

---

# Overview of MySQL Column Databases

*Robin Schumacher*  
*VP Products, Calpont*  
*March 2010*

# Agenda

---

- Overview of MySQL Row vs. Column Engines
- MySQL Column Database Options
- Benchmark Comparisons
- Conclusions

---

# Overview of Row vs. Column Engines

# MySQL Column Engines 'feel' like Normal MySQL



```
root@srvprodtest1:/usr/local/Calpont/bin

mysql> use tpch;
Database changed
mysql> show tables;
+-----+
| Tables_in_tpch |
+-----+
| customer        |
| lineitem        |
| mytable         |
| nation         |
| orders         |
| part           |
| partsupp       |
| region         |
| rms            |
| rms2           |
| supplier       |
| t1             |
| t2            |
+-----+
13 rows in set (0.05 sec)

mysql> select count(*) from lineitem;
+-----+
| count(*) |
+-----+
| 6001215 |
+-----+
1 row in set (0.31 sec)

mysql> desc t1;
+-----+-----+-----+-----+-----+-----+
| Field | Type          | Null | Key | Default | Extra |
+-----+-----+-----+-----+-----+-----+
| c1    | int(11)       | YES  |     | NULL    |       |
| c2    | varchar(100) | YES  |     | NULL    |       |
+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql>
```

