云数据管理系统：挑战与机遇

陆嘉恒

随着信息产业的发展，企业和公司产生的数据量快速增长，通常数据规模可以达到 TB 甚至 PB 级别。如何管理和分析海量数据是目前很多领域所面临的问题，例如在医疗、通信和互联网领域。传统的数据管理技术已经不能完全满足海量数据管理的需求，云计算技术的出现为海量数据管理带来了机遇，利用云平台来存储和管理海量数据是当前的研究热点之一。该 ppt 主要介绍了云数据管理方面的挑战和机遇，包括云数据管理的必要性和需求，数据管理的弱一致性研究，数据拷贝的一致性维护，数据的并发处理和版本模型等前沿数据库研究问题。

Cloud-based Data Management: Challenges & Opportunities

Jiaheng Lu

Renmin Universtiy of China

Research experience and interesting

- **National University of Singapore PhD**
  - XML query processing and XML keyword search
- **University of California, Irvine Postdoc**
  - Approximate string processing
  - Data integration and data cleaning
- **Renmin University of China**
  - Cloud data management
  - XML data management

Outline

- **Motivation: cloud data management**
- **Database Future and Challenges：**
  - Large-scale Data management & transaction processing
  - Cloud-based data indexing and query optimization

Motivation: Internet Chatter

BLOG Wisdom

- **"If you want vast, on-demand scalability, you need a non-relational database." Since scalability requirements:**
  - Can change very quickly and,
  - Can grow very rapidly.
- **Difficult to manage with a single in-house RDBMS server.**
- **Although RDBMS scale well:**
  - When limited to a single node.
  - Overwhelming complexity to scale on multiple sever nodes.

Current State

- **Most enterprise solutions are based on RDBMS technology.**
- **Significant Operational Challenges:**
  - Provisioning for Peak Demand
  - Resource under-utilization
  - Capacity planning: too many variables
  - Storage management: a massive challenge
  - System upgrades: extremely time-consuming

What is Cloud Computing?

- **Old idea: Software as a service (SaaS)**
  - Def: delivering applications over the internet
- **Recently: "[Hardware, infrastructure, Platform] as a service"**
  - Poorly defined so we avoid all "X as a service"
- **Utility Computing: pay-as-you-go computing**
  - Illusion of infinite resources
  - No up-front cost
  - Fine-grained billing (e.g. hourly)

Why Now?

- Experience with very large datacenters
  - Unprecedented economies of scale
- Other factors
  - Pervasive broadband internet
  - Pay-as-you-go billing model

Cloud Computing Spectrum

- Instruction Set VM (Amazon EC2, 3Tera)
- Framework VM
  - Google AppEngine, Force.com

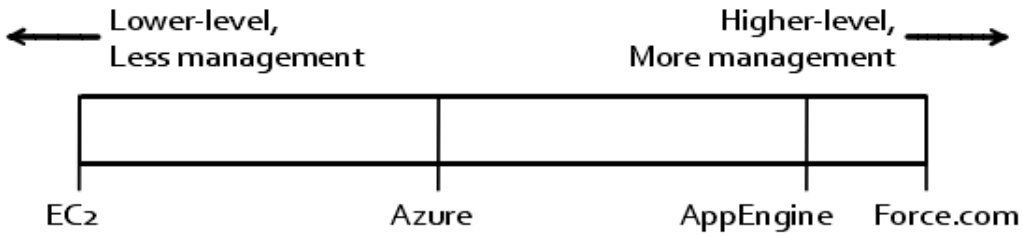

Cloud Killer Apps

- Mobile and web applications
- Extensions of desktop software
  - Matlab, Mathematica
- Batch processing/MapReduce

