



Western  
UNIVERSITY • CANADA

# DCSim

**A Data Centre Simulation Tool for Dynamic Virtualized Resource Provisioning**

Michael Tighe, Gaston Keller, Dr. Michael Bauer, Dr. Hanan Lutfiyya

April 13, 2012

# Agenda

Discuss the problem of dynamic resource management in the cloud. Introduce a new simulator, DCSim, to help develop new strategies.

# Dynamic Resource Management in the Cloud

# Cloud Computing & IaaS


- Computing resources on-demand, billed as a utility
- Hosted in large data centres
- Infrastructure as a Service (IaaS)
  - Clients rent servers with low-level access
  - Clients responsible for operating system and applications

# Provider Challenges

- Servers are typically underutilized
  - Average loads of a web server can be 30% of peak
- Wasted resources equates to:
  - More physical resources required
  - More power consumed
  - Higher maintenance costs

# Virtualization

- Run multiple “virtual machines” on a single physical machine
- Make better use of a physical server’s resources



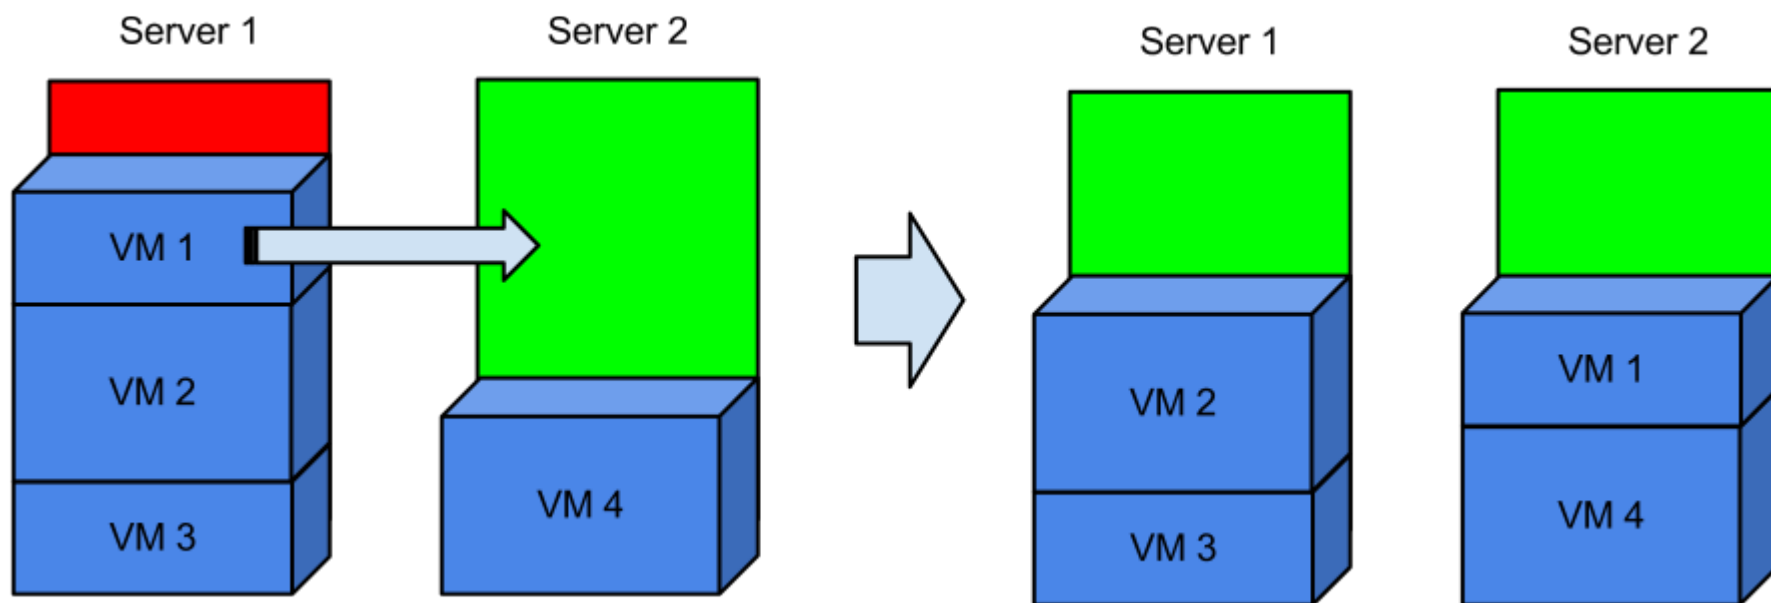
# Problem!