# MySQL Column Engines 'feel' like Normal MySQL

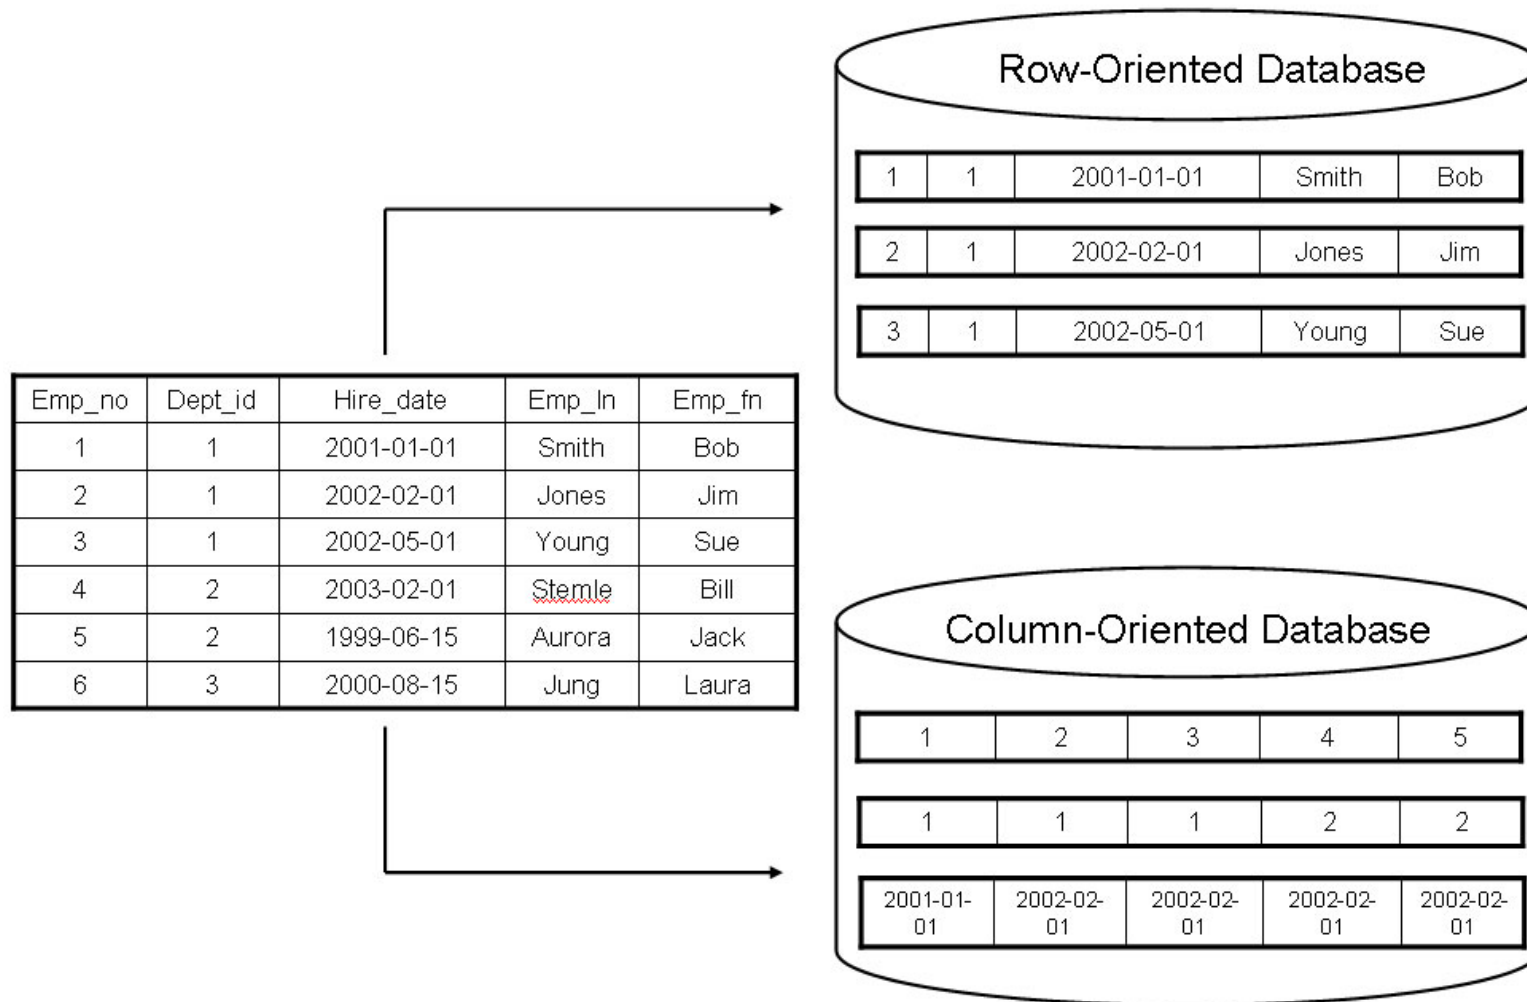


The screenshot shows the MySQL website interface. At the top, there is a navigation bar with links for 'Home', 'Products', 'Services', 'Customers', 'Why MySQL?', 'News & Events', 'About MySQL', and 'How to Buy'. Below this, there is a 'GET STARTED' section with a list of links: 'MySQL Enterprise 30-day Trial', 'MySQL Enterprise', 'Free Web Seminars', 'ISVs and OEMs', 'Test Drive MySQL Query Analyzer', and 'MySQL Training'. To the right of this list is a large banner for Symantec, featuring a stack of stones and the text 'symantec Relies on MySQL for Rock-Solid Quality' with a 'Learn More »' link.


- Standard MySQL connectors can be used – JDBC, ODBC, .NET, C++, etc.
- MySQL GUI tools support connection to/management of MySQL column databases
- BI and ETL tools that use normal MySQL connectors should work fine with MySQL column databases
- Installation for software-delivered column databases easy and quick (20MB download, RPM/TAR install)
- MySQL bundled in with column databases; no need to download or install MySQL separately. Must use bundled version of MySQL with each column database

# Column vs. Row Orientation

A column-oriented architecture looks the same on the surface, but stores data differently than legacy/row-based databases...



# Why a Column Database?



- Column databases only read the columns needed to satisfy a query vs. full rows
- If you are only selecting a subset of columns from a table and / or are using very wide tables, column DB's are a great choice
- Column databases (most of them...) remove the need for indexing because the column is the index
- Column databases automatically eliminate unnecessary I/O both logically and physically, so they do away with partitioning needs too as well as materialized views, etc.
- As a rule of thumb, column databases provide 5-10x (or more) the query performance of legacy RDBMS's

# Why a Column Database?



*"If you're bringing back all the columns, a column-store database isn't going to perform any better than a row-store DBMS, but analytic applications are typically looking at all rows and only a few columns. When you put that type of application on a column-store DBMS, **it outperforms anything that doesn't take a column-store approach.**"*


- Donald Feinberg, Gartner Group

# MySQL and Data Warehousing



- MySQL DW deployed at 28% of MySQL customers
- Strong DW ecosystem support

## What types of applications are you deploying on MySQL?




- Except for InnoDB and MySQL Cluster (which are not suitable for data warehousing), no MySQL storage engine can fully address more than 4 CPU's/cores
- The MySQL optimizer is widely known for its inefficiencies in subqueries, joins (e.g. no hash joins), unions, etc.
- No parallel SQL, DML, or load capabilities
- Once a database gets past 500GB, the issues noted above tend to become fairly noticeable
- Missing objects such as materialized views/summary tables, bitmap indexes, etc.

