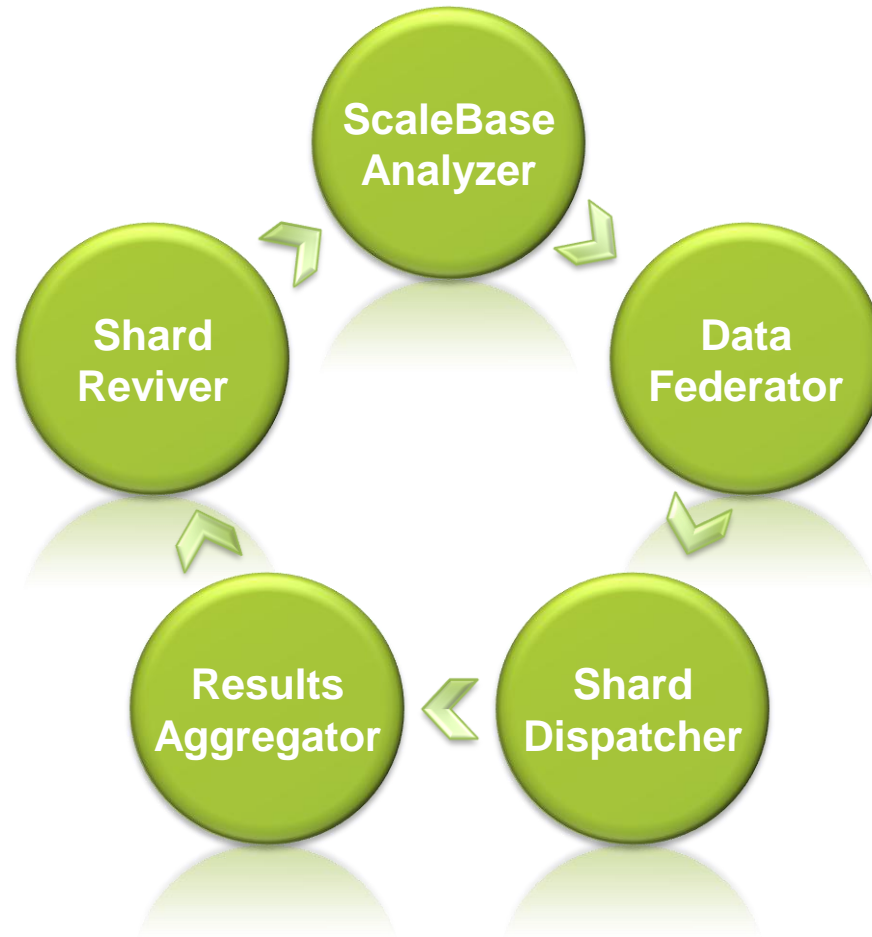
ID	First name	Last name
101	Neena	Kochhar
104	Bruce	Ernst


- Performance
 - Since each database is smaller in size: more caching, smaller indexes
 - Since each database gets less hits: less isolation
- Throughput
 - More simultaneous open connections
 - More simultaneous queries and DML
- Scale Out, not Scale Up
 - Return of investment

