Cloud Computing Imperatives

James Hamilton

2008-04-11

JamesRH@microsoft.com

http://research.microsoft.com/~JamesRH

blog: http://perspectives.mvdirona.com

Services Economies of Scale Inescapable

- Substantial economies of scale

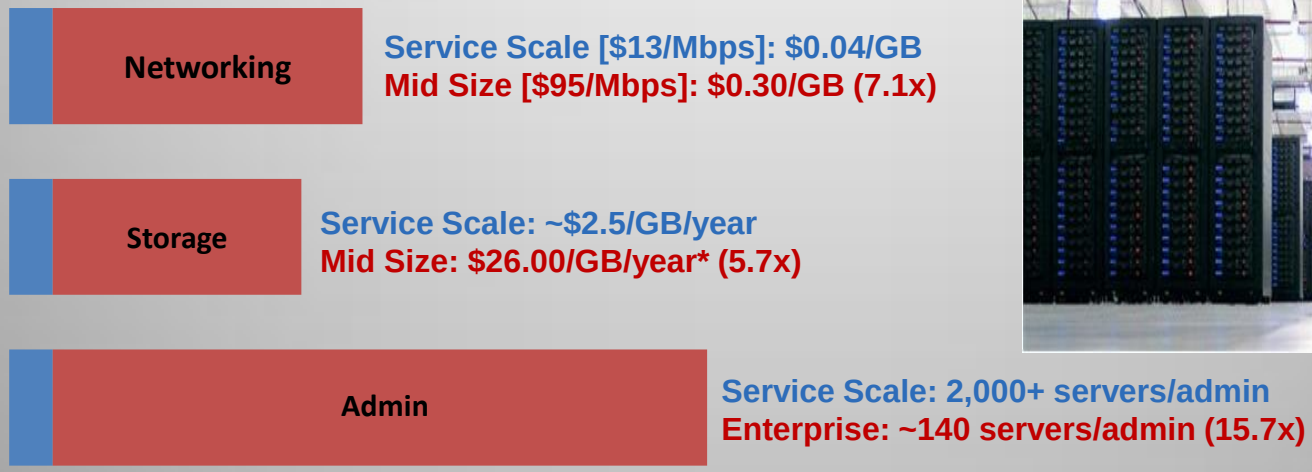

- High cost of entry
  - Physical plant expensive: 10MW roughly $200M
- Summary: significant economies of scale but at very high cost of entry
  - Small number of large players likely outcome

Power & Communications Limit

- Process core cycles are cheap & getting cheaper
- What limits application of infinite cores?
  - **Power**: cost rising and will dominate
  - **Data**: inability to get data to processor when needed
- DC power & mechanical trending up, servers down
- Most sub-Moore attributes require most innovation
  - Infinite processors require infinite power
  - Getting data to processors in time to use next cycle:
    - Caches, multi-threading, ILP,… consume power
- Latency bigger problem than bandwidth

	CPU	DRAM	LAN	Disk
Annual bandwidth improvement	1.5	1.27	1.39	1.28
Annual latency Improvement	1.17	1.07	1.12	1.11

Dave Patterson Graph

Yield Management, Optimization, & Data Analysis Dominate

- Yield mgmt first used in airline industry
  - Airplane more expensive than computation
- Heavily used in retail & Finance
  - Shelf space opt, supply chain optimization
  - 1000's of node financial analysis systems
- Declining cost of computing allows yield-management of less expensive resources
- Analysis systems dominate transactional systems
  - Transactional workload represents sales & changes in physical world
  - Analysis grows at rate of cost decline and potentially include data from ALL transactions

Resource Consumption Shaping

- Essentially yield mgmt applied to DC
- Network egress charged at 95th percentile:
  - Push peaks to troughs
  - Fill troughs for “free”
- Charged symmetrically so ingress also effectively free
- Power also charged at 95th percentile
  - Server idle to full-load : 158W to 230W (60% common)
  - S3 (suspend) or S5 (off) when server not needed
- Disks come with both IOPS capability & capacity in device fixed ratio
  - Mix hot and cold data
- Encourage urgency differentiation in charge-back model

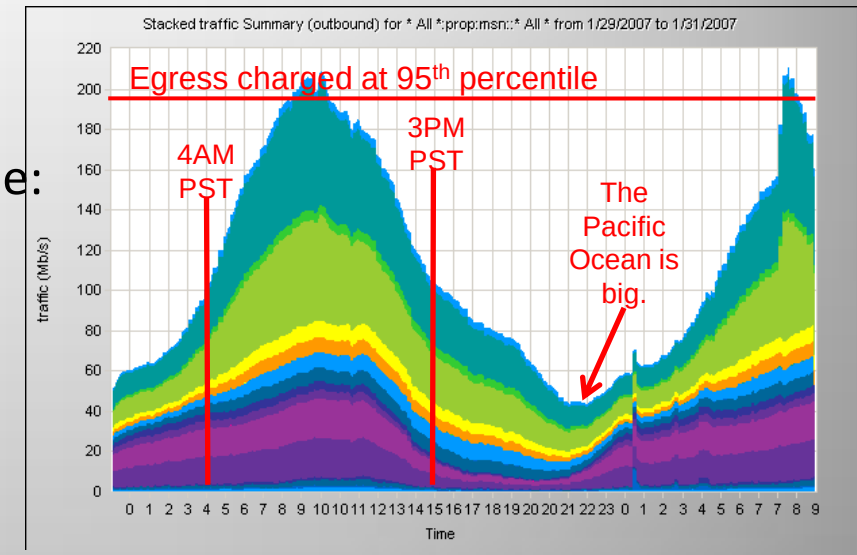

David Treadwell Graph

Mass Distribution & Mass Centralization

- Mass Distribution:
  - Device numbers exploding (cell phones +1B/yr)
  - Edge computing resources exceed those in core
  - Move computation closer to user
- Mass Centralization:
  - Yield management, optimization, & data analysis
  - Data is the asset
  - Move computation closer to data

Summary

- Five services imperatives:
  1. Services Economies of Scale Inescapable
  2. Power & Communications Limit
  3. Yield Management, Optimization, & Data Analysis Dominate
  4. Resource Consumption Shaping
  5. Mass Distribution & Mass Centralization
- TJ Watson appears to have been partly correct
  - Small number of very high scale services support vast majority of server-side computing
  - But, edge device count growing explosively large & with far more resources in aggregate