# Why Column Databases Now?




What problems will eventually drive you to replace your current primary data warehouse platform?  
(Select nine or fewer.)



Philip Russom, "Next Generation Data Warehouse Platforms", TDWI, 2009

# BI Database Technical Trends




Note: 'Plan to Use' means the responder has already written a check for the technology


# BI Database Technical Trends



Of the components of your data warehouse platform, which do you feel is the most critical to success?



Today, which processing architecture is your data warehouse on? Ideally, which processing architecture would you prefer for your data warehouse platform?



Philip Russom, "Next Generation Data Warehouse Platforms", TDWI, 2009



## IDC - Press Release

### **IDC Finds That a Third Generation of Database Technology Along with Vendors Are Shaking Up the Market**

12 Feb 2010

**FRAMINGHAM, Mass., February 12, 2010** – For database administrators (DBAs) and those who select and manage database management system (DBMS) technology that is based on conventional row-oriented, disk-based systems that drive queries in a linear fashion, a new generation of DBMS technology is sending a simple, clear message - 'Everything you know is wrong'.

Recent [IDC](#) research shows that at the current rate of development and adoption, it is likely that within five years:

- Most data warehouses will be stored in a columnar fashion
- Most OLTP databases will either be augmented by an in-memory database (IMDB) or reside entirely in memory
- Most large-scale database servers will achieve horizontal scalability through clustering
- Many data collection and reporting problems will be solved with databases that have no formal schema at all.

"Many of these new systems encourage you to forget disk-based partitioning schemes, indexing strategies, and buffer management, and embrace a world of large-memory models, many processors with many cores, clustered servers, and highly compressed columnwise storage," says [Carl Olofson](#), research vice president for [Database Management and Data Integration](#) software research at IDC. "These technologies are no longer on the fringes, used only to solve exotic data warehouse or other intensive number-crunching problems. They are becoming mainstream, and are offered by a wide variety of vendors – some new, some quite familiar – to solve common, real-world data management problems."

## Why Not?

- If you routinely have SELECT \* queries or queries that request the majority of columns in a table
- If you constantly are doing lots of singleton inserts and deletes. As these are row-based operations they will normally run somewhat slower on a column DB than a row-oriented DB (more block touches are needed). Updates tend to run OK as they are a column operation
- If you want to do pure OLTP work. Some column DB's are transactional (so data integrity is ensured), but they are not suited for straight OLTP work
- If you have a small database: such a DB eclipses the benefit column databases offer over row DB's

# MySQL Column Database Options

---

# Main MySQL Column Databases



## INFOBRIGHT

- Strengths: Great data compression, good metadata layer, overcomes many MySQL optimizer issues, easy to install and maintain
- Weaknesses: Single-server centric, no MPP capability, a query is tied to one CPU only (so no SMP or MPP – scale up or up)



