VIRTUAL MACHINE MIGRATION

K M Sabidur Rahman

Jan 29, 2015

Friday Group Meeting, Netlab

UC Davis

Agenda

- Who is interested?
- How VM migration works?
- Related research

What is VM Migration?

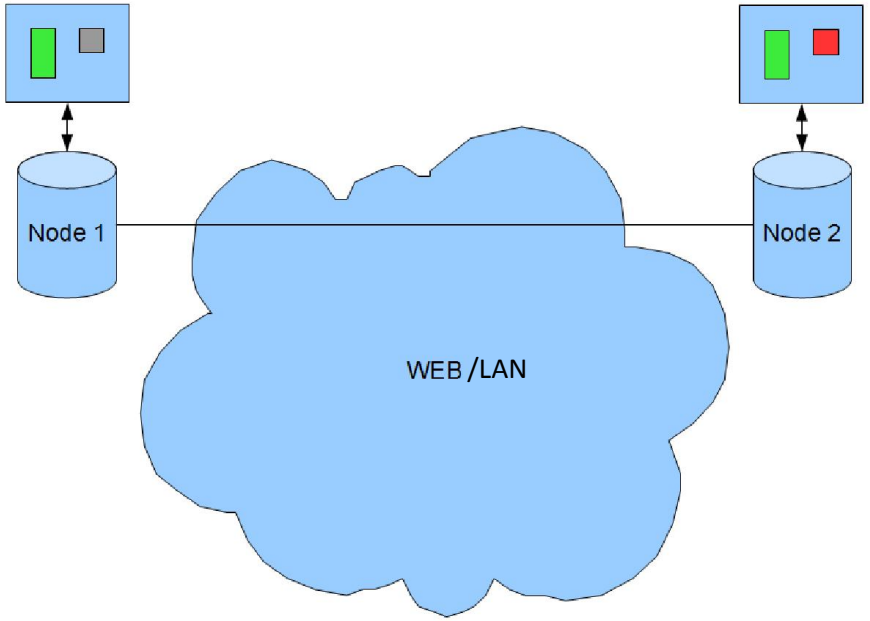

Image courtesy: [1]

Why VM Migration?

Online migration[2]:

- Vacating a VM Host system
- Targeting a particular VM Host
- Balancing Host workloads
- Optimizing physical resource utilization

Why VM Migration?

Offline migration[2]:

- Less resources (like memory and CPUs) used on the source and target VM Hosts
- Different processor types preventing online migration
- Source VM may not support Online VM Migration
- VM may just have stopped, the configuration information needs to be moved offline

Amazon-AWS

- VM Migration from one EC2 to another
- VM migration from EC2 to Virtual Private Cloud(VPC)

VMWare-vMotion

- Moving VM from one physical server to another

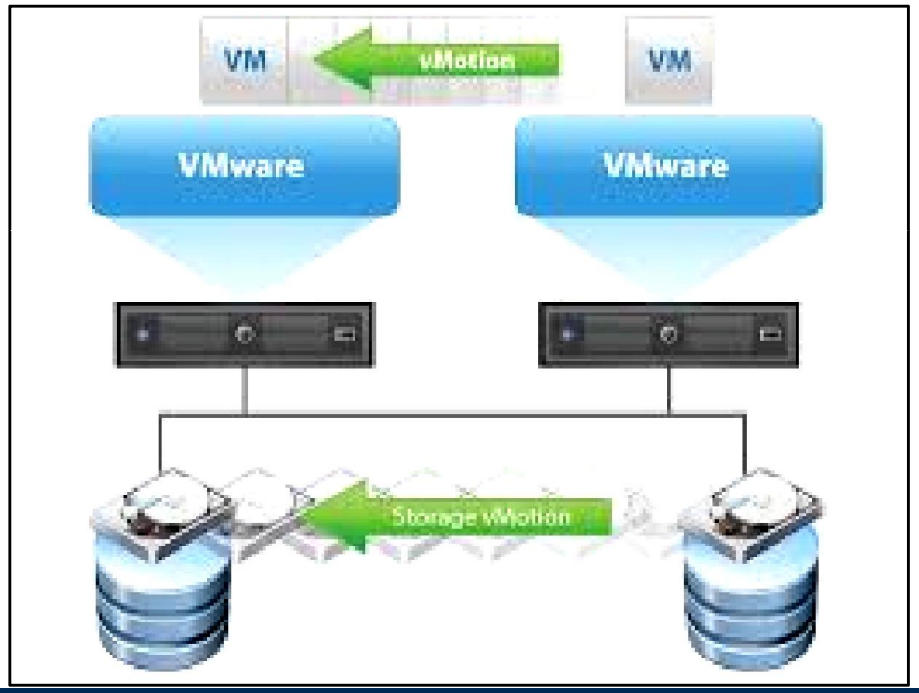

RackSpace-Migration services

- Migration from cloud service provider to in-house server
- From one cloud service provider to other cloud service provider