Economics of Cloud Users

- Pay by use instead of provisioning for peak

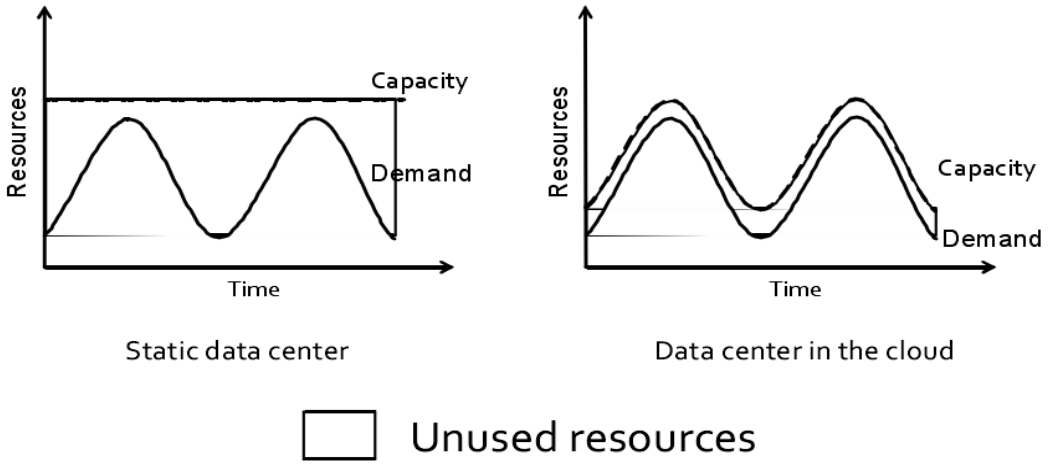

Economics of Cloud Users

- Risk of over-provisioning: underutilization

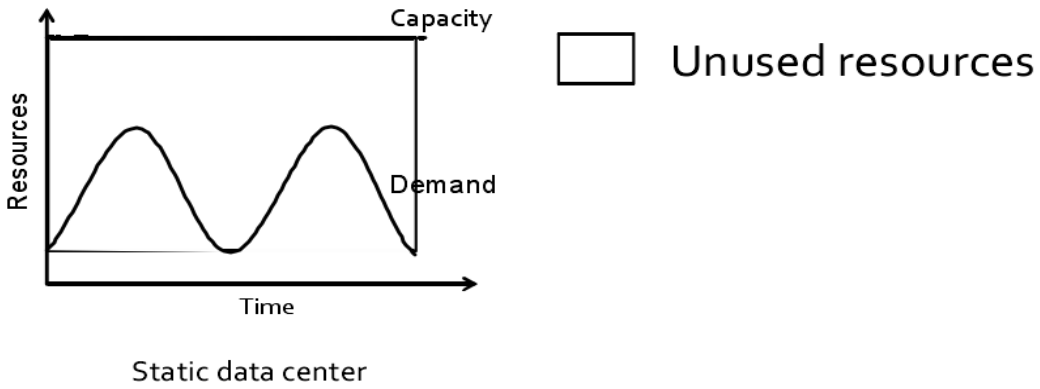

Economics of Cloud Users

- Heavy penalty for under-provisioning

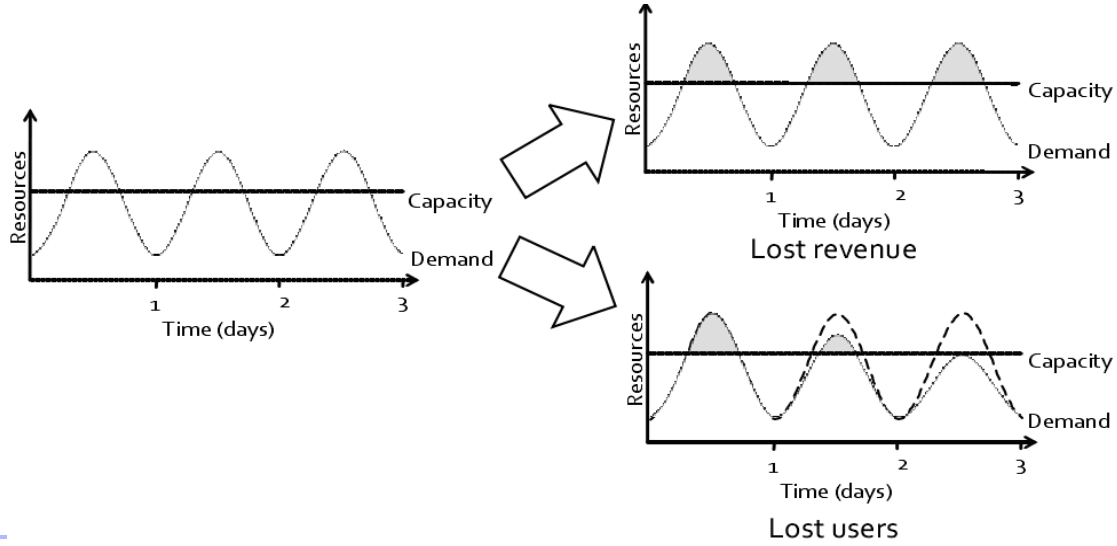

Engineering Definition

- Providing services on virtual machines allocated on top of a large physical machine pool.