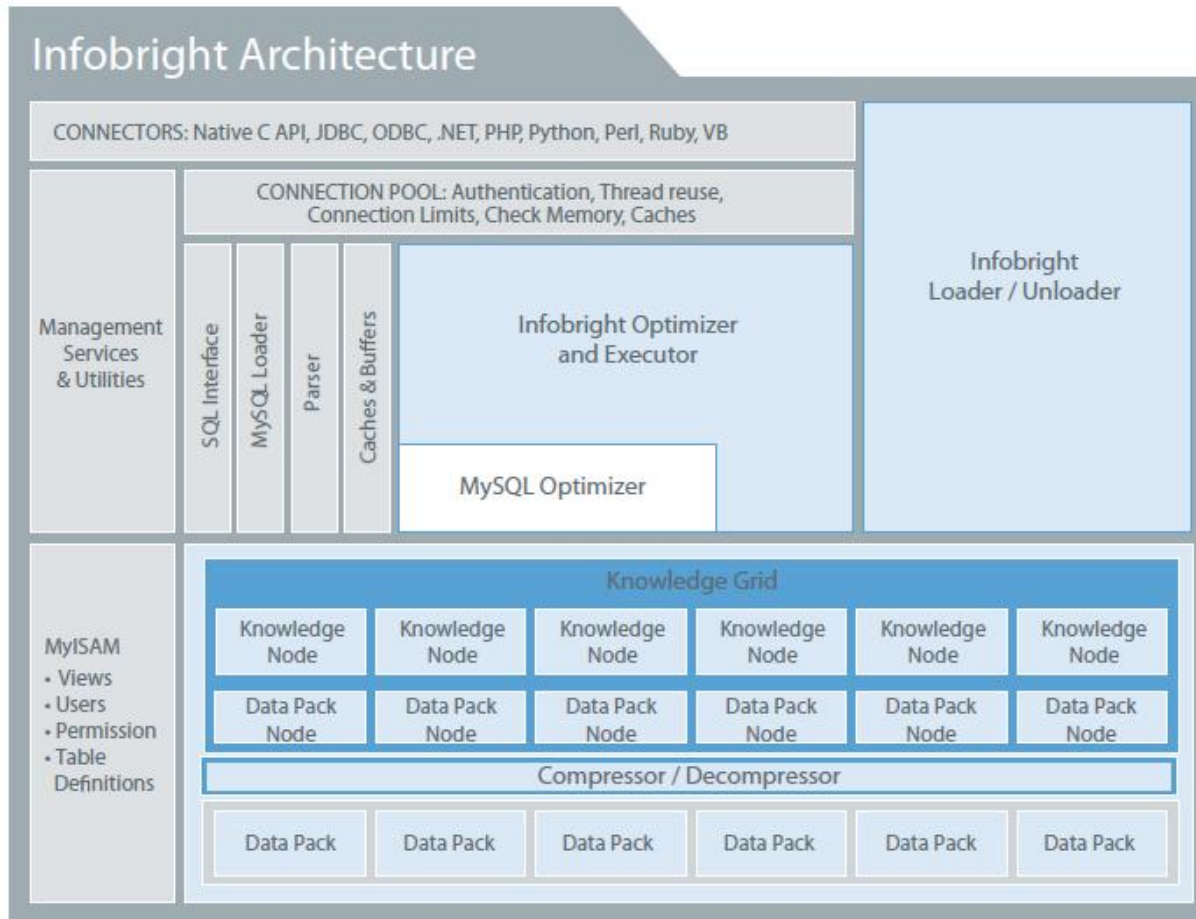
## KICKFIRE

- Strengths: Strong feature set, good data compression, overcomes many MySQL optimizer issues, good TPC-H benchmarks
- Weaknesses: Hardware/appliance approach, no MPP capability



- Strengths: Strong feature set, Scale up/out (SMP and MPP capable), overcomes many MySQL optimizer issues, good comparison benchmarks
- Weaknesses: SQL syntax support limitations, no physical compression yet