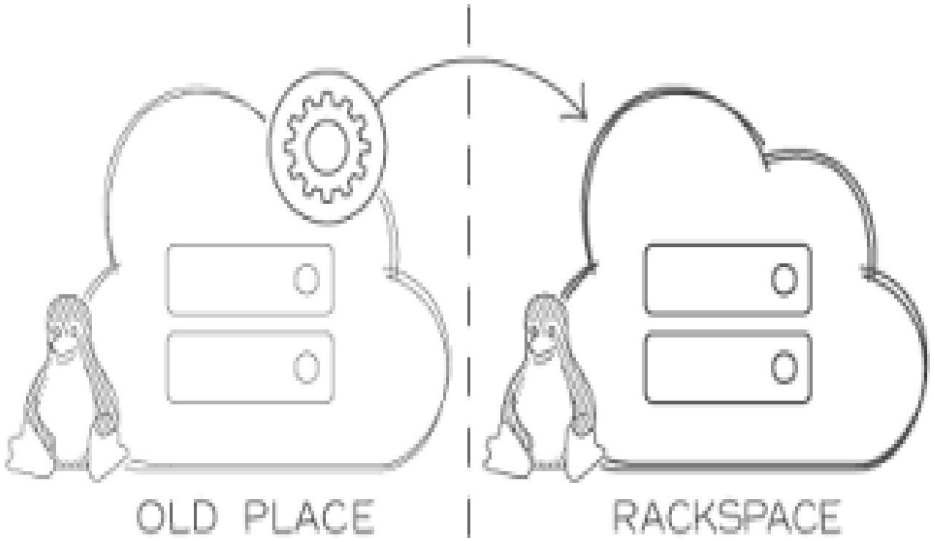

Citrix-XenServer

- Supports live migration
- VM on shared storage
- VM not on shared storage (XenMotion)

Cisco

- Teamed up with VMWare to enable Virtual Server mobility

Intel and AMD

- VM migration between heterogeneous processors

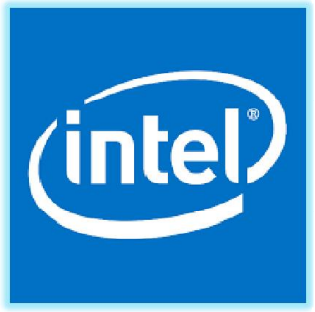

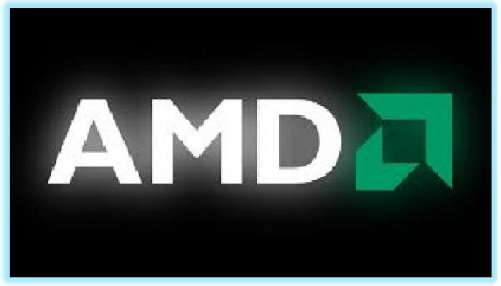

Microsoft-Hyper-V

- Hyper-V is Microsoft's solution for virtualization in Windows Server
- Migrating VM running on "Virtual Server 2005 R2 SP1" to Hyper-V

OpenStack and Red Hat

- OpenStack supports live and non-live migrations for cloud administrators
- Red Hat can live migrate VMs for maintenance, load balancing or power saving

CloudNet: Dynamic Pooling of Cloud Resources by Live WAN Migration of Virtual Machines

- Timothy Wood and Prashant Shenoy, University of Massachusetts Amherst
- K.K. Ramakrishnan and Jacobus Van der Merwe, AT&T Labs - Research

VM Migration over the WAN

- Move from one server to multiple other servers across the internet
- Cloud computing platform linked by a VPN based network infrastructure
- Minimize the cost of transferring storage and virtual machine memory
- Enabling migrations over low bandwidth and high latency Internet links

Things to consider

- Seamless secure cloud connections
- Resource pool that spans data centers
- Efficient WAN Migration

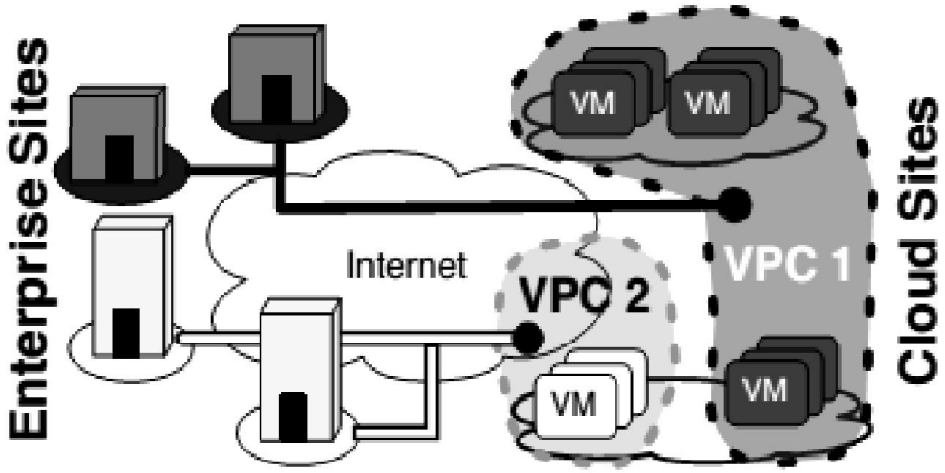

Image courtesy: [3]

