



# CIDR2015 CLOUD SESSION

---

Alan Fekete (University of Sydney): session chair

# PUBLIC VS PRIVATE CLOUDS

- Private Cloud has a subset of the difficult challenges, and a subset of the potential benefits
- I will focus here on research targeting public clouds like AWS, Microsoft Azure etc

# Consumer-provider interface

- The consumer who runs some activity in the cloud is distinct from the provider who owns/operates the physical resources
- Their interests can diverge
  - Trust is limited
- There is a restricted communication channel
  - For the consumer to find out useful information about what they are getting (monitoring) or about their choices
  - For the consumer to let provider know what they desire (control)

*Research issue: choose effective level of abstraction in consumer-provider information exchange*

# Multi-tenancy and performance

- The set of activities that are running changes frequently
- The provider has a lot of flexibility in allocating activities among physical resources
- Different consumers' activities may share physical resources in ways that change frequently
  - The performance properties obtained by a given consumer change in ways that are complex (from interactions between the activities)
- The provider has a difficult optimization problem
- The consumer operates with limited knowledge of the actual performance they will get

*Research issue: optimization across tenants;  
also good performance models for the interactions*

# Security issues

- Consumer does not fully trust provider (and does not at all trust other consumers who may sit on same physical resources), but the consumer needs to give provider the consumer's data and have the provider run computations on it
- Perhaps consumer will use multiple providers

*Research issues: discovering side channels;  
allowing computation that preserves privacy; verifiable computation*

# Restricted functionality

- Many cloud-based data management platforms offer reduced functionality compared to traditional DBMS
  - Give up in-platform-joins-across-data-sets, consistent indices, multi-item transactions, constraint maintenance
  - In order to do well on scalability, fault-tolerance, flexibility

*Research issues: building stronger properties at application level on weak platforms; finding slightly stronger functionality that can still be supported with good performance*

# Heterogeneity

- Since consumer and provider do not share all interests, there is great desire among consumers to avoid lock-in and be able to migrate from one provider to another, or even to make use of multiple providers at one time

*Research issues: live migration, accommodating heterogeneity of interface*

# Geo-replication and geo-distribution

- It is now common for a cloud provider to have data centers in many parts of the world
  - Communication across data centers incurs very high latency
- A consumer may wish to have information stored across several of these regions
  - To tolerate failures of a whole data center
  - To bring data and computing closer to a consumer
  - To support legal requirements of various jurisdictions
- Coordination across data centers is so expensive that special solutions are needed
  - *Research issue: new coordination algorithms to minimize impact of very large latencies (esp high variability of latency); new choice of API to allow less coordination; optimization of placement and processing*