Business Definition

- A method to address scalability and availability concerns for large scale applications.

Data Management in the Cloud?

The Vision

- R&D Challenges at the macro level:
  - Where and how does the DBMS fit into this model.
- R&D Challenges at micro level:
  - Specific technology components that must be developed to enable the migration of enterprise data into the clouds.

Data and Networks: Attempt Ⅰ

- Distributed Database (1980s):
  - Idealized view: unified access to distributed data
  - Prohibitively expensive: global synchronization
- Remained a laboratory prototype:
  - Associated technology widely in-use: 2PC

Data and Networks: Attempt II

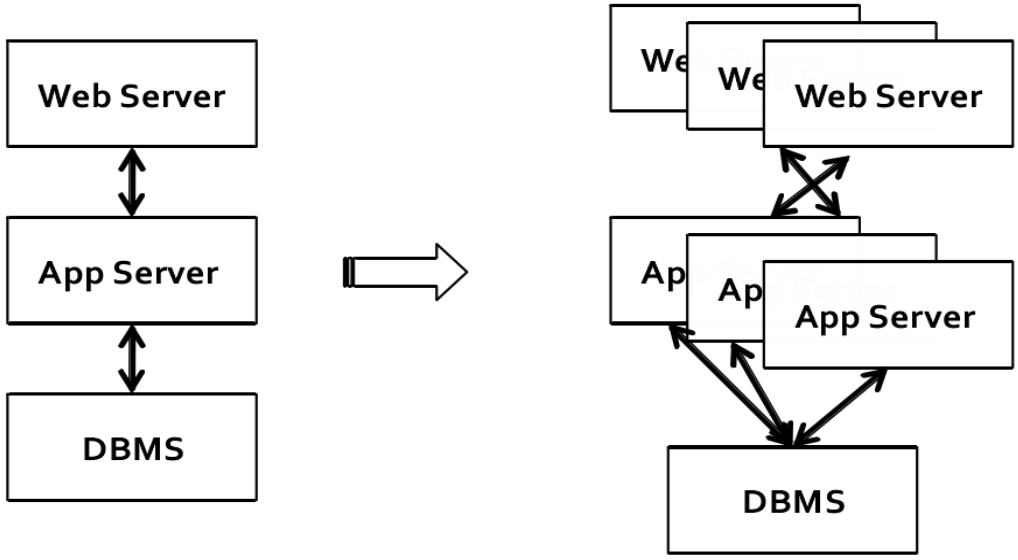

Data and Networks: Pragmatics

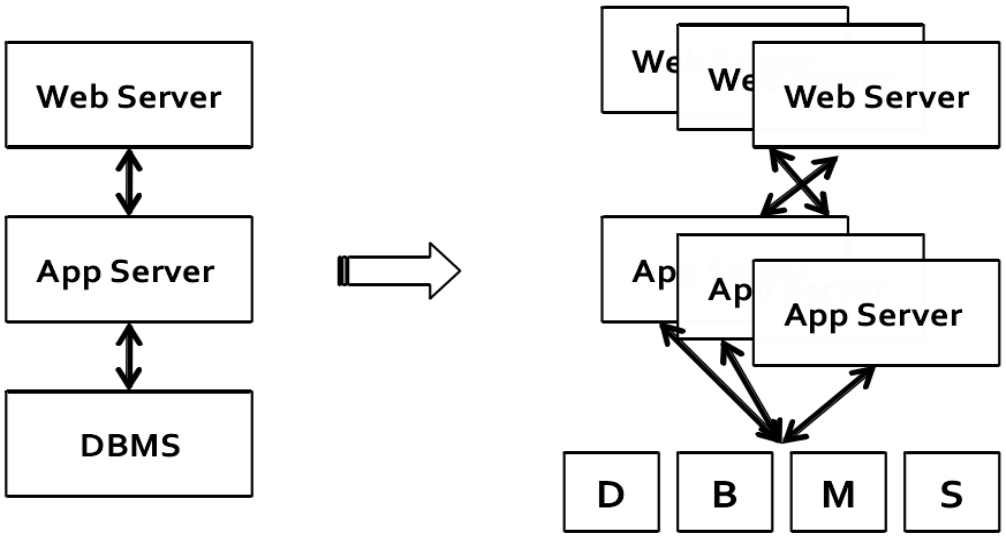

Database on S3: SIGMOD'08

- Amazon's Simple Storage Service(S3):
  - Updates may not preserve initiation order
  - No "force" writes
  - Eventual guarantee
- Proposed solution:
  - Pending Update Queue
  - Checkpoint protocol to ensure consistent ordering
  - ACID: only Atomicity + Durability