Things to consider

- Seamless secure cloud connections
- Resource pool that spans data centers
- Efficient WAN Migration

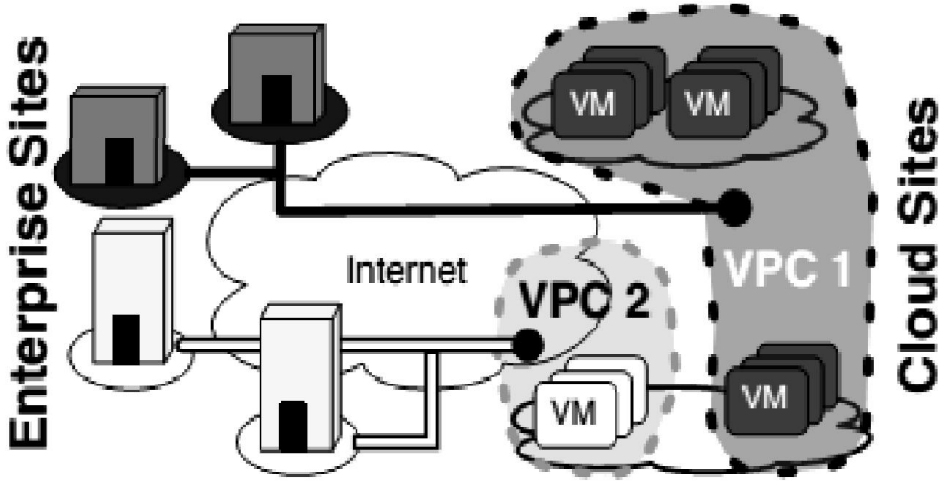

Image courtesy: [3]

Things to consider

- Seamless secure cloud connections
- Resource pool that spans data centers
- Efficient WAN Migration

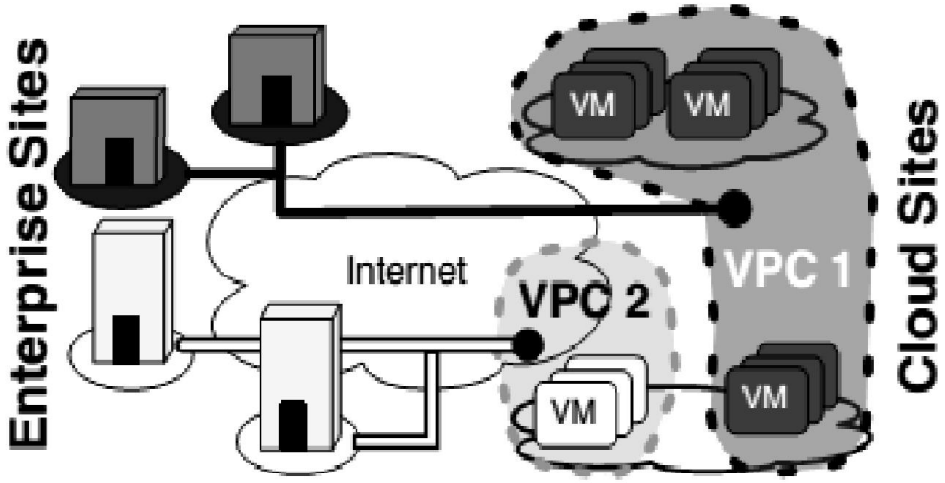

Image courtesy: [3]

Steps in Migration

Step 1: Establish layer-2 connectivity between data centers, if needed.

Step 2: If storage is not shared, transfer the application's disk state.

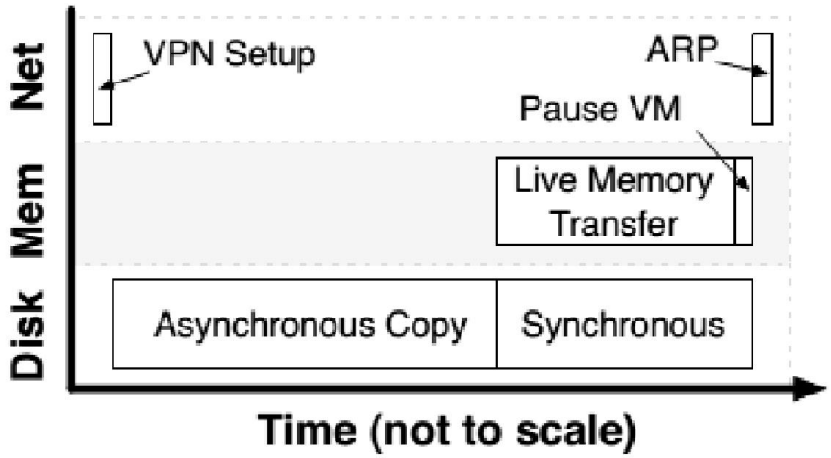

Image courtesy: [3]

Steps in Migration

Step 3: Transfer the memory state of the application to a server in Data Center B, as it continues running without interruption.