## INFOBRIGHT




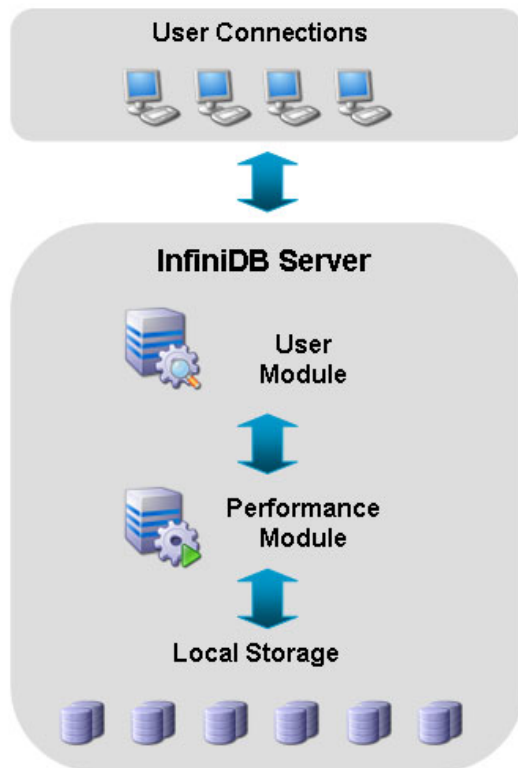


| Model                          | 2200                      | 2300                      | 2400                       | 3200   |
|--------------------------------|---------------------------|---------------------------|----------------------------|--|
| Rack Space                     | 2U                        | 2U                        | 3U                         | 5U   |
| Raw Disk                       | 8 2.5" 146 GB SAS drives  | 8 2.5" 146 GB SAS drives  | 8 2.5" 300 GB SAS drives   | 8 2.5" 300 GB SAS drives + 12 3.5" 1 TB SAS drives |
| User Data Range*               | 750 GB - 1.5 TB           | 750 GB - 1.5 TB           | 1.5 - 3 TB                 | 3 - 6 TB   |
| Query Processing Module Memory | 32 GB                     | 64 GB                     | 128 GB                     | 256 GB   |
| Power                          | 600 W                     | 600 W                     | 700 W                      | 900 W  |
| Weight                         | 62 lbs (28.1kg)           | 62 lbs (28.1kg)           | 77 lbs (34.9kg)            | 130 lbs (59kg)                                     |
| Mounting                       | 2 or 4 post               | 2 or 4 post               | 2 or 4 post                | 4 post   |
| Operating Temperature          | 0 - 40° C                 | 0 - 40° C                 | 0 - 40° C                  | 5 - 35° C  |
| Dimensions                     | 16.33"W x 25.76"D x 3.5"H | 16.33"W x 25.76"D x 3.5"H | 16.33"W x 25.76"D x 5.25"H | 17.33"W x 25.76"D x 8.75"H                         |

# Calpont's InfiniDB



## Scale Up (SMP) and Scale Out (MPP)



# MySQL Column Database Feature Comparison



|  | Infobright<br>Community | Infobright<br>Enterprise | Kickfire   | InfiniDB<br>Community | InfiniDB<br>Enterprise |
|--|-------------------------|--------------------------|------------|-----------------------|------------------------|
| License  | GPL                     | Commercial               | Commercial | GPL                   | Commercial             |
| Software/hardware (appliance) solutions                            | Software                | Software                 | Hardware   | Software              | Software               |
| Multi-threaded engine (queries/writes uses all CPU's/cores on box) | No                      | No                       | Yes        | Yes                   | Yes                    |
| Logical data compression   | Yes                     | Yes                      | Yes        | Yes                   | Yes                    |
| Physical column compression  | Yes                     | Yes                      | Yes        | No                    | No                     |
| High-Speed bulk loader w/ no blocking queries while loading        | No                      | Yes                      | Yes        | Yes                   | Yes                    |
| Crash-recovery   | No                      | No                       | Yes        | Yes                   | Yes                    |
| Transaction support (ACID compliant w/ DML)                        | No                      | Yes                      | Yes        | Yes                   | Yes                    |
| INSERT/UPDATE/DELETE (DML) support                                 | No                      | Yes                      | Yes        | Yes                   | Yes                    |
| Support for Hash Joins   | No                      | No                       | No         | Yes                   | Yes                    |
| No indexing necessary  | Yes                     | Yes                      | No         | Yes                   | Yes                    |
| Primary and Foreign Key support                                    | No                      | No                       | Yes        | No                    | No                     |
| MVCC support – snapshot read (readers don't block writers)         | No                      | No                       | No         | Yes                   | Yes                    |
| Alter Table with online add column capability                      | No                      | No                       | No         | Yes                   | Yes                    |
| Concurrency/Locking support  | Table                   | Table                    | Table      | Page                  | Page                   |
| Massive parallel processing capable (MPP – multi-node parallelism) | No                      | No                       | No         | No                    | Yes                    |
| Non-MySQL supplied SQL tracing/diagnostics                         | No                      | No                       | No         | Yes                   | Yes                    |
| Replication support  | No                      | No                       | No         | No                    | No                     |
| Terabyte database capable  | Yes                     | Yes                      | Yes        | Yes                   | Yes                    |

