FreeWheel®

SERVING THE VIDEO REVOLUTION.

Big Data Migration Practice using Column Database

Sean FreeWheel

2011/12/03

Agenda

- About FreeWheel
- Challenge of Growth
- Column-based Database
- The Migration
- The Performance
- Thoughts

FreeWheel is focused on media and entertainment

- We're on a mission to help the largest media companies make the most money possible from their content
  - Founded 2007
  - 130+ industry veterans based in: Silicon Valley, New York, Beijing, London
  - Backed by Steamboat, Turner, Battery Ventures, and Foundation Capital
- Proud to call these companies clients:

ComScore: Top 10 U.S. Online Video Content Properties by Unique Viewers (Sep, 2011)

Top U.S. Online Video Content Properties Ranked by Unique Video Viewers
Sep-11
Total U.S. – Home/Work/University Locations
Content Videos Only (Ad Videos Not Included)
Source: comScore Video Metrix

Property	Total Unique Viewers (000)	Videos (000)*	Minutes per Viewer
Total Internet : Total Audience	*181,915*	*39,813,781*	*1,167.90*
Google Sites	161,406	18,609,393	378
VEVO	57,348	748,163	60.2
Microsoft Sites	53,972	735,911	39.3
Viacom Digital	53,424	633,598	53.8
Facebook.com	49,891	252,521	16.1
Yahoo! Sites	46,104	526,745	49.4
AOL, Inc.	41,850	408,157	55.9
NBC Universal	29,772	95,518	15.3
Turner Digital	28,594	240,980	33.2
Hulu	27,051	642,405	180.3

