

# What eBay really wants to buy.

Charles Kalko

Operations Strategist, eBay Technical Operations

25<sup>th</sup> Sept 2009



# Copyright Notice

- © 2009 eBay Inc. All rights reserved.
- No part of these materials may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording, or otherwise) without the prior permission of eBay Inc.
- eBay and the eBay logo are registered trademarks of eBay Inc.
- PayPal and the PayPal logo are registered trademarks of PayPal, Inc.
- Other trademarks and brands are the property of their respective owners.
- Please do not take our picture or record the class/session without asking permission.

# Background

- In 2006 we realized added that our exponential growth of our site was unsustainable
  - We where going to have to start building multiple data center/year
  - That cost of running our site was going affect our business
- We also realized that technology was now limiting the growth rate of our business
  - The business was literally waiting for operations to build out new capacity.
- We needed to change our mind set from always build new to a more mature managed approach.
- 2006 is the year that eBay needed Green Metrics.

# eBay – Some Interesting Numbers

- 241+ million registered users
- 160+ million plus Items
- 8 million plus new items per day
- 58.5 billion SQL transactions per day
- 800+ production database instances (inc replicas)
- 100+ clusters
- ~4PB production data
- ~8PB data warehouse

# Parallel Efficiency

Our definition of Parallel Efficiency is: A even distribution of workload across a parallel processing environment.

- An Example: Evenly distributing a workload across application servers in many disparate locations.
- We generate the metric by sampling the max average and min average of pool members at peak load. □
- A pools capacity is based on its worst members.
  - Since pools might have three or four generations of hardware in them
- We had a homogenous work load on heterogonous hardware.
- For example:
  - If the Max average CPU of the pool = 90% (old servers)
  - and the Min average CPU of the pool = 60% (new servers)
  - then Pool Utilization is said to be 90% (old servers dominate)

# Our main metric Transactions/Watt

- We have seen significant growth in CPU capabilities
- Nearly 12:1 improvement for some of our workloads
- Cost of servers are largely flat
- Multi-core is paying off in Transactions/Watt
- Cost of operating a server far outweighs the initial outlay