Unbundling Txns in the Cloud

- **Research results:**
  - CIDR'09 proposal to unbundle Transactions Management for Cloud Infrastructures
  - Attempts to refit the DBMS engine in the cloud storage and computing

Analytical Processing

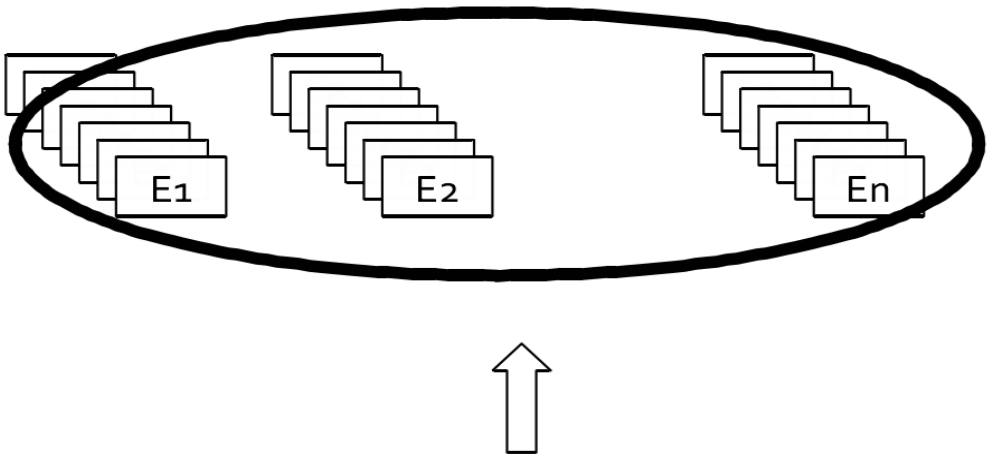

Analysis Queries: Distributed Processing

Architectural and System Impacts

- **Current state:**
  - MapReduce Paradigm for data analysis
- **What is missing:**
  - Auxiliary structures and indexes for associative access to data (i.e., attribute-based access)
  - Caveat: inherent inconsistency and approximation
- **Future projection:**
  - Eventual merger of databases (ODSs) and data warehouses as we have learned to use and implement them.

Underlying Principles: CIDR'2009

- **Business data may not always reflect the state of the world or the business:**
  - Inherent lack of perfect information
- **Secondary data need not be updated with primary data:**
  - Inherent latency
- **Transactions/Events may temporarily violate integrity constraints:**
  - Referential integrity may need to be compromised

Data Security & Privacy

- Data privacy remains a show-stopper in the context of database outsourcing.
- Encryption-based solutions are too expensive and are projected to be so in the foreseeable future:
  - Private Information Retrieval (Sion'2008)
- Other approaches:
  - Information-theoretic approaches that uses data-partitioning for security (Emekci'2007)
  - Hardware-based solution for information security

Self management and self tuning in cloud-based data management

- Self management and self tuning

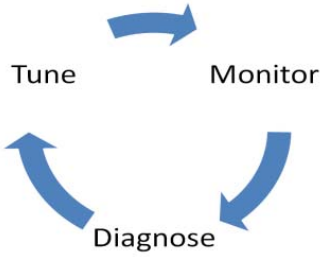

- Query optimization on thousands of nodes

Remarks

- **Data Management for Cloud Computing poses a fundamental challenge to database researchers:**
  - Scalability
  - Reliability
  - Data Consistency
- **Radically different approaches and solution are warranted to overcome this challenge:**
  - Need to understand the nature of new applications

References

- Life Beyond Distributed Transactions: An Apostate's Opinion by P.Helland, CIDR'07
- Building a Database on S3 M.Brartner, D.Florescu, D.Graf, D.Kossman, T.Kraska, SIGMOD'08
- Unbundling Transaction Services in the Cloud D.Lo,et, A.Fekete, G.Weikum, M.Zwilling, CIDR'09
- Principles of Inconsistency S.Finkelstein, R.Brendle, D.Jacobs, CIDR'09
- VLDB Database School (China) 2009 http://www.sei.ecnu.edu.cn/~vldbschool2009/VLDBSchool2009English.htm