FW Client FW Partner

FreeWheel Delivers 1 in 3 US Video Ads*

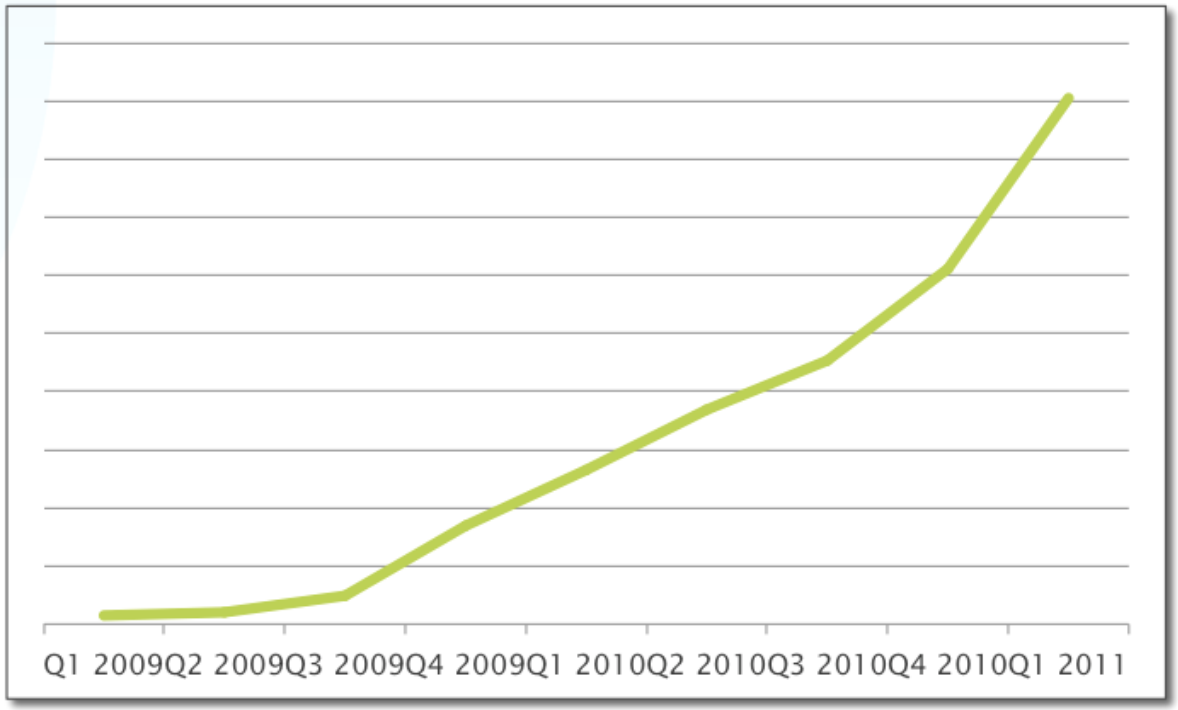

Partners, Integrations and Distributors

Content Distributors	20+
Player Platform and Content Management Systems	10+
Ad Networks	20+
Advertising Tools/Technologies/Servers	10+
Mobile Runtime Environments	5
Device types	10+

FreeWheel has a central role and holistic view across apps, partners, and devices

*Derived from monthly FreeWheel data and comScore Video Metrix data

Evolution

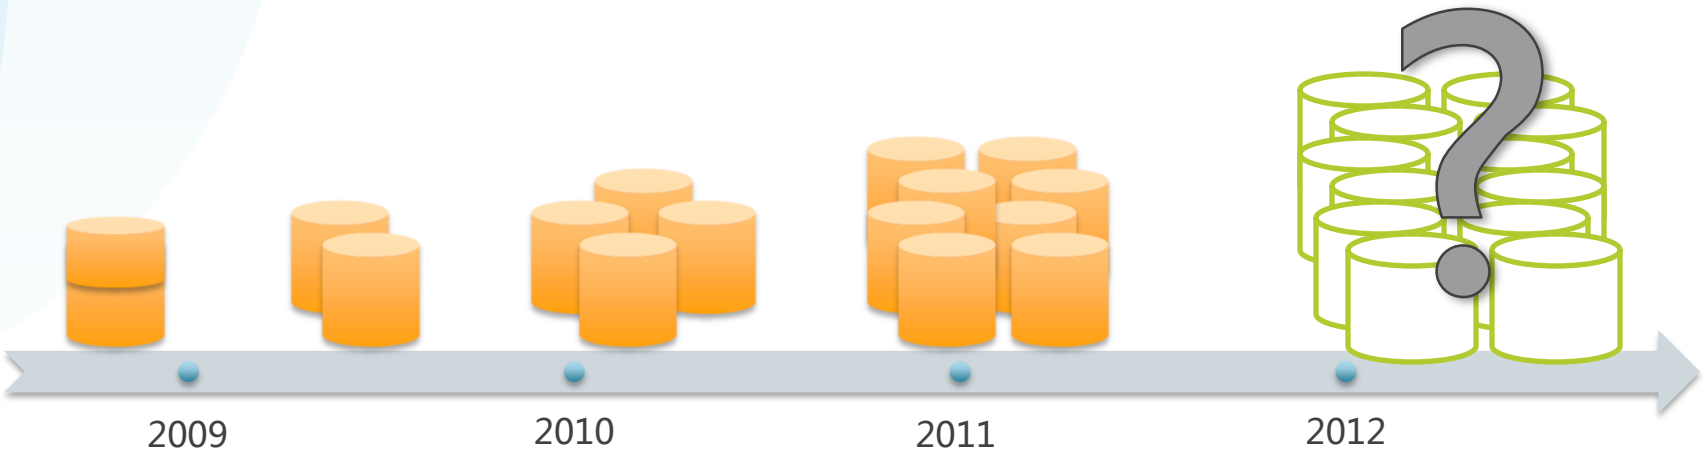

- Rollup/Pivoting
- Partitioning by event_date
- Sharding by clients
- Innodb Compression
- Summary Tables

Problems

- Some #, Mid 2011
  - Keep 19 Months data
  - 4 Shards, 8 Servers
  - ~4TB as Total
  - ~50MM Rows (~60GB)/Day
- Query/Load Performance
  - Query hours to return
  - Index never satisfy
- Storage/Server Cost
- Maintenance Cost
  - Schema change
  - Archive/Backup

