# CLOUD COMPUTING

## SLA - Tutorial

**PROF. SOUMYA K. GHOSH**
**DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING**
**IIT KHARAGPUR**

# What is Service Level Agreement?

- A formal contract between a Service Provider (SP) and a Service Consumer (SC)

- SLA: foundation of the consumer's trust in the provider

- Purpose : to define a formal basis for performance and availability the SP guarantees to deliver

- SLA contains Service Level Objectives (SLOs)
  – Objectively measurable conditions for the service
  – SLA & SLO: basis of selection of cloud provider

# Problem-1

Cloud SLA: Suppose a cloud guarantees service availability for 99% of time. Let a third party application runs in the cloud for 12 hours/day. At the end of one month, it was found that total outage is 10.75 hrs.

Find out whether the provider has violated the initial availability guarantee.

# Problem-2

Consider a scenario where a company X wants to use a cloud service from a provider P. The service level agreement (SLA) guarantees negotiated between the two parties prior to initiating business are as follows:

- Availability guarantee: 99.95% time over the service period
- Service period: 30 days
- Maximum service hours per day: 12 hours
- Cost: $50 per day

Service credits are awarded to customers if availability guarantees are not satisfied. Monthly connectivity uptime service level are given as:

| Monthly Uptime Percentage | Service Credit |
| --- | --- |
| <99.95% | 10% |
| <99% | 25% |

However, in reality in was found that over the service period, the cloud service suffered five outages of durations:

5 hrs, 30 mins, 1 hr 30 mins, 15 mins, and 2 hrs 25 mins, each on different days, due to which normal service guarantees were violated.

If SLA negotiations are honored, compute the effective cost payable towards buying the cloud service.

# Thank You!

# Cloud Computing : Economics Tutorial

**Prof. Soumya K Ghosh**

**Department of Computer Science and Engineering**

**IIT KHARAGPUR**

# Cloud Properties: Economic Viewpoint

- **C**ommon Infrastructure
  - pooled, standardized resources, with benefits generated by statistical multiplexing.

- **L**ocation-independence
  - ubiquitous availability meeting performance requirements, with benefits deriving from latency reduction and user experience enhancement.

- **O**nline connectivity
  - an enabler of other attributes ensuring service access. Costs and performance impacts of network architectures can be quantified using traditional methods.

# Cloud Properties: Economic Viewpoint *(contd...)*

- **U**tility pricing
  - usage-sensitive or pay-per-use pricing, with benefits applying in environments with variable demand levels.

- on-**D**emand Resource**s**
  - scalable, elastic resources provisioned and de-provisioned without delay or costs associated with change.

# Utility Pricing in Detail

| D(t) | demand for resources 0<t<T |
|------|---------------------------|
| P | max (D(t)) : Peak Demand |
| A | Avg (D(t)) : Average Demand |
| B | Baseline (owned) unit cost<br>*[$B_T$ : Total Baseline Cost]* |
| C | Cloud unit cost<br>*[$C_T$ : Total Cloud Cost]* |
| U (=C/B) | Utility Premium<br>*[For rental car example, U=4.5]* |

$$C_T = \int_0^T U \times B \times D(t)\,dt = A \times U \times B \times T$$

$$B_T = P \times B \times T$$

- Because the baseline should handle peak demand

When is cloud cheaper than owning?

$$C_T < B_T \rightarrow A \times U \times B \times T < P \times B \times T$$

$$\rightarrow U < \frac{P}{A}$$

- When utility premium is less than ratio of peak demand to Average demand

# Utility Pricing in Real World

- In practice demands are often highly spiky
  - News stories, marketing promotions, product launches, Internet flash floods, Tax season, Christmas shopping, etc.
- Often a hybrid model is the best
  - You own a car for daily commute, and rent a car when traveling or when you need a van to move
  - Key factor is again the ratio of peak to average demand
  - But we should also consider other costs
    - Network cost (both fixed costs and usage costs)
    - Interoperability overhead
    - Consider Reliability, accessibility