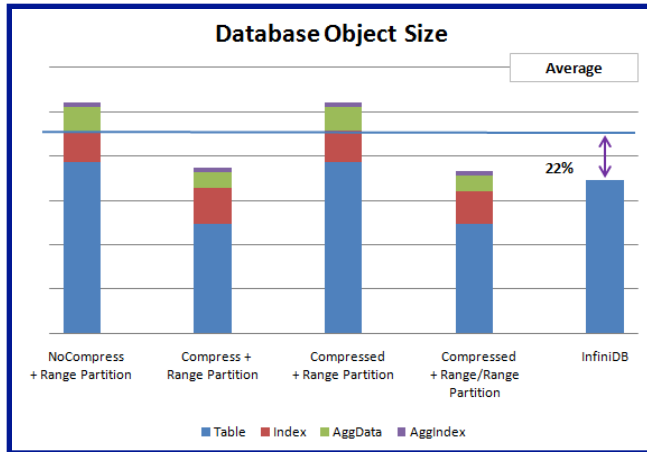
# Benchmark Examples

---

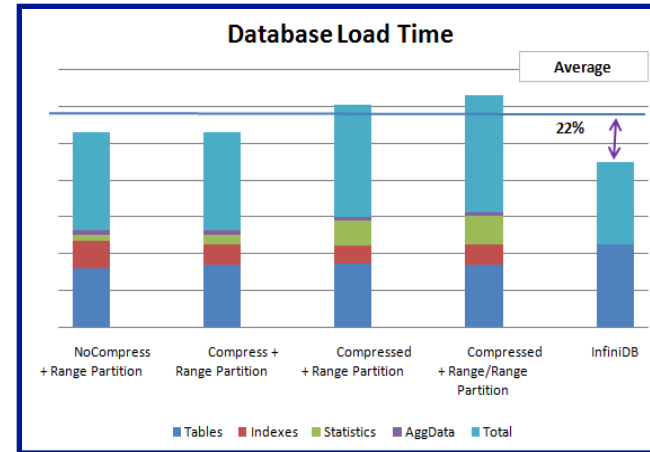
# InfiniDB vs. "Leading" Row Database



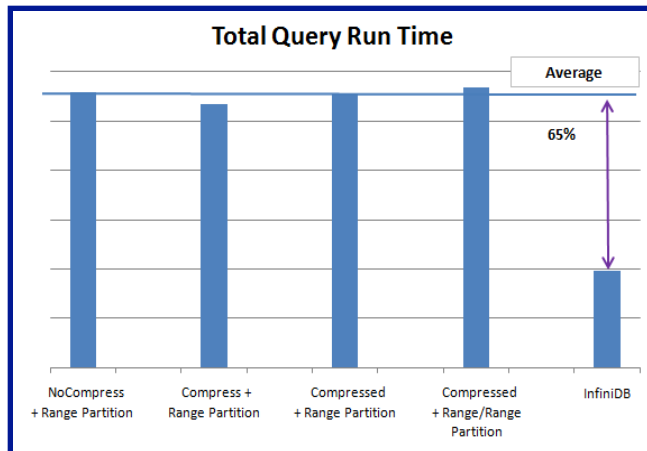
**InfiniDB takes up 22% less space**




**InfiniDB loaded data 22% faster**



**InfiniDB total query times were 65% less**




**InfiniDB average query times were 59% less**



Notice not only are the queries faster, but also more predictable

\* Tests run on standalone machine: 16 CPU, 16GB RAM, CentOS 5.4 with 2TB of raw data