Why Column-based Database

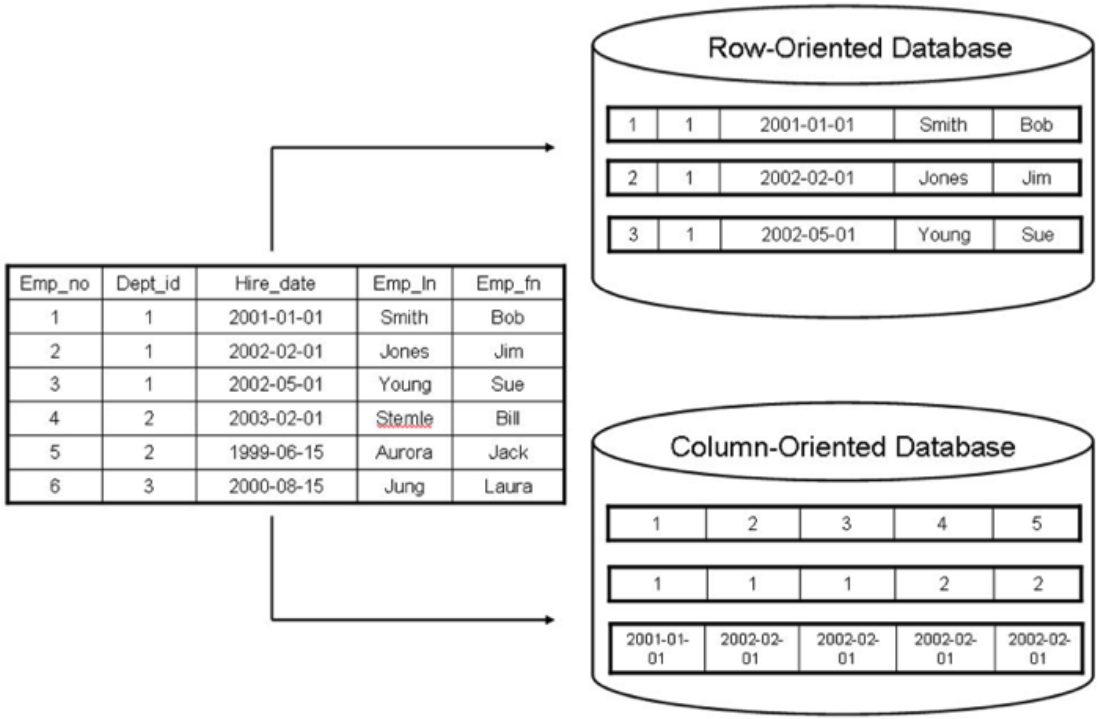

Emp_no	Dept_id	Hire_date	Emp_ln	Emp_fn
1	1	2001-01-01	Smith	Bob
2	1	2002-02-01	Jones	Jim
3	1	2002-05-01	Young	Sue
4	2	2003-02-01	Stemle	Bill
5	2	1999-06-15	Aurora	Jack
6	3	2000-08-15	Jung	Laura

- Read needed column VS full rows
- Compression with same data (type)
- No more index, better statistics
- Reduce unnecessary I/O