# Value of on-Demand Services

- Simple Problem: When owning your resources, you will pay a penalty whenever your resources do not match the instantaneous demand
    I. Either pay for unused resources, or suffer the penalty of missing service delivery

D(t) – Instantaneous Demand at time t

R(t) – Resources at time t

$$\text{Penalty Cost } \alpha \int |D(t) - R(t)| dt$$

- *If demand is flat, penalty = 0*

- *If demand is linear periodic provisioning is acceptable*

# Penalty Costs for Exponential Demand

- Penalty cost $\propto \int |D(t) - R(t)| dt$

- If demand is exponential ($D(t)=e^t$), any fixed provisioning interval ($t_p$) according to the current demands will fall exponentially behind

- $R(t) = e^{t-t_p}$

- $D(t) - R(t) = e^t - e^{t-t_p} = e^t(1 - e^{t_p}) = k_1 e^t$

- Penalty cost $\propto c.k_1 e^t$

Demand $D(t) = e^t$

$t_p$

Unserved Demand

Resources $R(t)=e^{(t-t_p)}$

$t_p$     $t_1$     $t_2$     Time

Exponential Growth with Continuous Monitoring And Non-Zero Provisioning Interval

IIT KHARAGPUR

NPTEL ONLINE
CERTIFICATION COURSES
NPTEL

# Assignment 1

Consider the peak computing demand for an organization is 120 units. The demand as a function of time can be expressed as:

$$D(t) = \begin{cases} 50\sin(t), & 0 \le t < {}^{\pi}/_{2} \\ 20\sin(t), & {}^{\pi}/_{2} \le t < \pi \end{cases}$$

The resource provisioned by the cloud to satisfy current demand at time t is given as:

$$R(t) = D(t) + \delta.\left(\frac{dD(t)}{dt}\right)$$

where, $\delta$ is the delay in provisioning the extra computing recourse on demand

The cost to provision unit cloud resource for unit time is 0.9 units.
Calculate the penalty.
*[Assume the delay in provisioning is ${}^{\pi}/_{12}$ time units and minimum demand is 0]*
*(Penalty: Either pay for unused resource or missing service delivery)*

# Assignment 2

Consider that the peak computing demand for an organization is *100 units*.
The demand as a function of time can be expressed as

$$D(t) = 50(1 + e^{-t})$$

Baseline (owned) unit cost is *120* and cloud unit cost is *200*.

In this situation is cloud cheaper than owning for a period of *100* time units?

# Assignment 3

A company X needs to support a spike in demand when it becomes popular, followed potentially by a reduction once some of the visitors turn away. The company has two options to satisfy the requirements which are given in the following table:

| Expenditures | In-house server (INR) | Cloud server |
|---|---|---|
| Purchase cost | 6,00,000 | - |
| Number of CPU cores | 12 | 8 |
| Cost/hour (over three year span) | - | 42 |
| Efficiency | 40% | 80% |
| Power and cooling (cost/hour) | 22 | - |
| Management cost (cost/hour) | 6 | 1 |

- Calculate the price of a core-hour on in-house server and cloud server.

- Find the cost/effective-hour for both the options.

- Calculate the ratio of the total cost/effective-hour for in-house to cloud deployment.

- If the efficiency of in-house server is increased to 70%, which deployment will have now better total cost/effective-hour?

# Cloud Computing :
## *MapReduce - Tutorial*

**Prof. Soumya K Ghosh**

**Department of Computer Science and Engineering**

**IIT KHARAGPUR**

# Introduction

- MapReduce: programming model developed at Google

- Objective:
  - Implement large scale search
  - Text processing on massively scalable web data stored using BigTable and GFS distributed file system

- Designed for processing and generating large volumes of data via massively parallel computations, utilizing tens of thousands of processors at a time

- Fault tolerant: ensure progress of computation even if processors and networks fail

- Example:
  - Hadoop: open source implementation of MapReduce (developed at Yahoo!)
  - Available on pre-packaged AMIs on Amazon EC2 cloud platform

# MapReduce Model

- Parallel programming abstraction

- Used by many different parallel applications which carry out large-scale computation involving thousands of processors

