

# ACE in Clouds: Availability Changes Everything

**Chunming Qiao, IEEE Fellow**  
**Computer Science and Engineering, SUNY Buffalo**

**Collaborators: T. Furlani, R. Ramesh, S. Smith (SUNY Buffalo) and G. Lazsewski (Indiana University)**


# Cloud Technologies

- Basic infrastructure components:
  - Physical servers (and virtual machines, aka VMs), racks, clusters
  - Power distribution units (PDUs) and cooling infrastructures
  - Switches, routers and datacenter networks
- Increasing adoption/reliance
  - Providers: Amazon, Google, Microsoft, Rackspace, Salesforce...
  - Clients: individuals, and small to large companies/institutions
- Availability/reliability is a top concern
  - $\text{availability} = \text{uptime} / (\text{total period}) = 1 - \text{downtime} / (\text{total period})$
  - cited by 67%, followed by device based security (66%) and cloud application performance (60%).

# Failures are all too common

- Frequent small-scale failures and infrequent large-scale failures
- Typical first year for a new cluster (Jeff Dean, Google)
  - ~0.5 overheating (power down most machines in <5 mins, ~1-2 days to recover)
  - ~1 PDU failure (~500-1000 machines suddenly disappear, ~6 hours to come back)
  - ~20 rack failures (40-80 machines instantly disappear, 1-6 hours to get back)
  - ~5 racks go wonky (40-80 machines see 50% packet loss)
  - ~3 router failures (have to immediately pull traffic for an hour)
  - ~dozens of minor 30-second blips for DNS
  - ~1000 individual machine failures
  - ~thousands of hard drive failures

# Failures cost too much



# Why Current Cloud Services Are Flawed

- Current Service Level Agreement (SLA) is loosely defined in terms of availability/reliability measurements.
  - SLA is a contract between a user and the service provider (price, service/duration, and penalty etc.)
- Penalty term is not user-friendly. The refund is usually issued in the form of *credit* with a lot of exclusions.
  - Amazon EC2 will refund the user in the form of credit if fail to meet the SLA.
  - Rackspace will credit the user 5% month fee for each 30 mins network/infrastructure downtime, up to 100% monthly fee of the affected server.
- Lack of high availability/reliability guarantee for critical services
  - Cannot guarantee 3-9's (99.9%), let alone 5-9's as in Telco networks.

# Key Challenges and Solutions

A user/app may request:

- # of VMs  $\underline{n}$  (e.g., 100) to achieve certain response-time performance
- Minimum desirable availability  $\underline{\alpha}$  (e.g., 99.9%)
- Desirable contract duration  $\underline{t}$  (e.g., 3 months)

The cloud SP performs the following:

- Downtime prediction based on failure models
  - Model component failures
  - Determine downtime distributions
- Availability-aware cloud resource provisioning and allocation
  - Determine the optimal (minimal) # of backup VMs,  $\underline{k}$ , to be allocated
  - Both risk and energy minimizing placement of  $\underline{n+k}$  VMs
- SLA contract design\*
  - Determine its costs: Capex ( $\sim h(n; k)$ ) and Opex ( $\sim$ energy consumption)
  - A price list (schedule) for  $\langle$ duration, availability-guarantee, penalty $\rangle$