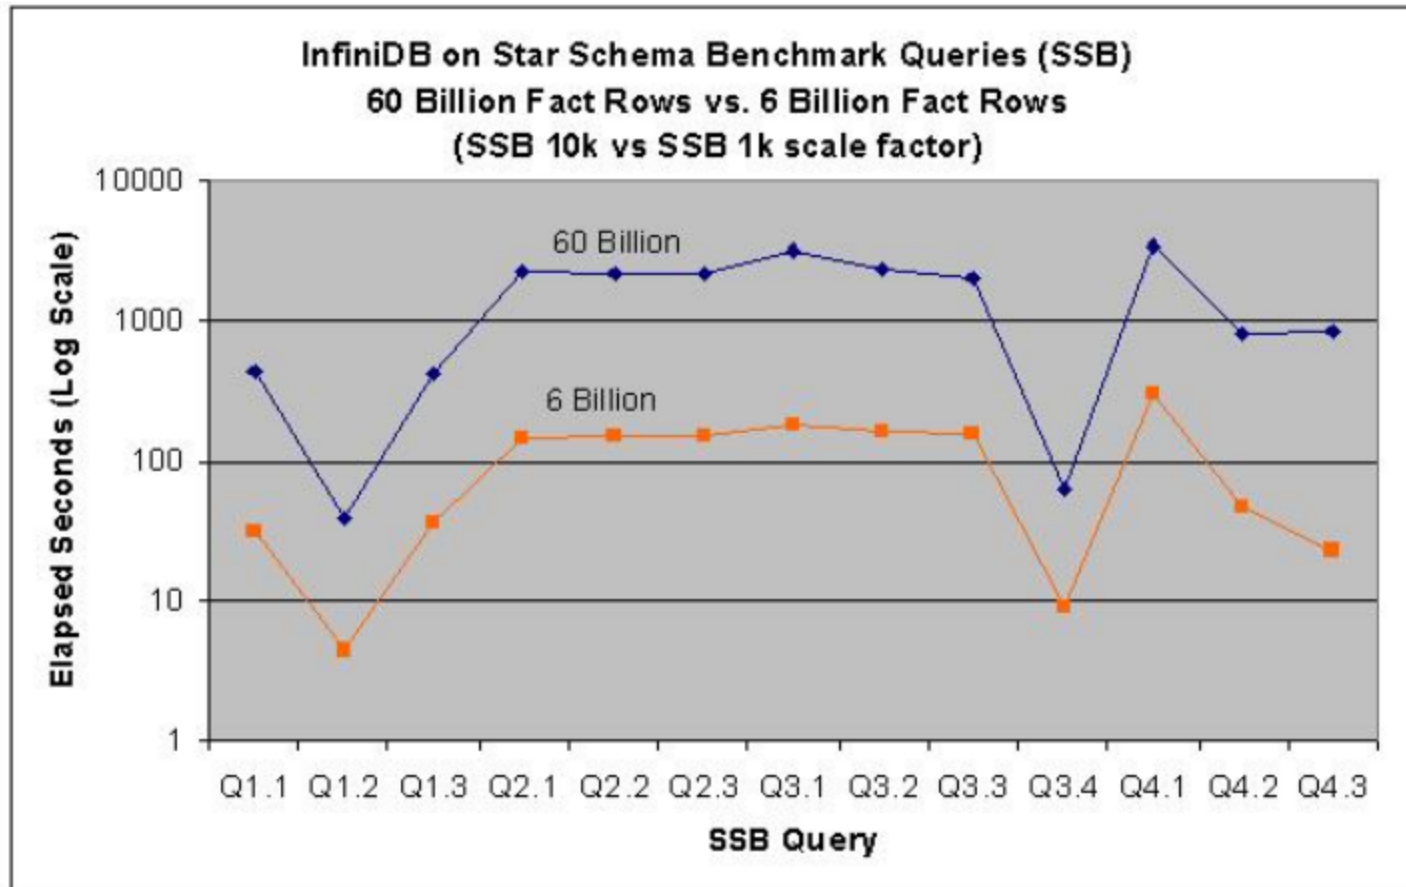
## SSB queries



610 GB of raw data; 8 Core Machine

<http://www.mysqlperformanceblog.com/2010/01/07/star-schema-bechmark-infobright-infinidb-and-luciddb/>

# Internal MPP Stress Test using SSB Benchmark



Six Performance Modules against data size of Percona tests and then 10x the size. Shows predictable performance, MPP linear speed gains over standalone implementation, and capabilities that other open source databases don't have.

<http://www.infinidb.org/infinidb-blog/mysql-parallel-query-processing-of-60-billion-rows-via-infinidb.html>

- Infobright:
  - [www.infobright.org](http://www.infobright.org)
  - [www.infobright.com](http://www.infobright.com)
- Kickfire:
  - [www.kickfire.com](http://www.kickfire.com)
- Calpont/InfiniDB:
  - [www.infinidb.org](http://www.infinidb.org)
  - [www.calpont.com](http://www.calpont.com)



---

# Overview of MySQL Column Databases

*Robin Schumacher*  
*VP Products, Calpont*  
*Thanks!*