Why Infobright

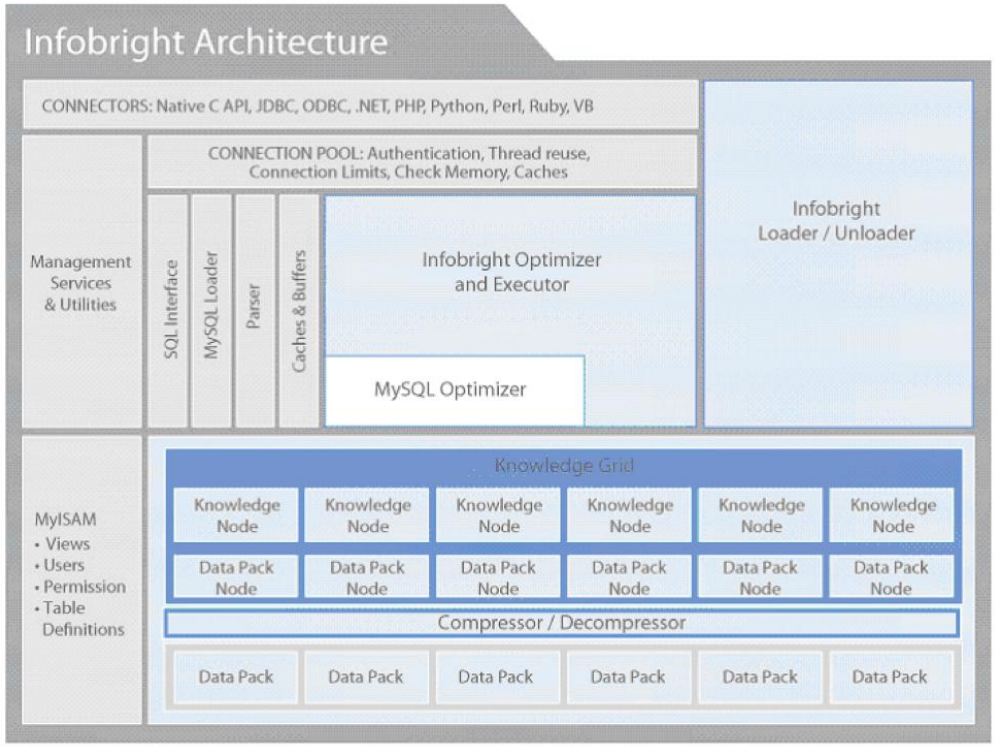

- High compression ratio (at least 10:1)
- Fast response times for ad hoc analytic queries
- *Like* MySQL

Evaluation

- One Fact table, 123GB, 344MM rows

	ICE 3.5.2	MySQL 5.1.49 w Innodb
Data Compression	2.6GB, **1/47**	241GB, 1.95
Load Performance	03:45:00	08:52:00
Query Performance	00:10:35	00:14:58

- Differences
  - Lack of DML
  - Data Type
  - Etc.

The Migration – challenges!

- Migration requirement
  - 7x24 online service, global clients
  - QoS, SLA
  - Data: Consistency, Correctness
  - Service: Availability
- Planning: SMOOTH
  - Parallel run w/ production data
  - Reply request using query log
    - Cover all the query patterns
    - Correctness
    - Performance analysis
  - Cross Validation

Issues Found and Resolved

- Not fully support UTF-8
- Not auto truncate when load data
- Not support store function with parameters
- Bugs* or Limitations
  - UNION ALL in sub query makes mysqld crash
  - Query with LEFT JOIN and a non-join condition will hang there
  - A query with DISTINCT make the server restart
  - Not support query two alias of same column from sub query (from VIEW)
  - Both operator and operand must be string data types for LIKE comparison
  - LEFT JOIN with OR in condition doesn't work
  - JOIN changes the order of rows (thus lost the order in sub query)
  - No order after GROUP BY
  - HAVING support only GROUP BY key or aggregate fun metrics
  - Doesn't support IS NULL in LEFT JOIN (only for 3.5.2)
- → Change query or workflow to resolve