# AQUA: An Analytic approach to Quantifying Availability for Cloud Resource Provisioning and Allocation

CSR Medium Collaborative  
1409809 & 1409256

Lead: SUNY Buffalo

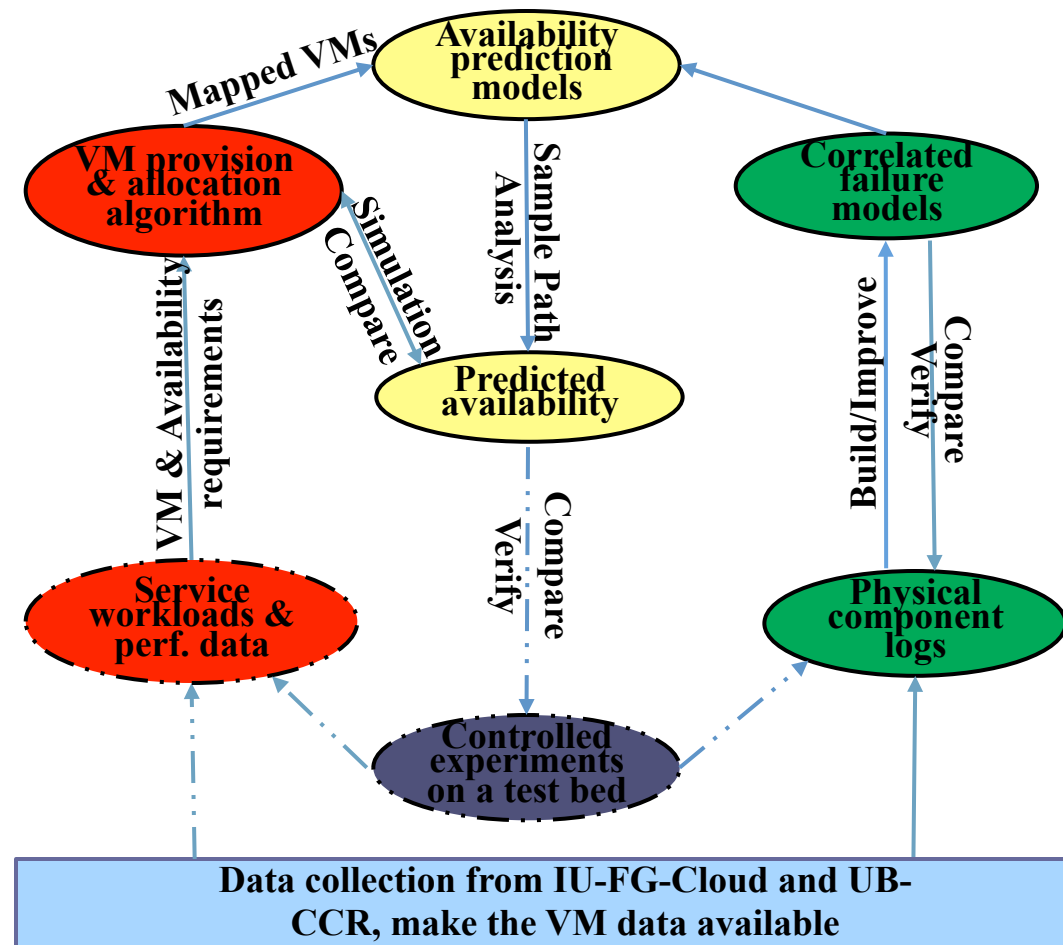
- PI: Chunming Qiao
- Co-I: S. Smith, R. Ramesh, T. Furlani

Collaborator: Indiana Univ.

- PI: G. Laszewski

<http://www.cse.buffalo.edu/AQUA/index.html>

(the scope of the NSF project includes neither the work in the dashed ovals, nor SLA Contract Design)



# More Information

- Z. Yang ,L. Liu, C. Qiao, S. Das, R. Ramesh and A.Y. Du, “Availability-aware and Energy-efficient Virtual Machine placement algorithm” accepted ICC 2015 (New)
- A. Y. Du, S.Das, Z. Yang, C. Qiao and R. Ramesh, “Predicting Transient Downtime in Virtual Server Systems: An Efficient Sample Path Randomization Approach”, to appear in IEEE Trans. on Computers (New)
- Yuan, S., Das, S., Ramesh, R. and Qiao, C., “Availability-aware Resource Provisioning, Pricing, and Allocation Adjustment in the Cloud”, Conference on Information Systems and Technology (CIST 2014), San Francisco, CA.
- Yuan, S., Das, S., Du, A.Y., Ramesh, R. and Qiao, C., “Cloud Resource Provisioning and Contract Adjustment in the Backdrop of SLA Violation Risk Mitigation”, Conference on Information Systems and Technology (CIST 2013), Minneapolis, MN.
- A.Y. Du, S. Das, C. Qiao, R. Ramesh and Z. Yang, “Downtime Predictions for Virtual Servers: A Study under Two Checkpointing Scenarios,” in Conf. on Info. Systems and Technology (CIST), 2012.
- A.Y. Du, S. Das, C. Qiao, R. Ramesh and Z. Yang, “Reliability in Cloud Computing: Downtime Predictions for Virtual Servers,” in 21st Workshop on Information Technologies and Systems (WITS), 2011.

- Contact:

Chunming Qiao  
CSE Department, SUNY Buffalo  
qiao@computer.org