- Leverages a common underlying fault-tolerant implementation

- Two phases of MapReduce:
  - Map operation
  - Reduce operation

- A configurable number of M 'mapper' processors and R 'reducer' processors are assigned to work on the problem

- The computation is coordinated by a single master process

# MapReduce Model Contd…

- Map phase:
    - Each mapper reads approximately *1/M* of the input from the global file system, using locations given by the master
    - Map operation consists of transforming one set of key-value pairs to another:

    $$\text{Map:} \quad (k_1, v_1) \rightarrow [(k_2, v_2)].$$

    - Each mapper writes computation results in one file per reducer
    - Files are sorted by a key and stored to the local file system
    - The master keeps track of the location of these files

# MapReduce Model Contd...

- **Reduce phase:**
  - The master informs the reducers where the partial computations have been stored on local files of respective mappers
  - Reducers make remote procedure call requests to the mappers to fetch the files
  - Each reducer groups the results of the map step using the same key and performs a function $f$ on the list of values that correspond to these key value:

$$\text{Reduce:} \quad (k_2, [v_2]) \rightarrow (k_2, f([v_2])).$$

  - Final results are written back to the GFS file system

# MapReduce: Example
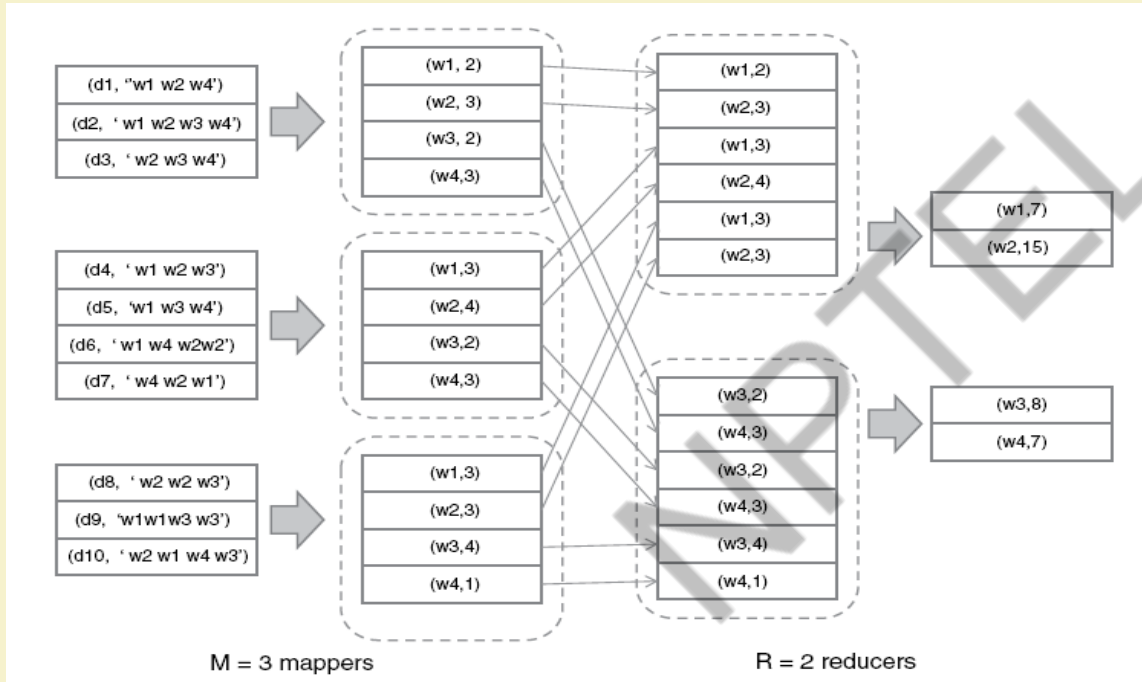


- 3 mappers; 2 reducers
- Map function:

$$(d_k, [w_1 \ldots w_n]) \rightarrow [(w_i, c_i)].$$

- Reduce function:

$$(w_i, [c_i]) \rightarrow \left(w_i, \sum_i c_i\right)$$

# Problem-1

In a MapReduce framework consider the HDFS block size is 64 MB. We have 3 files of size 64K, 65Mb and 127Mb. How many blocks will be created by Hadoop framework?