*Some bugs already resolved in later version

Real Query Replay

- Much faster, ~50% queries cost 1/10 time

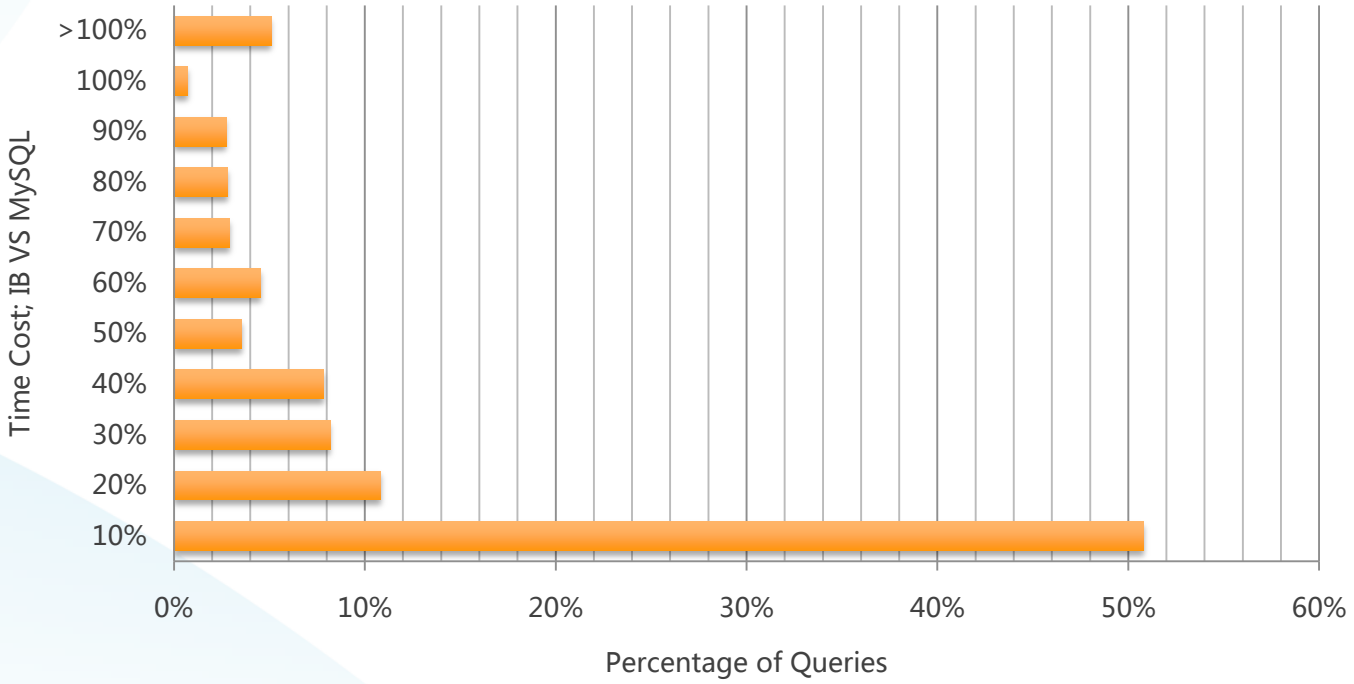

The Migration

- Gradually replace the production

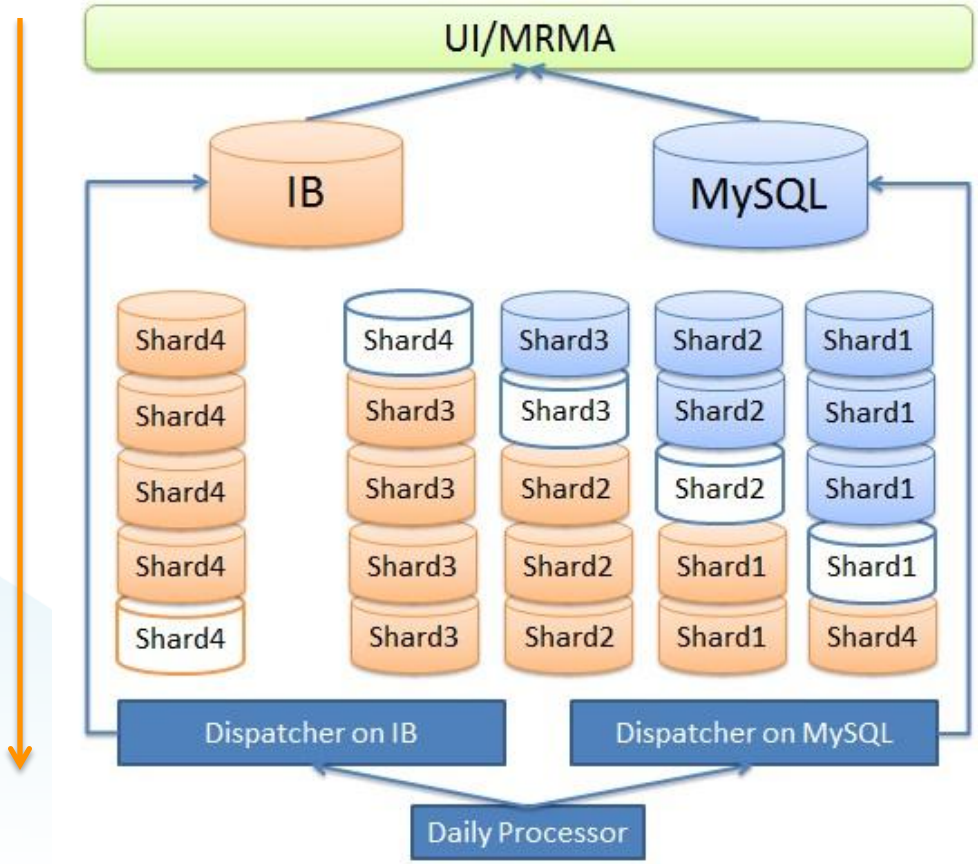

Compression

- Avg. ratio 1/40 (1/30 – 1/70)