Step 4: Once the disk and memory state have been transferred, briefly pause the application for the final transition of memory and processor state to Data Center B.

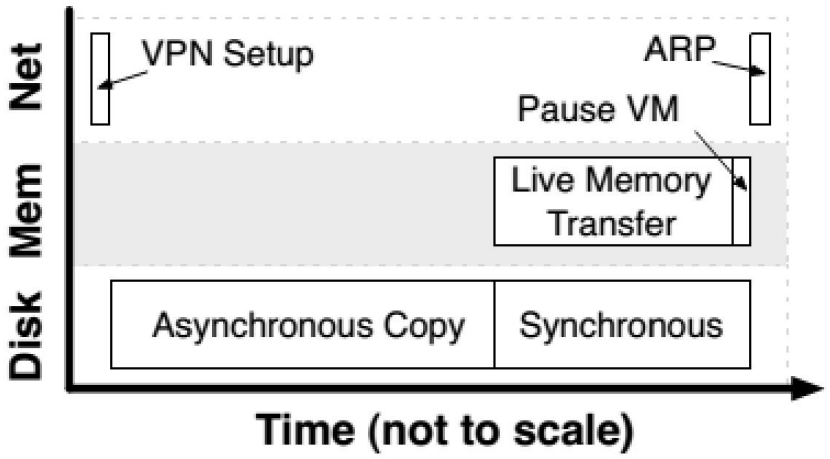

Image courtesy: [3]

Disk State Migration

- Initial copy is made asynchronously
- Synchronous replication, once the remote disk is stable
- Disk updates continuously propagated to the remote disk

Transferring Memory State

- “Pre-copy” mechanism to iteratively copy the memory contents
- Finally, VM is paused to copy the final memory state
- Time required to transfer a VM’s memory depends on its RAM allocation, working set size and write rate, and available bandwidth

Traditional approach to Memory transfer

The default Xen migration algorithm will iterate until either

- very small number of pages remain to be sent
- it has already sent more than 3 times the VM's total memory
- limit of 30 iterations is reached

At that point, the VM is paused, and all remaining pages are sent

Smart Stop and Copy

- Large number of iterations only extends the total migration time and increases the total data transferred
- The migration algorithm could intelligently pick when to stop iterating

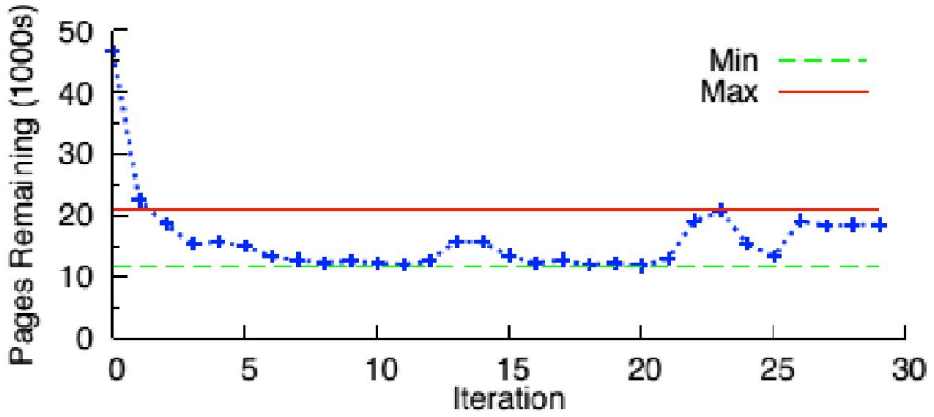

Image courtesy: [3]

Smart Stop and Copy

- A heuristic that tracks the number of pages remaining to be sent
- Looks for local minima in the current history window
- Greedy approach works best for window size 5

Using page Deltas

- After the first iteration, most of the pages are transferred
- Why don’t we send only the changes that happened in the pages
- Reduces bandwidth consumed during migration

Results and improvement

Default Xen code vs CloudNet's optimizations

	Data Tx (GB)		Tot Time (s)		Pause Time (s)	
TPC-W	1.5	0.9	135	78	3.7	2.3
Kernel	1.5	1.1	133	101	5.9	3.5
SPECjbb	1.2	0.4	112	35	7.8	6.5

Image courtesy: [3]

