



PERCONA  
Performance Consulting Experts

# Breaking the Boundaries

Percona Performance  
Conference

April 22-23 2009

Santa Clara, CA

by Peter Zaitsev, Percona Inc

# World is Different now

---

- Different Scale
- Different Economics
- Different Hardware
- Different Technology needs

# Difference Scale

- Internet Applications often larger than Enterprise
  - Performing more transactions per second
  - Having larger data set
- Visa does 7000 transactions/sec at peak
  - Compare this to interactions with FaceBook
  - How many documents Google crawls a second ?
- Most internal Enterprise applications designed for thousands/tens thousands of users
  - Internet applications may have millions and tens of millions of users.

# Different Economics

- Internet Applications have a lot of data and high transaction rate
  - But little money which comes from this activity
- Compare revenue Visa has per transaction or Airline per booking to Facebook per page view or Youtube per movie watched

# Power Requirements

- More important on the high end
  - If you rent 1-2 servers you do not care how much power they take.
- Power/Power Density not the space is leading limiting factor for older centers
  - Many data centers are half empty due to power budget
- Power is significant contributor to TCO
- Internet applications need large scale systems requiring a lot of power.
- Performance per Watt/Transactions per Watt metrics

# High Load with Little money

- Become Creative in your solutions
- Optimize Performance
- Use Open Source Software
- Use Commodity Architectures
- Sacrifice on requirements
- Look for simplicity
  - Development and operations cost a lot
- Wikipedia, YouTube, Craigslist
  - Small teams; great resources.

# Scalability and Efficiency

- Scalable solution may not be efficient one
  - Efficiency aspect is often forgotten
- You can run complex tasks with desired response time
  - If you supply enough servers
- But how this matches to performance you could reach ? To performance of single node task
- You can find “scalable” solutions which are order of magnitude different in their efficiency.

# Different Hardware

- The exponential performance growth of single core is no more.
  - Though now we're getting more and more “cores” and “threads” per chip
- How many OLTP applications “scale” to 256 threads
  - How many can effectively use such hardware to run a single job ?
- Cheap hardware allows using networks of hundreds and thousands of computers for processing

# Different Hardware

- Memory continues its drop in prices
  - You can get “commodity” systems with 512GB these days
  - This allows to run many applications in single box memory.
    - Happily “internet population” grows slower than hardware improvements.
- SSD – breakthrough in cost per IOP
  - Feasible replacement for Hard drives
    - 64GB Intel X25-E SSD about \$800
    - 146GB 2.5” 15K RPM about \$400
  - Earlier in its technology S-Curve
    - More rapid improvements are expected
- SSD still very expensive for large storage.

# Different Hardware

- SSD Means
  - You can allow larger “miss rate” and operate with smaller memory.
  - Random IO is not as bad any more
    - Nested loops join is cool again
- Databases spent years optimizing for sequential IO
  - And now need to make changes
- The optimization of IO path becomes very important
  - RAID cards, File Systems, Innodb are often not ready to handle tens of thousands of IO operations per device

# Flash is not SSD

- “Hard Drive Replacement” is only one of the uses of Flash memory technology
- PCI-E cards by FusionIO and others
- Violin Memory “appliance” model
- Virident using Flash for DRAM chip replacements

# Other hardware innovations

---

- **Fast Interconnects**
  - 10Gb Ethernet getting traction
  - Dolphin Interconnect
- **Data flow processing**
  - Kickfire doing for databases what happened to 3D graphics processing 10 years ago

# Technology Needs

# Data Consistency

- Low consistency requirement
  - Typically it is not money
  - Or not a lot of money at least
- Ready to pay consistency for performance
  - Loose durability settings
  - Limited use of Foreign Keys
  - “Dirty” synchronization algorithms
  - Non exact result algorithms
    - Such as sampling

# System multi use

- We want to have the cake and eat it
- We want to use same system for everything
  - OLTP transactions
  - Information Retrieval/Full text search
  - Reporting Queries
- Additional systems, ELT process make development and operations more complicated

# Global Audience

- Geo Distribution
  - High availability and response time
- Can't pay “classical” cost of distributed transactions
- Synchronous replication is too slow
- Asynchronous replication makes handling conflicts tricky
  - Or poses visibility challenges

# Data Visibility

- Not every session/operation requires the same view of data
  - Helps use of async replication locally and globally
- Visibility options go well beyond traditional SQL isolation modes and MVCC
- Visibility often applies to user Web sessions
  - Which are not represented by DB transaction or session
- Example:
  - You post a forum message. You should see it at once
  - Different user may see it in several seconds

# Dynamic Sizing

- Capacity requirements may be unforeseen
  - Successful application experience fast exponential growth
- Capacity requirements can be non uniform
  - Event specific resources
  - Advertising campaigns and unexpected promotions
- Cloud Computing as a response
- Few modern “database” technologies work very well with cloud.

# ACID vs BASE

- Alternative to the “ACID” concept of standard databases
- **BASE**
  - **B**asically **A**vailable
  - **S**oft State
  - **E**ventually Consistent
- One of the interesting alternative modes
  - Published by ACM

# What does it tell us ?

---

- Old Technologies do not work well as they could
- New technologies are not mature
- Great time to Innovate !
- Great time to learn new technologies
  - We bring some of them at Percona Performance Conference