# Problem-2

Write the pseudo-codes (for map and reduce functions) for calculating the average of a set of integers in MapReduce.

Suppose A = (10, 20, 30, 40, 50) is a set of integers. Show the map and reduce outputs.

# Problem-3

Compute total and average salary of organization XYZ and group by based on gender (male or female) using MapReduce. The input is as follows

*Name, Gender, Salary*

*John, M, 10,000*

*Martha, F, 15,000*

*----*

NPTEL ONLINE
CERTIFICATION COURSES

# Problem-4

Write the *Map* and *Reduce* functions (pseudo-codes) for the following **Word Length Categorization** problem under *MapReduce* model.

*Word Length Categorization*: Given a text paragraph (containing only words), categorize each word into following categories. Output the frequency of occurrence of words in each category.

Categories:
 **tiny:** 1-2 letters;  **small:** 3-5 letters;   **medium:** 6-9 letters;   **big:** 10 or more letters

# Different Resources in Computing



Source: http://www.cse.hcmut.edu.vn/~ptvu/gc/2012/GC-pp.pdf

# Resources types

- **Physical resource**
  - ❑ Computer, disk, database, network, scientific instruments.

- **Logical resource**
  - ❑ Execution, monitoring, communicate application .

IIT KHARAGPUR

NPTEL ONLINE
CERTIFICATION COURSES

# Resources Management

- The term ***resource management*** refers to the operations used to control how capabilities provided by Cloud resources and services cane be made available to other entities, whether users, applications, services in an *efficient* manner.

# Data Center Power Consumption

- Currently it is estimated that servers consume 0.5% of the world's total electricity usage.

- Server energy demand doubles every 5-6 years.

- This results in large amounts of $CO_2$ produced by burning fossil fuels.

- Need to reduce the energy used with minimal performance impact.

*Ref: Efficient Resource Management for Cloud Computing Environments, by Andrew J. Younge, Gregor von Laszewski, Lizhe Wang, Sonia Lopez-Alarcon, Warren Carithers,*

# Motivation for Green Data Centers

**Economic**

- New data centers run on the Megawatt scale, requiring millions of dollars to operate.

- Recently institutions are looking for new ways to reduce costs

- Many facilities are at their peak operating stage, and cannot expand without a new power source.

**Environmental**

- Majority of energy sources are fossil fuels.

- Huge volume of $CO_2$ emitted each year from power plants.

- Sustainable energy sources are not ready.

- Need to reduce energy dependence

# Green Computing ?

- Advanced scheduling schemas to reduce energy consumption.

  - Power aware

  - Thermal aware

- Performance/Watt is not following Moore's law.

- Data center designs to reduce Power Usage Effectiveness.

  - Cooling systems

  - Rack design

# Research Directions

How to conserve energy within a Cloud environment.

- Schedule VMs to conserve energy.

- Management of both VMs and underlying infrastructure.

- Minimize operating inefficiencies for non-essential tasks.

- Optimize data center design.