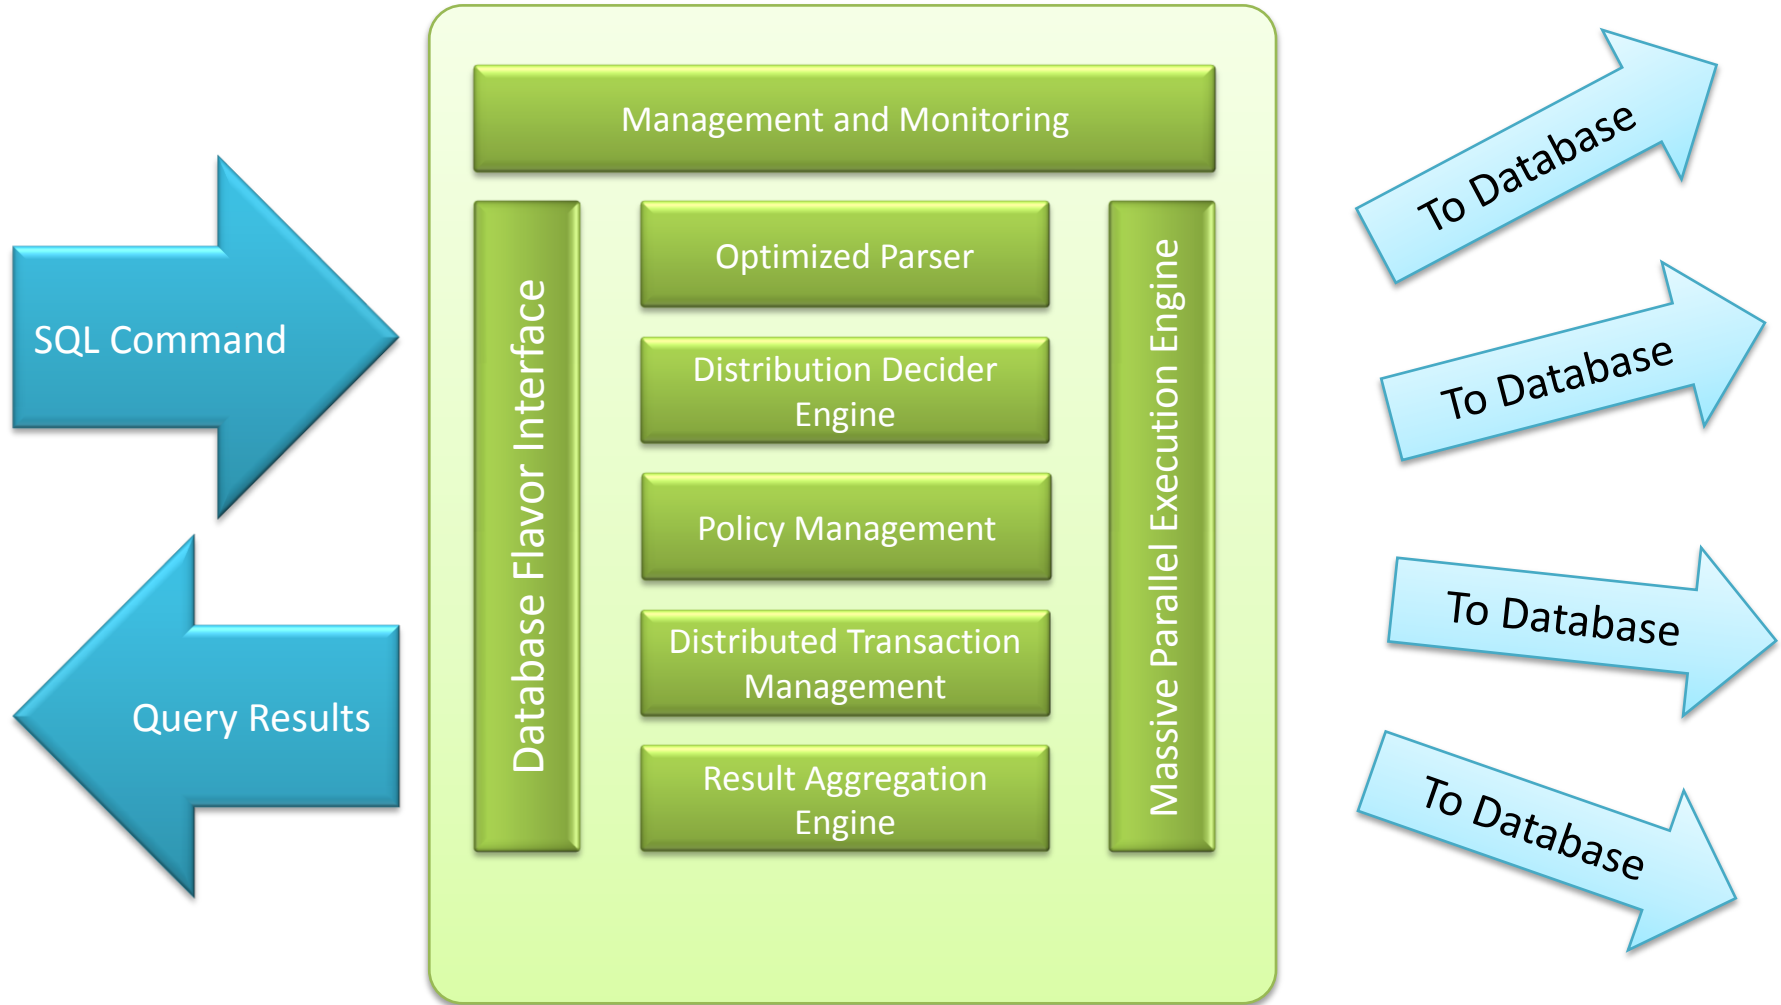
- Complexity
 - In development
 - In administration

Pros	Cons
Provides performance improvements. The only way to really improve read and also writes	Requires substantial development effort
Scale out the database	Results in major management costs, as DBAs, BI experts and System Administrators need to be familiar with the Sharding solution.

- Transparent Sharding
- All the sharding benefits – minus the downsides







[Q [& A]]

Thank You!