Summary of the paper

- CloudNet Connects resources at multiple data center and enterprise sites
- Handles persistent storage, network connections, and memory state transfer with minimal downtime
- Minimizes the total migration time, application downtime, and the volume of data transferred
- Evaluates how different application types(web/client-server/development) impact migration performance

Cost-efficient live VM migration based on varying electricity cost in optical cloud networks

Abhishek Gupta, Uttam Mandal, Pulak Chowdhury, Massimo Tornatore and Biswanath Mukherjee

The idea

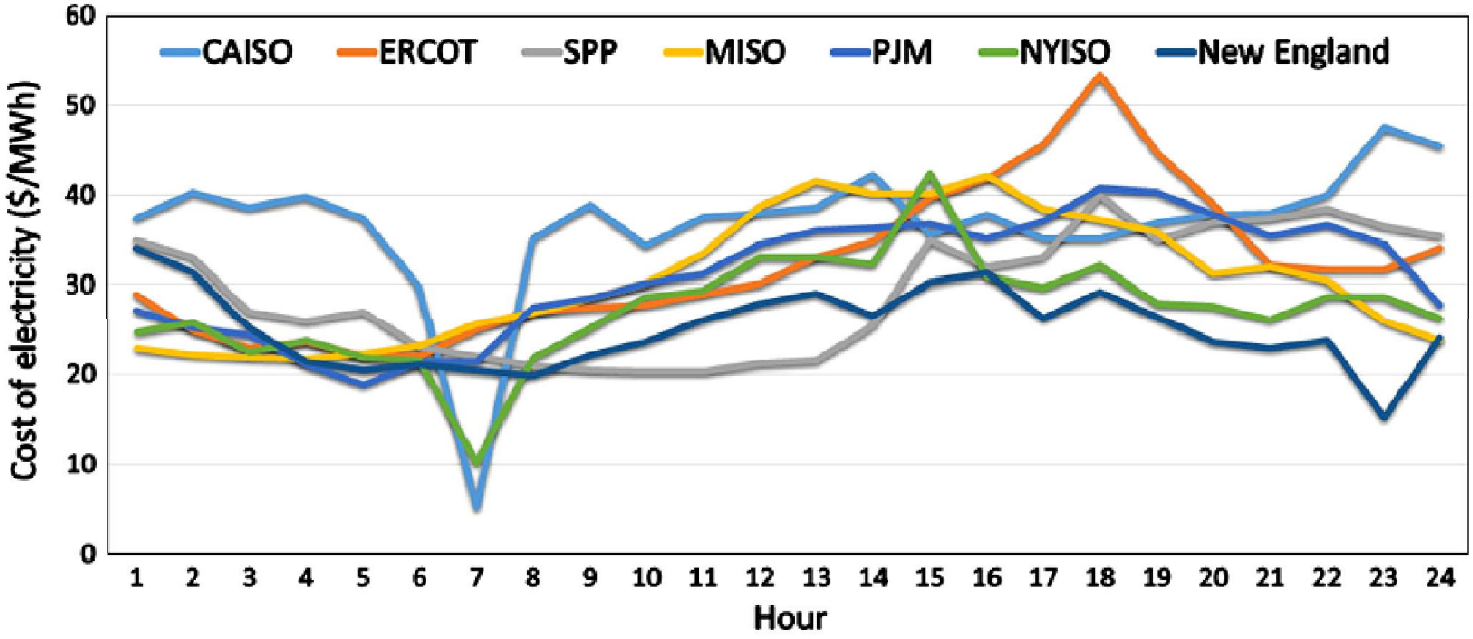

Image courtesy: [4]

The idea

Migration of VMs to Data Centers with cheaper electricity prices

Considering multiple parameters:

- bandwidth for migration
- cost of migration
- duration of migration
- number of servers and racks to be switched on/off

Problem : High level view

- Minimizing the operating cost of VMs
- Live VM migration
- Use of mixed-integer linear program (MILP) formulation