# Steps towards Energy Efficiency
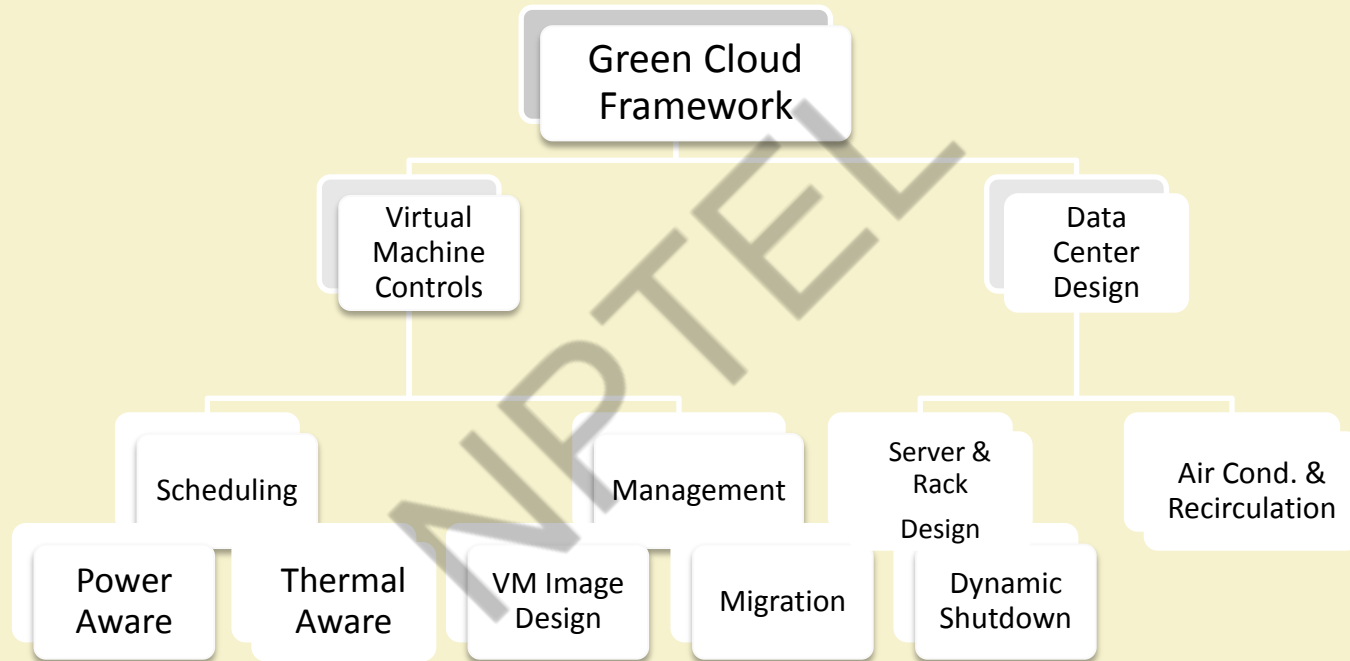
# VM scheduling on Multi-core Systems

- There is a nonlinear relationship between the number of processes used and power consumption

- We can schedule VMs to take advantage of this relationship in order to conserve power

Scheduling



*Power consumption curve on an Intel Core i7 920 Server (4 cores, 8 virtual cores with Hyperthreading)*

# Power-aware Scheduling

- Schedule as many VMs at once on a multi-core node.

  - Greedy scheduling algorithm

  - Keep track of cores on a given node

  - Match VM requirements with node capacity

**Algorithm 1** Power based scheduling of VMs

FOR $i = 1$ TO $i \leq |pool|$ DO
  $pe_i$ = num cores in $pool_i$
END FOR

WHILE ($true$)
  FOR $i = 1$ TO $i \leq |queue|$ DO
  $vm = queue_i$
    FOR $j = 1$ TO $j \leq |pool|$ DO
    IF $pe_j \geq 1$ THEN
      IF check capacity $vm$ on $pe_j$ THEN
        schedule $vm$ on $pe_j$
        $pe_j - 1$
      END IF
    END IF
    END FOR
  END FOR
  wait for interval $t$
END WHILE

Scheduling

# 485 Watts vs. 552 Watts !



Node 1 @ 170W

Node 2 @ 105W

Node 3 @ 105W

Node 4 @ 105W

VS.