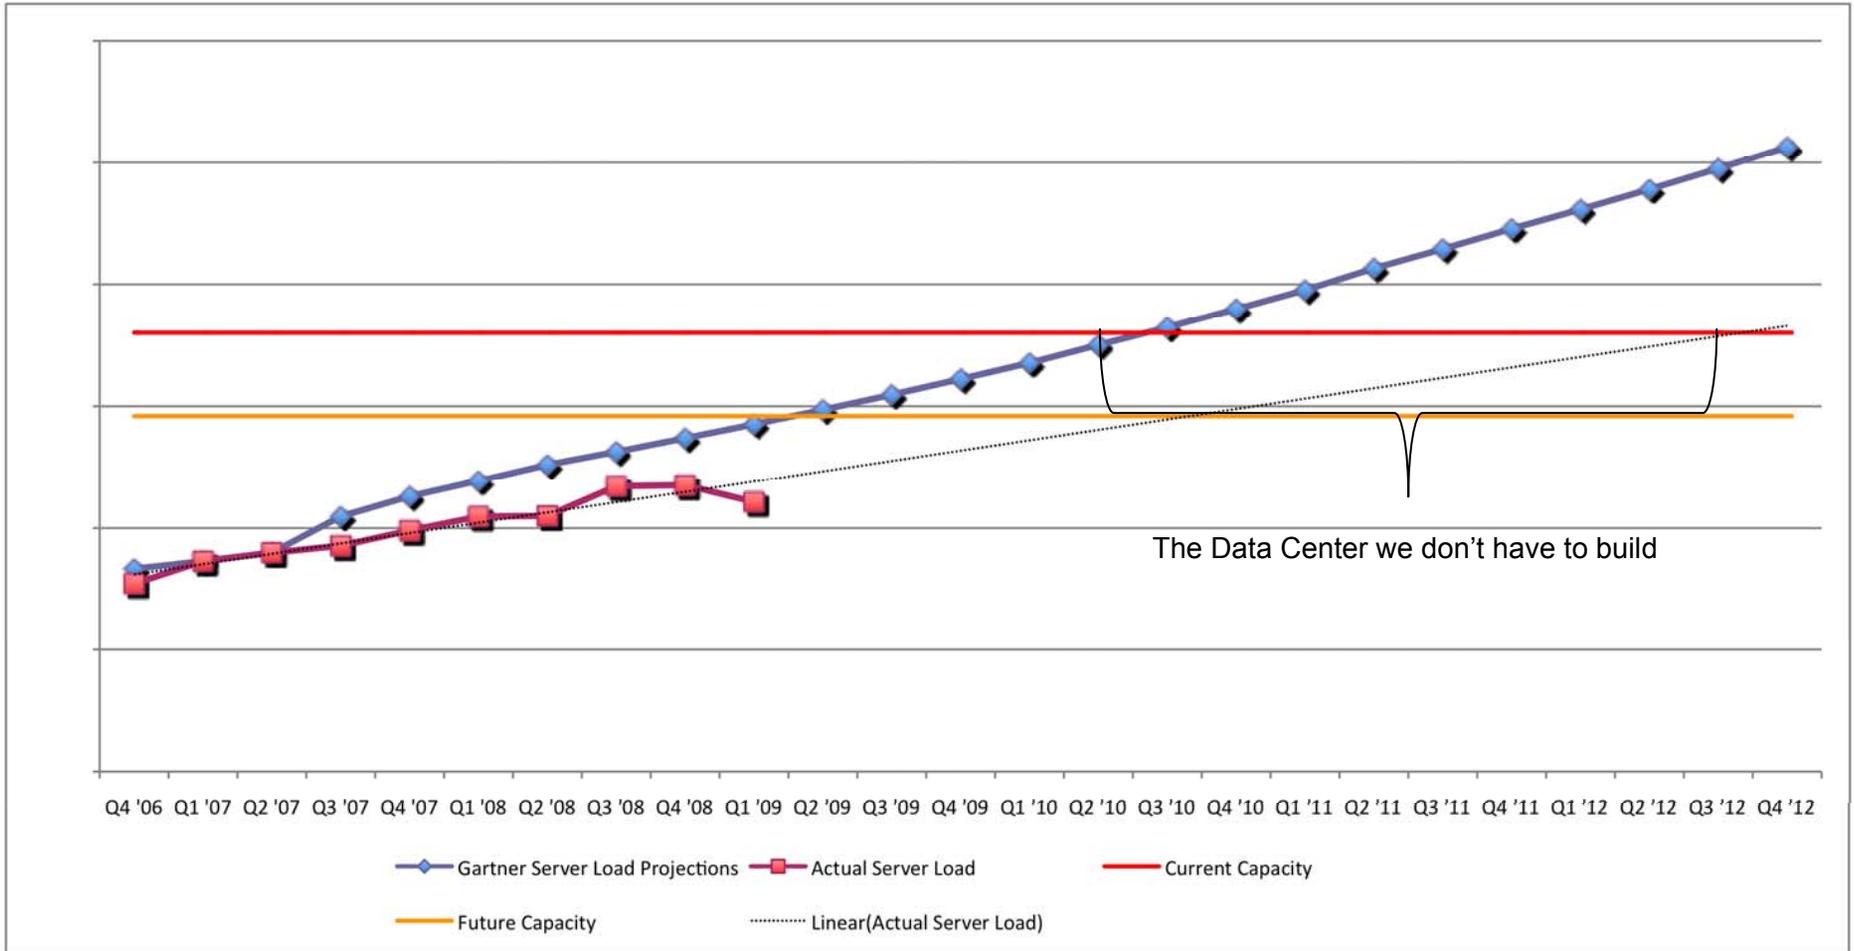
	2 x 2.6Ghz	2 x 3.0 Ghz	4 x 2.2 Ghz	8 x 2.5 Ghz	8 x 2.26 Ghz
2 x 2.6 Ghz	1	0.7	0.5	0.28	0.08
2 x 3.0 Ghz	1.4	1	0.7	0.39	0.17
4 x 2.2 Ghz	2.0	1.4	1	0.55	0.33
8 x 2.5 Ghz	3.6	2.6	1.8	1	0.61
8 x 2.26 Ghz	12	6.0	3.0	1.65	1

# Managing Transaction/Watt with Tech Refresh

- We stopped purchasing servers and started leasing them.
- Using PE in conjunction with hardware upgrades gave us a finer level of control over our IT load.
  - We could now manage pool workload horizontally and vertically
- We could now leverage our IT load by tuning our pools based on seasonal demand.
- We are staying inside of Moore's Law
  - We can double our site capacity every 18 months and still stay in the same foot print.

# The results of our efforts

## Volume Trend



## What we expect from our OEM's today.

- Our primary metrics are Transaction/Watt and, \$/Watt
  - We benchmark all new platforms to determine performance to cost
  - How do we optimize our power consumption?
  - Right sized power supplies
  - Lower power CPU's
  - Smarter power management (larger and smarter fans)
  - Flexible leasing return options.
- We are also asking for other things
  - A SKU to order less wasteful packaging and filler.
  - Shipping a fully built out and burned in rack of systems.

# What we are expecting from our OEM's tomorrow.

- We are starting to track our Green House Gas emissions.
  - What is the embedded cost of energy in manufacturing and disposal.
- A more generic performance to cost metric, than our workloads.
- More efficient power supplies at lower utilization
  - With on board battery backup similar to the Google model.
- Higher inlet temperatures (28C ~ 30C)
- Higher exhaust temperatures (for better  $\Delta T$ )
- Liquid cooling
- Reduced components on the system.
  - Delete USB and LED's
  - Stripped down video cards
  - SSD vs. Spindles?