WAN Migration

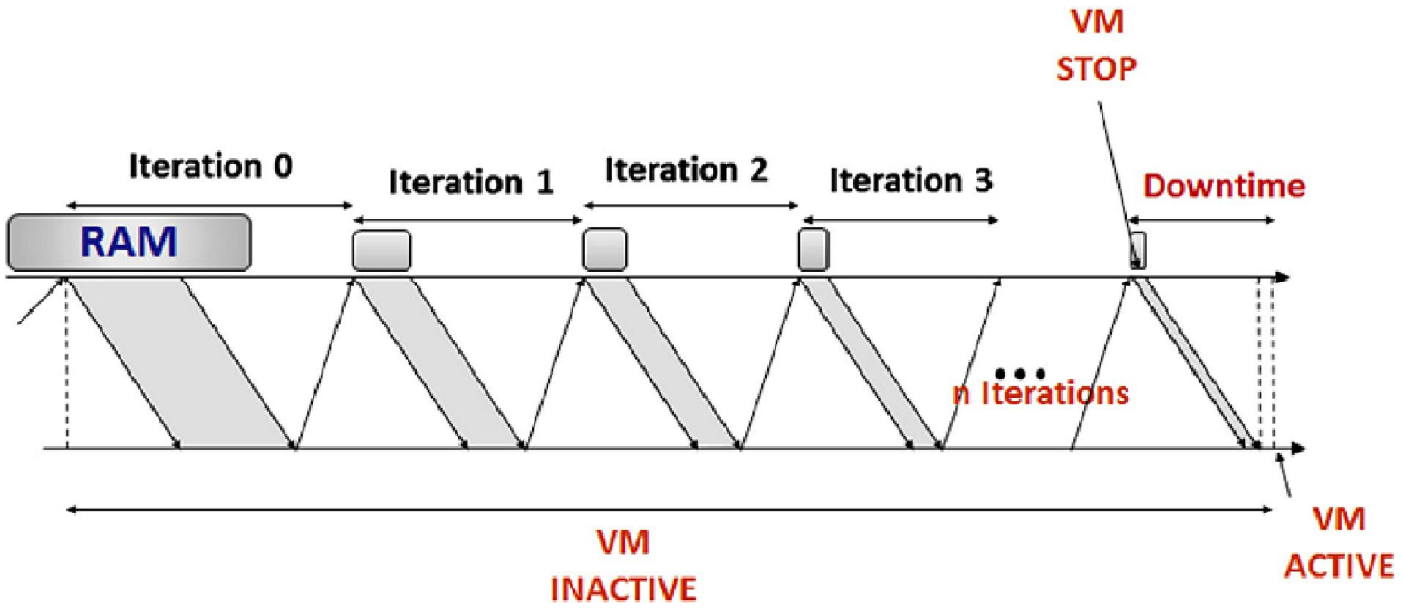

Image courtesy: [4]

Power Model

Power consumptions to consider:

- Source DC power consumption
- Destination DC power consumption
- Data Migration power consumption

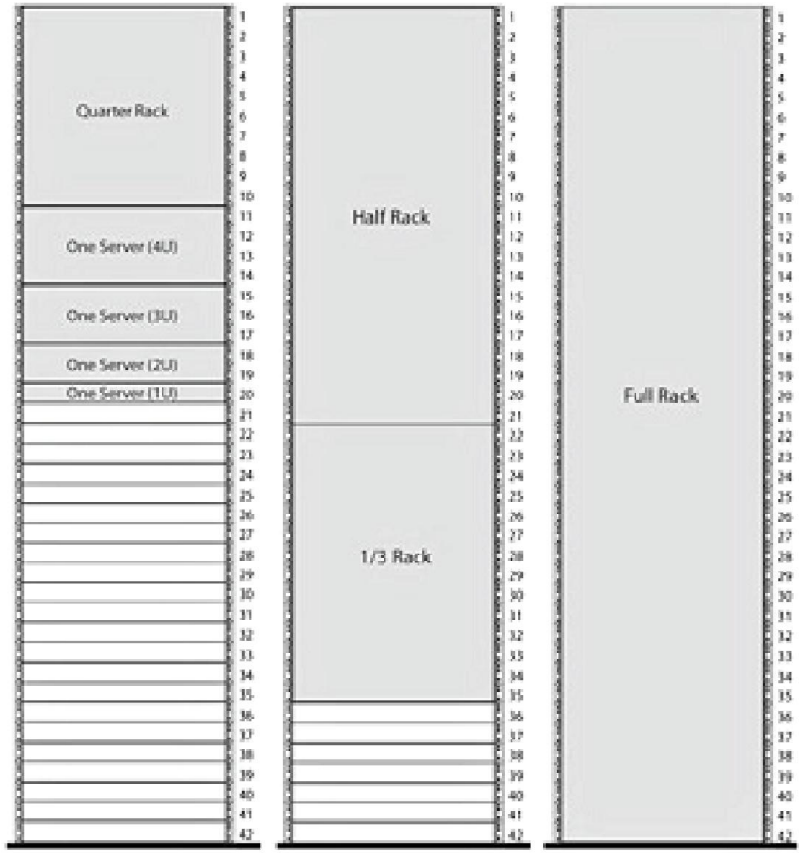

Image courtesy: [4]

Power Model : On rack and Server

Base power in idle state

Additional power ∞ Load

$$P_{DC} = y * P_R + x * P_S \quad \text{where} \quad x \leq \beta * y$$

Power Model : Migration

Migrating a VM from a source DC to a destination DC consume power, in the backbone network

Migration power consumption =

*the total number of bits transferred in a VM migration **

*the power consumed by a core router in transmitting one bit **

the cost of electricity at the core router.

Formal Optimization-problem Statement

Given optical backbone network topology, a set of DC nodes, initial locations of each VM, hourly prices of electricity at each node, link capacities, maximum number of VMs a DC can host, and a multi-hour period, our objective is to minimize the operating cost of the VMs over this period by deciding whether (including when and where) or not to migrate the VMs.