- Resources still underutilized = overcommit!
- Workload of VMs are dynamic, therefore so are its resources needs
- What happens if a server doesn't have enough resources to meet demand?



# Solution?

- Dynamically reallocate and provision resources to meet changing demand
- Virtual Machine Live Migration



# Algorithms

- Determining which VM to migrate where and when is difficult
  - Similar to multi-dimensional bin packing problem (NP-Hard), except more complex
  - Algorithms must execute quickly to respond to dynamic resource demands
  - Balance needs of clients (SLA) with Cloud Provider costs (i.e. power, infrastructure)

# Evaluating Algorithms

- How can we experiment with and evaluate dynamic resource provisioning algorithms?
- Physical test-bed
  - Implementation is complex
  - Long experiments
  - An entire data centre to experiment with?

# Evaluating Algorithms

- Simulation

- Quickly evaluate algorithms
- Simulate large scale over long time periods
- Needs to be supplemented with small scale real-world experiments

# Existing Simulators

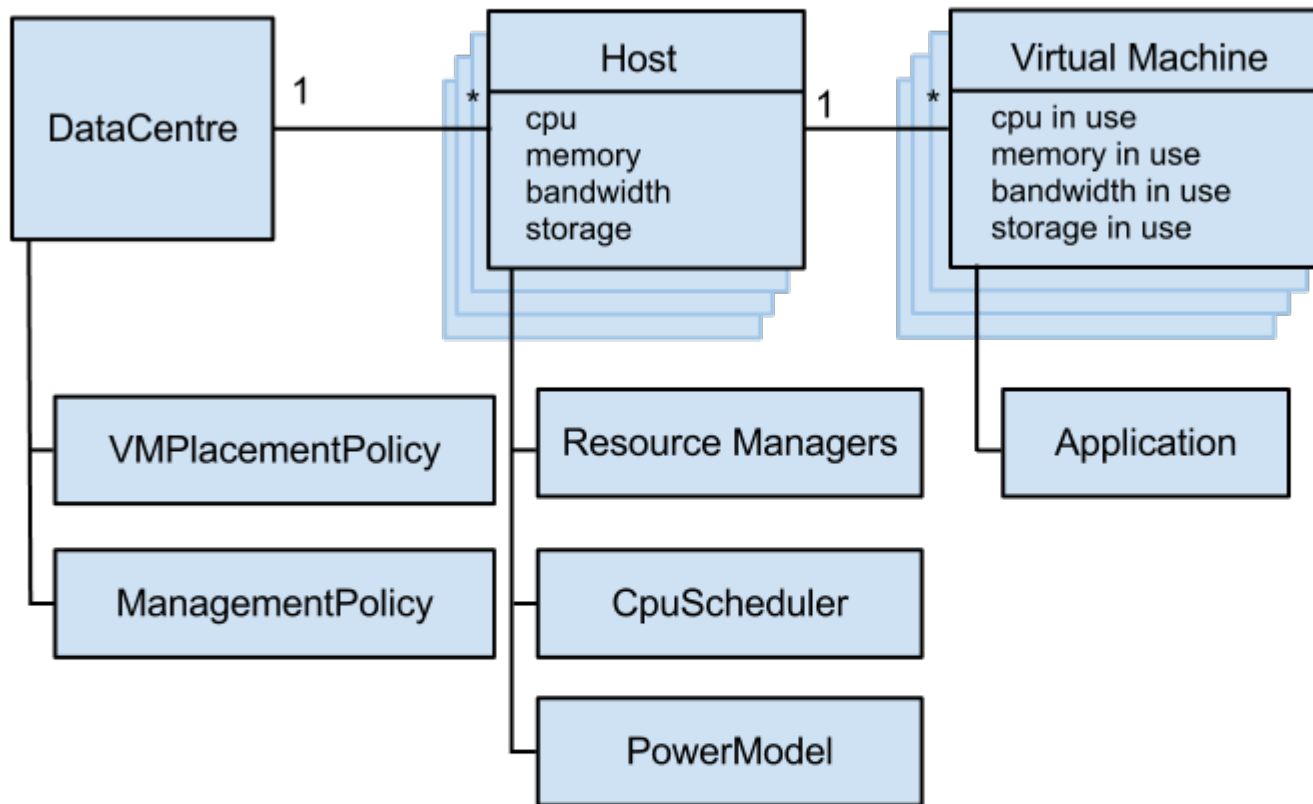
- GreenCloud, MDCCSim, GDCCSim
  - No virtualization, single tenant, commercial, focus on physical data centre design...
- CloudSim
  - Close to what we require
  - Originally designed for grid-style workloads
  - Missing some features (VM replication, VM dependencies, work conserving CPU...)