Node 1 @ 138W
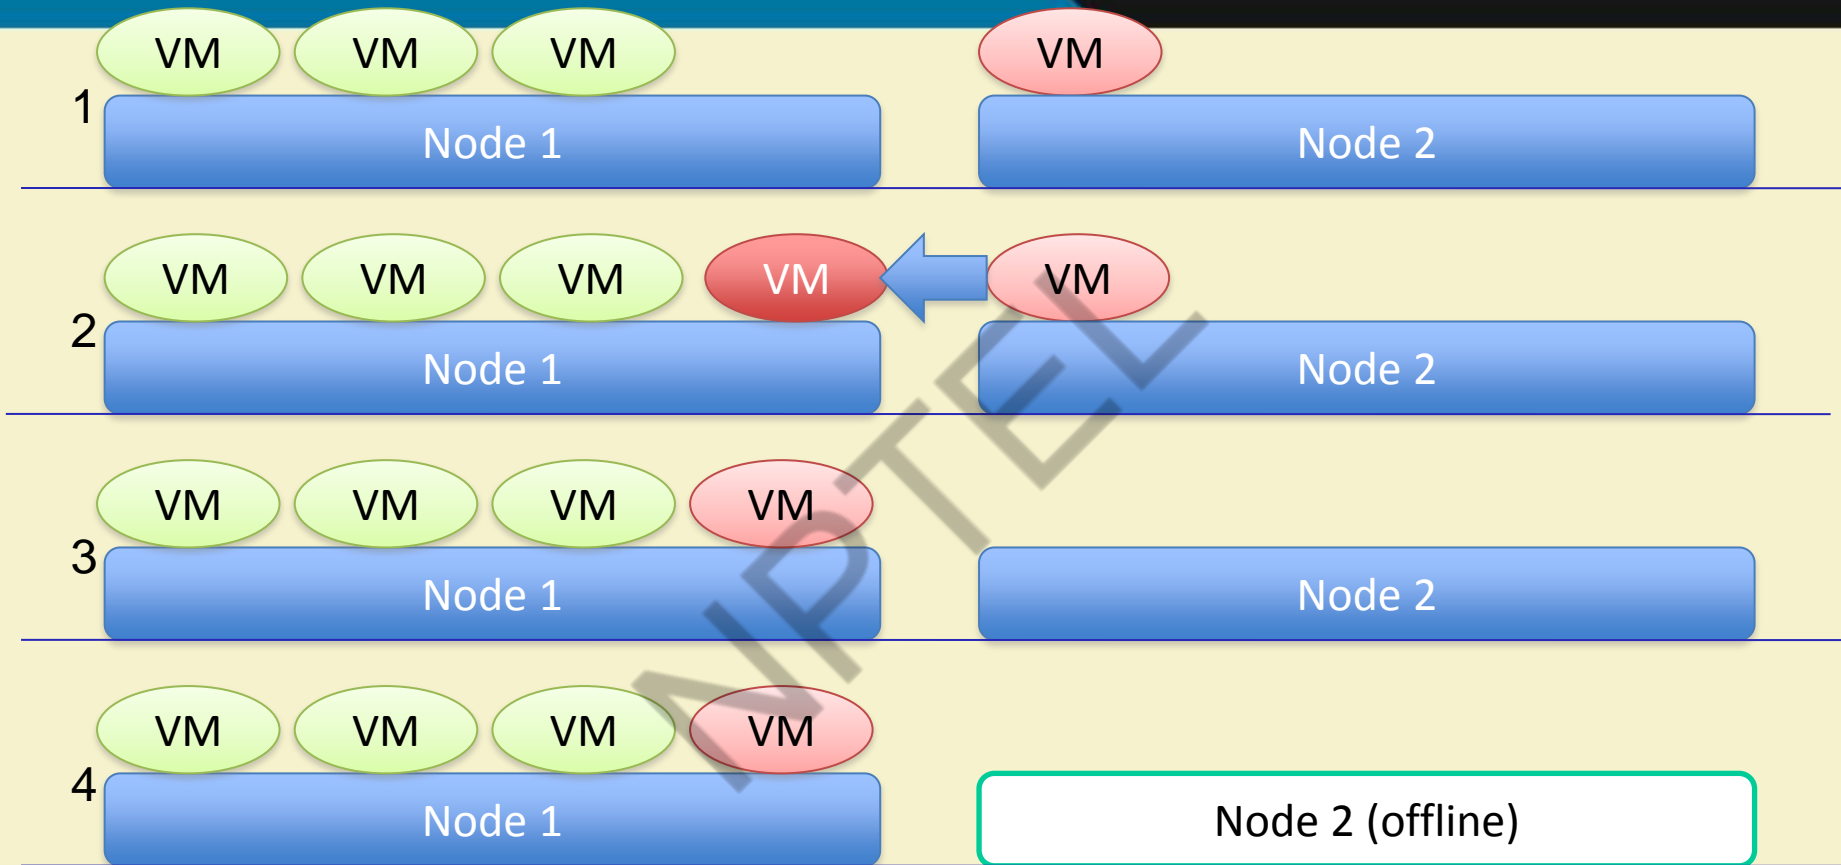
Node 2 @ 138W

Node 3 @ 138W

Node 4 @ 138W

# VM Management

- Monitor Cloud usage and load.

- When load decreases:

    - Live migrate VMs to more utilized nodes.

    - Shutdown unused nodes.

- When load increases:

    - Use WOL to start up waiting nodes.

    - Schedule new VMs to new nodes.

Management

# Minimizing VM  Instances

- Virtual machines are loaded!

  - Lots of unwanted packages.

  - Unneeded services.

- Are multi-application oriented, not service oriented.

  - Clouds are based off of a Service Oriented Architecture.

- Need a custom lightweight Linux VM for service oriented science.

- Need to keep VM image as small as possible to reduce network latency.

Management

# Typical Cloud Linux Image

- Start with Ubuntu 9.04.
- Remove all packages not
- required for base image.
  - No X11
  - No Window Manager
  - Minimalistic server install
  - Can load language support on demand (via package manager)
- Readahead profiling utility.
  - Reorder boot sequence
  - Pre-fetch boot files on disk
  - Minimize CPU idle time due to I/O delay
- Optimize Linux kernel.
  - Built for Xen DomU
  - No 3d graphics, no sound, minimalistic kernel
  - Build modules within kernel directly

**VM Image Design**

**Boot chart for ubuntu-minimal (Fri May 8 15:01:26 EDT 2009)**
uname: Linux 2.6.28-11-generic #42-Ubuntu SMP Fri Apr 17 01:58:03 UTC 2009 x86_64
release: Ubuntu 9.04
CPU: Intel(R) Core(TM)2 Duo CPU    T9300  @ 2.50GHz (1)
kernel options: root=UUID=042a98cc-dab1-4c5d-a45f-9088b7067ad9 ro quiet splash quiet