Simulation

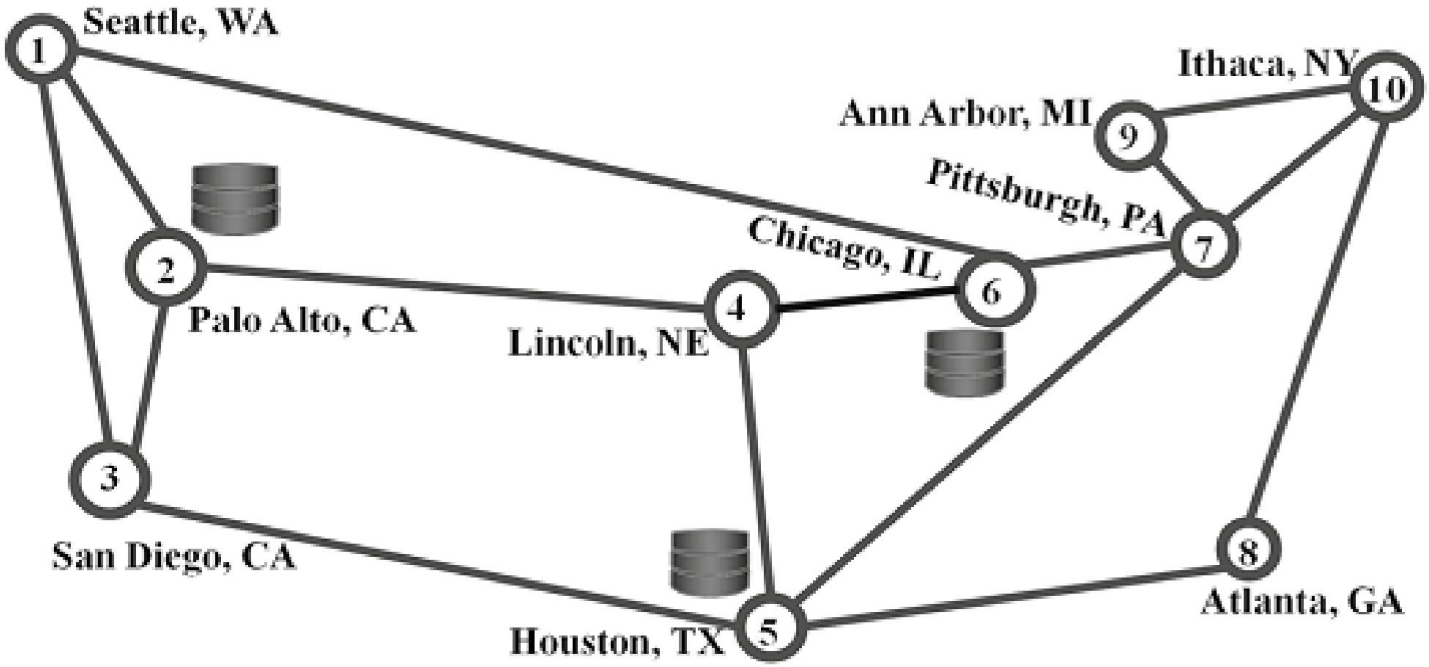

Image courtesy: [4]

VM distribution

Table 1 Given initial VM distribution

VMs	Node 2	Node 5	Node 6
100	36	29	35
200	71	57	72
300	106	96	98

Image courtesy: [4]