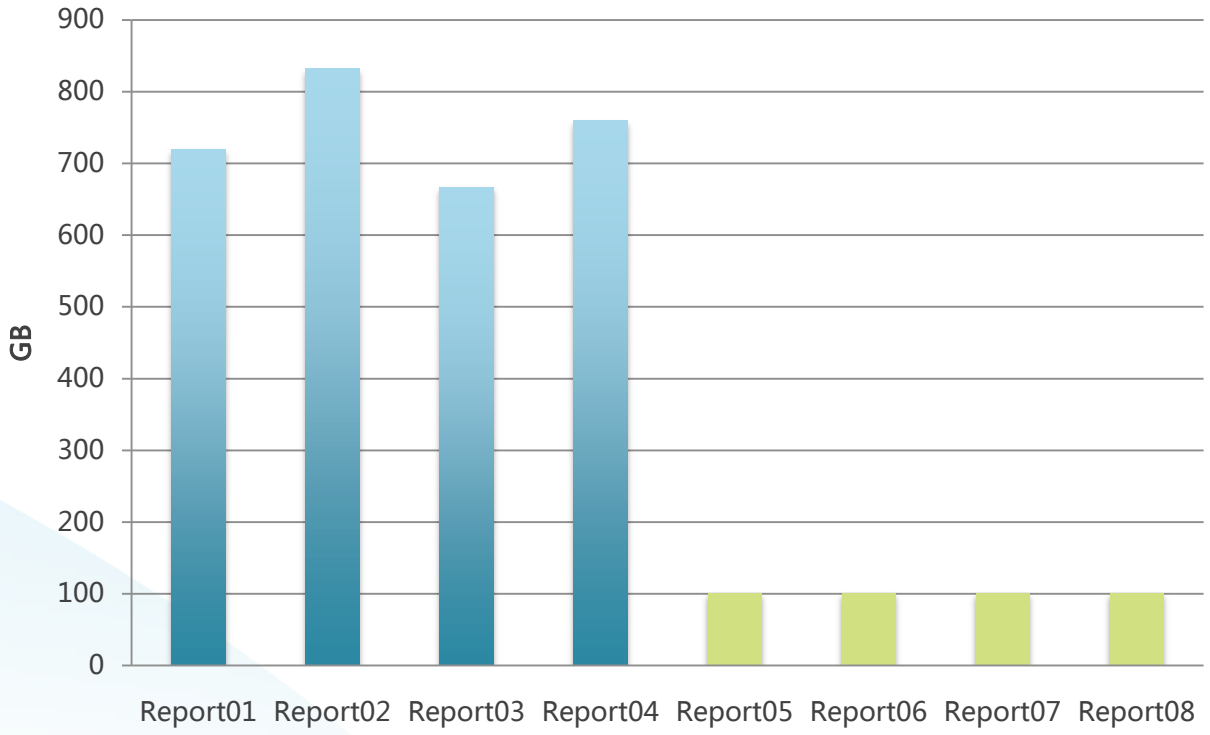

Note: Report05-08 migrated to Infobight

Load Performance

- Remains constant as the size of the database grows