# Introducing DCSim

# DCSim: A New Simulator

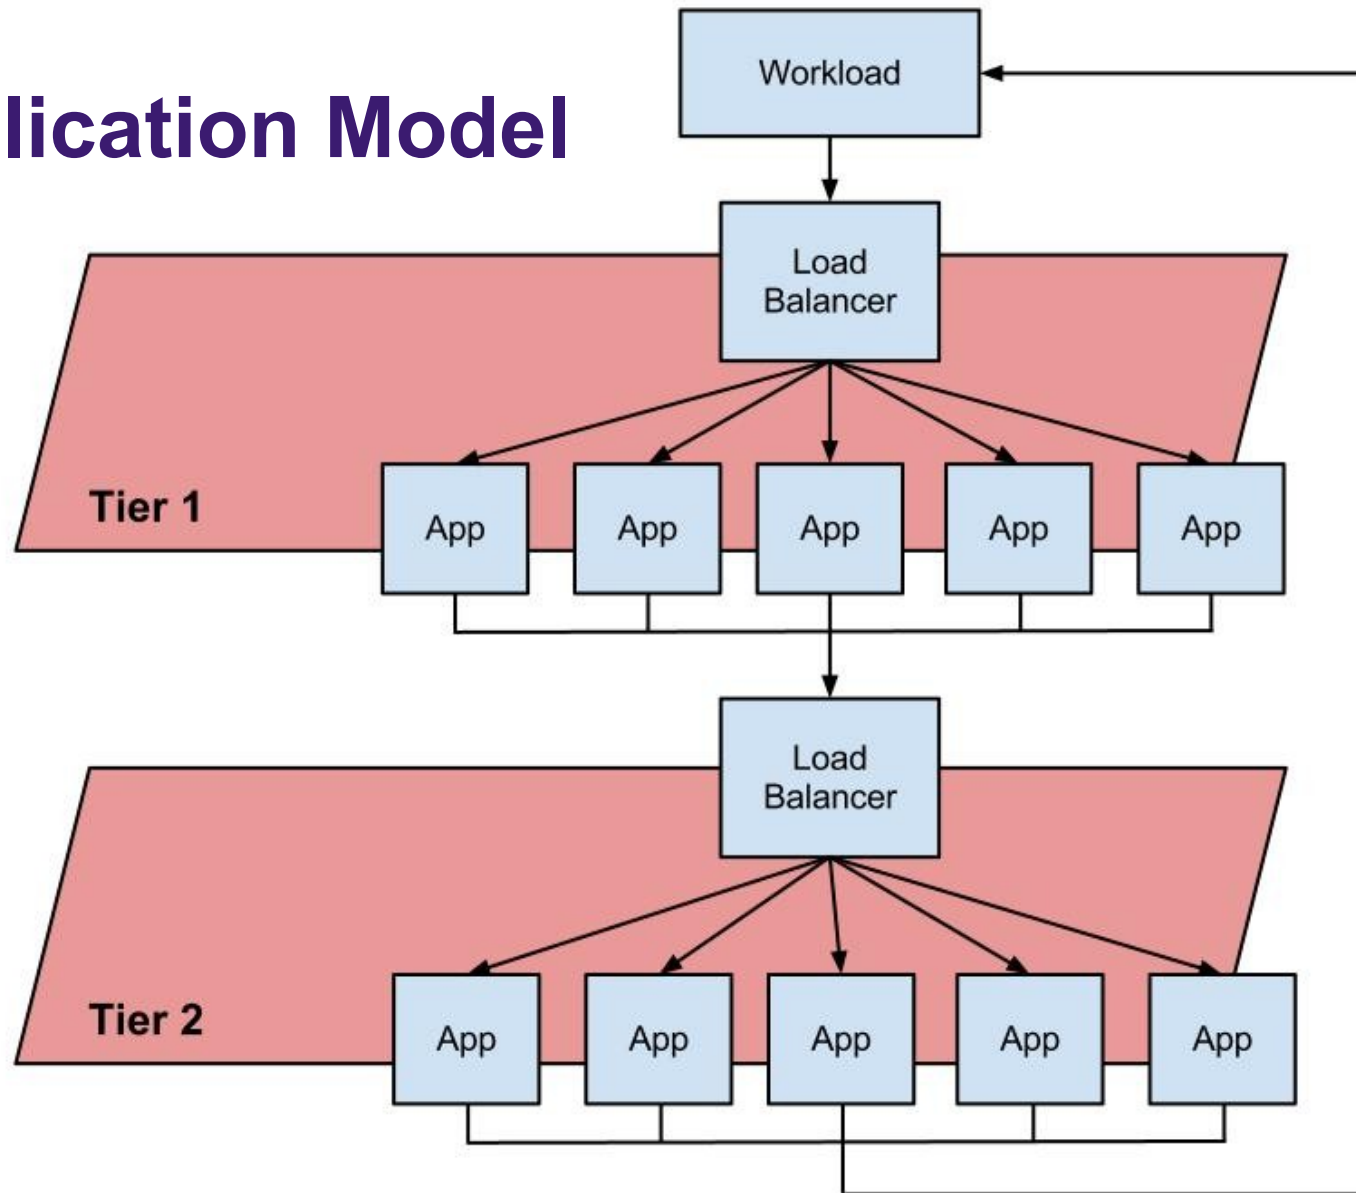
- Work conserving CPU (KVM, Xen, VMWare)
- Overcommitting of CPU
- Application Model (transactional, continuous)
- VM Dependencies
- VM Migration
- VM Replication
- Server power-on/suspend/power-off time
- *Extensible*