Operational cost vs number of VMs

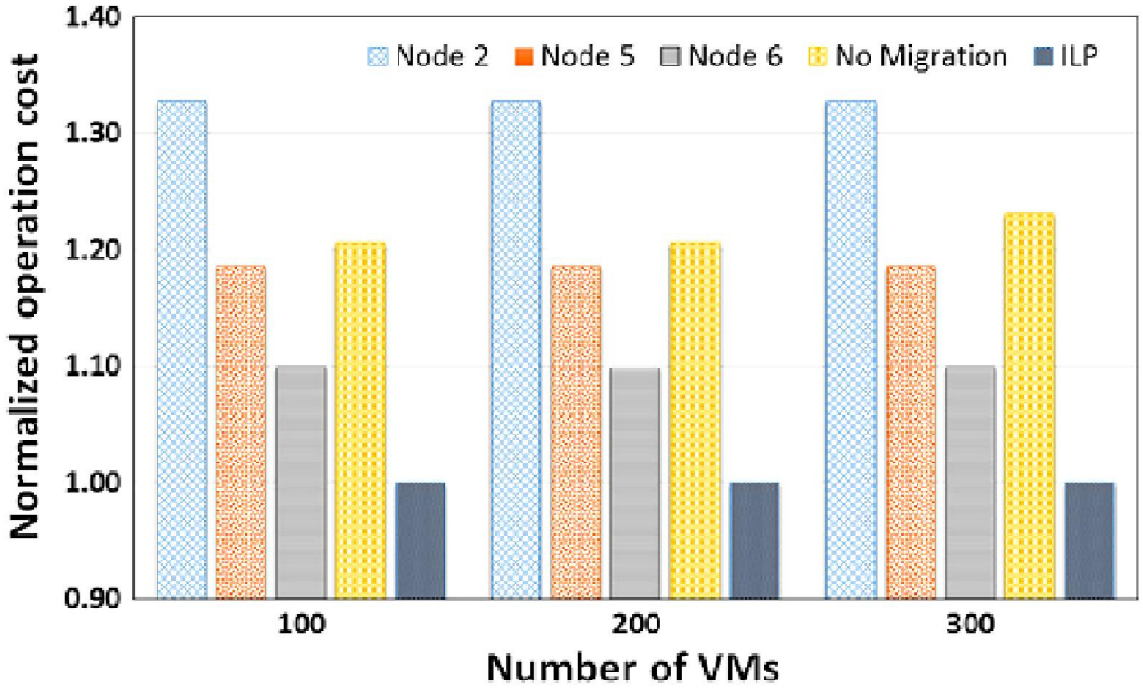

Image courtesy: [4]

24 hour simulation

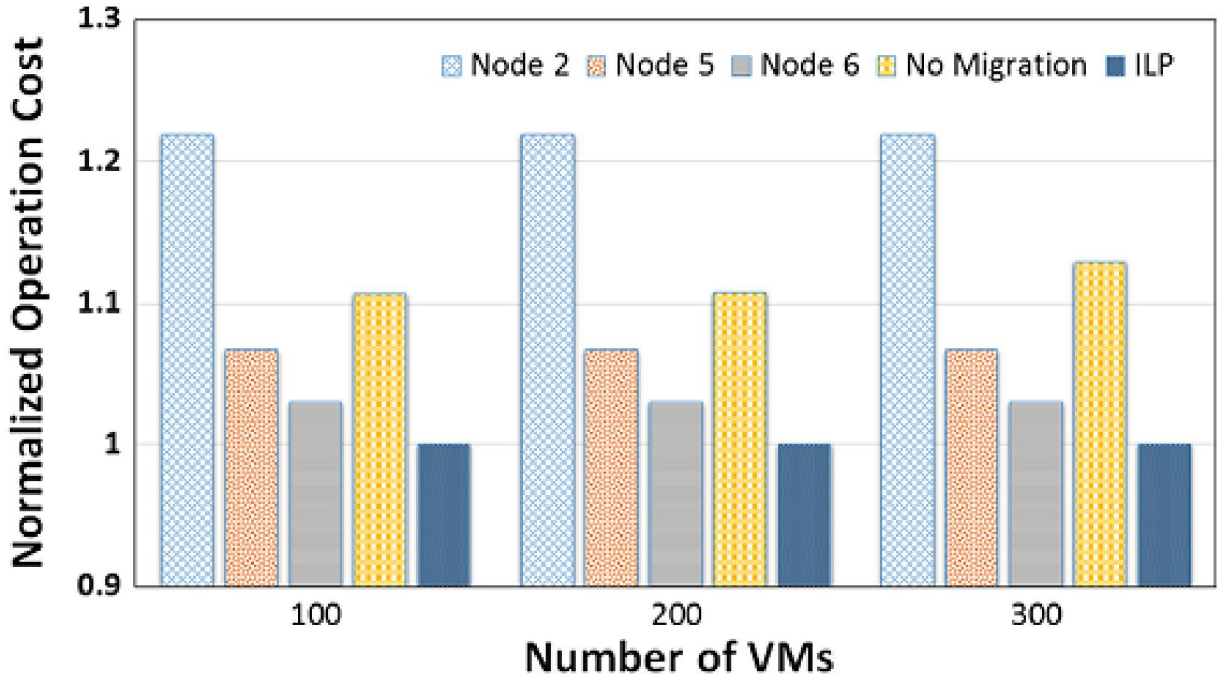

Image courtesy: [4]

Dynamic scenario for the problem

- Service requests coming in Dynamically
- Requests are queued and then served
- Request life can be short or long
- Power modeling can consider more dynamic inputs
- Heterogeneous types and performance of VMs
- Selection of rack to place the VM

Reference

- [1] www.greenstarnetwork.com
- [2] https://sort.symantec.com/public/documents/sfha/6.0/hp-ux/productguides/html/sfha_virtualization/ch03s02.htm
- [3] T.Wood, P. Shenoy, K.K. Ramakrishnan and J. V. der Merwe, ”CloudNet: Dynamic Pooling of Cloud Resources by Live WAN Migration of Virtual Machines”.
- [4] A. Gupta, U. Mandal, P. Chowdhury, M. Tornatore and B. Mukherjee , “Cost-efficient live VM migration based on varying electricity cost in optical cloud networks”.


Thank You