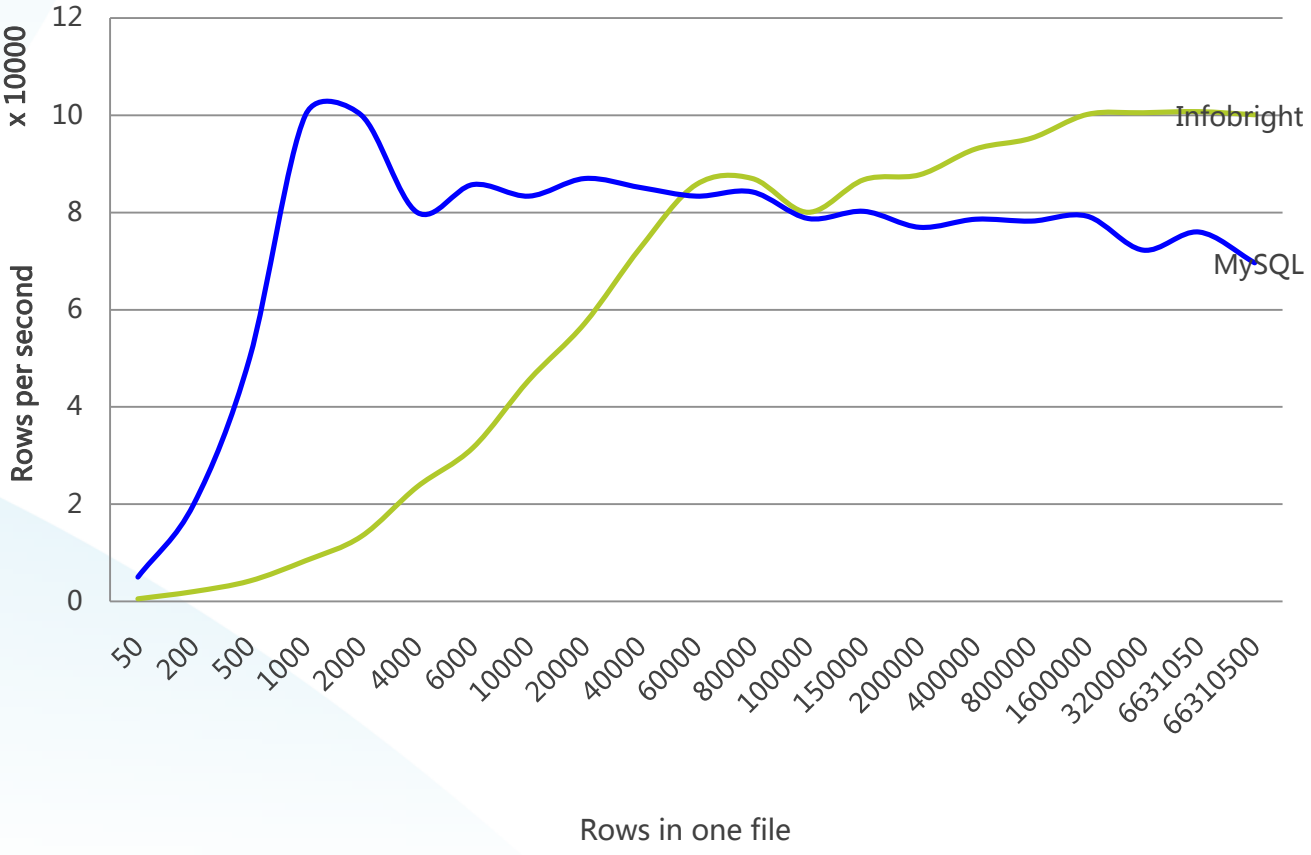

Query performance

- Return more stably

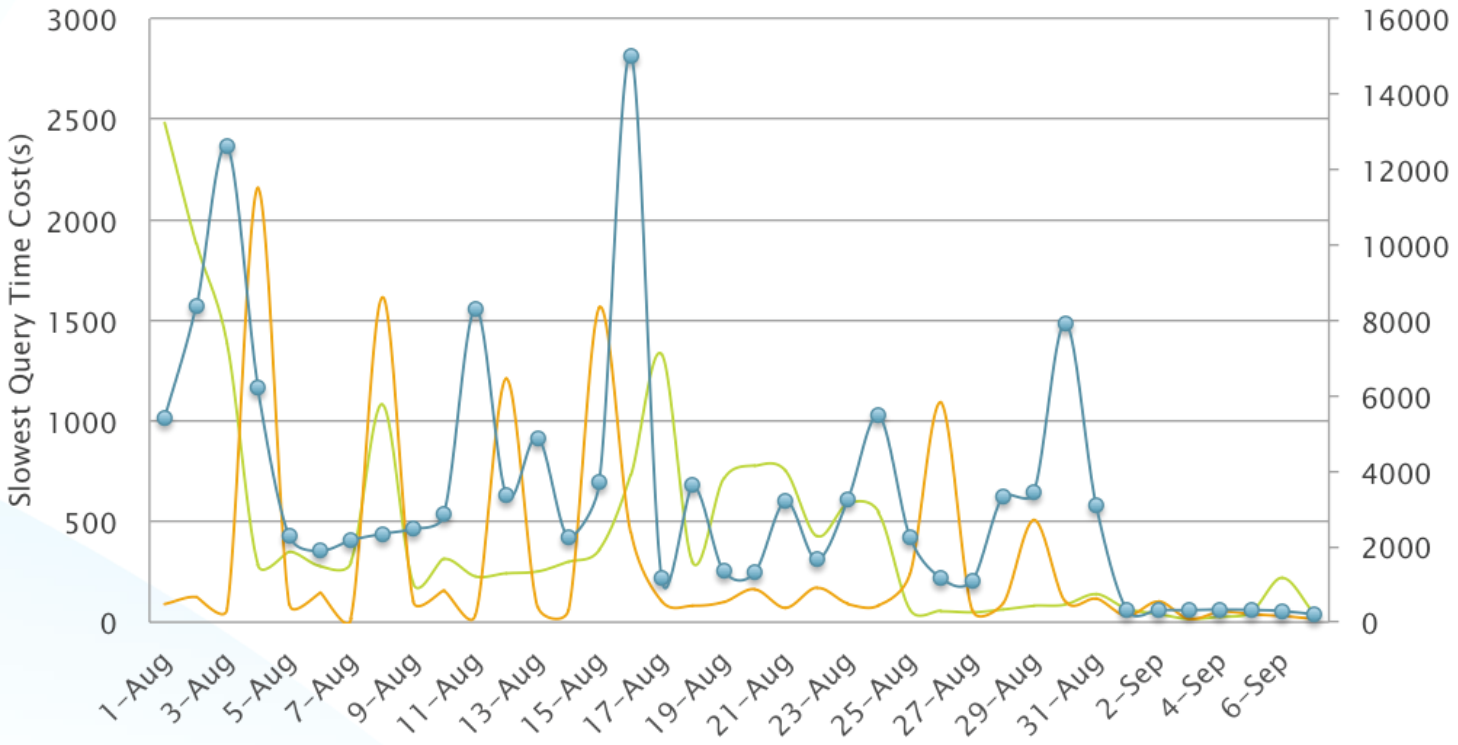

Thoughts

- Ready for change
  - Right for the biz need
  - Keep simple and fast
  - Use production data
  - Parallel running: rollback
  - → Smooth Migration

FreeWheel®

SERVING THE VIDEO REVOLUTION.

Q & A

Thank you!