# Architecture





# Application Model



# Metrics

- SLA Violation
- Power Consumed (kWh)
- Number of Migrations
- Active Hosts
- Host-hours
- Active Host Utilization
- ...

# Evaluation

- Compare three methods of VM allocation
  - Static allocation for peak load
  - Static allocation for average load
  - Dynamic reallocation
- 200 Hosts, 400 VMs
- Workload traces from Internet Traffic Archive
- VMs initially placed in random order
- 10 simulated days

# Evaluation Results

	Static Peak	Static Avg	Dynamic
# Migrations	0	0	18547
Avg. Hosts	58.2	24	36.4
Host Util.	48.8%	94.6%	77%
Power (kWh)	2804.1	1431.1	2009.9
SLA Violation	0%	21.4%	0.8%

- Dynamic VM allocation achieves some of the advantages of both static alternative
- The simulator is demonstrated to be a useful tool to evaluate VM allocation strategies

# Performance

- Simulate 10 days with increasing numbers of hosts and VMs
- Dynamic VM reallocation performed
- Represents typical usage

	1	2	3	4	5	6
# Hosts	100	1000	10000	1000	1000	1000
# VMs	400	4000	40000	5000	6000	7000
Time	9s	~2 min	~1 hour	~3 min	~4 min	~5 min

# Work Using DCSim

- Comparison of First-fit Heuristics for VM Relocation <sup>[1]</sup>
- Dynamic Management Strategy Switching <sup>[under submission]</sup>
- Service Tier Auto-scaling
- Distributed Algorithms for VM Relocation & Consolidation

[1] G. Keller, M. Tighe, H. Lutfiyya, M. Bauer, “An Analysis of First Fit Heuristics for the Virtual Machine Relocation Problem” in SVM Proceedings, 6<sup>th</sup> Int. DMTF Academic Alliance Workshop on, Oct. 2012

# Future Work

- Data centre organization (racks, clusters)
- Networking
- Thermal impact of hosts & cooling costs
- Obtain traces/workloads with dynamic memory and bandwidth
- Memory Overcommitting



Western  
UNIVERSITY • CANADA