time: 0:08

CPU (user+sys)    I/O (wait)

Disk throughput    Disk utilization
158 MB/s

Running (%cpu)    Unint.sleep (I/O)

# Energy Savings

- Reduced boot times from 38 seconds to just **8** seconds.
  - 30 seconds @ 250Watts is 2.08wh or .002kwh.
- In a small Cloud where 100 images are created every hour.
  - Saves .2kwh of operation @ 15.2c per kwh.
  - At 15.2c per kwh this saves $262.65 every year.
- In a production Cloud where 1000 images are created every minute.
  - Saves 120kwh less every hour.
  - At 15.2c per kwh this saves over 1 million dollars every year.
- Image size from 4GB to 635MB.
  - Reduces time to perform live-migration.
  - Can do better.

# Summary - 1

- Cloud computing is an emerging topic in Distributed Systems.

- Need to conserve energy wherever possible!

- Green Cloud Framework:

  - Power-aware scheduling of VMs.

  - Advanced VM & infrastructure management.

  - Specialized VM Image.

- Small energy savings result in a large impact.

- Combining a number of different methods together can have a larger impact then when implemented separately.

# Summary - 2

- Combine concepts of both Power-aware and Thermal-aware scheduling to minimize both energy and temperature.

- Integrated server, rack, and cooling strategies.

- Further improve VM Image minimization.

- Designing the next generation of Cloud computing systems to be more efficient.

# Thank you!

# CLOUD COMPUTING

## Resource Management - II

**PROF. SOUMYA K. GHOSH**
**DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING**
**IIT KHARAGPUR**

# Different Resources in Computing
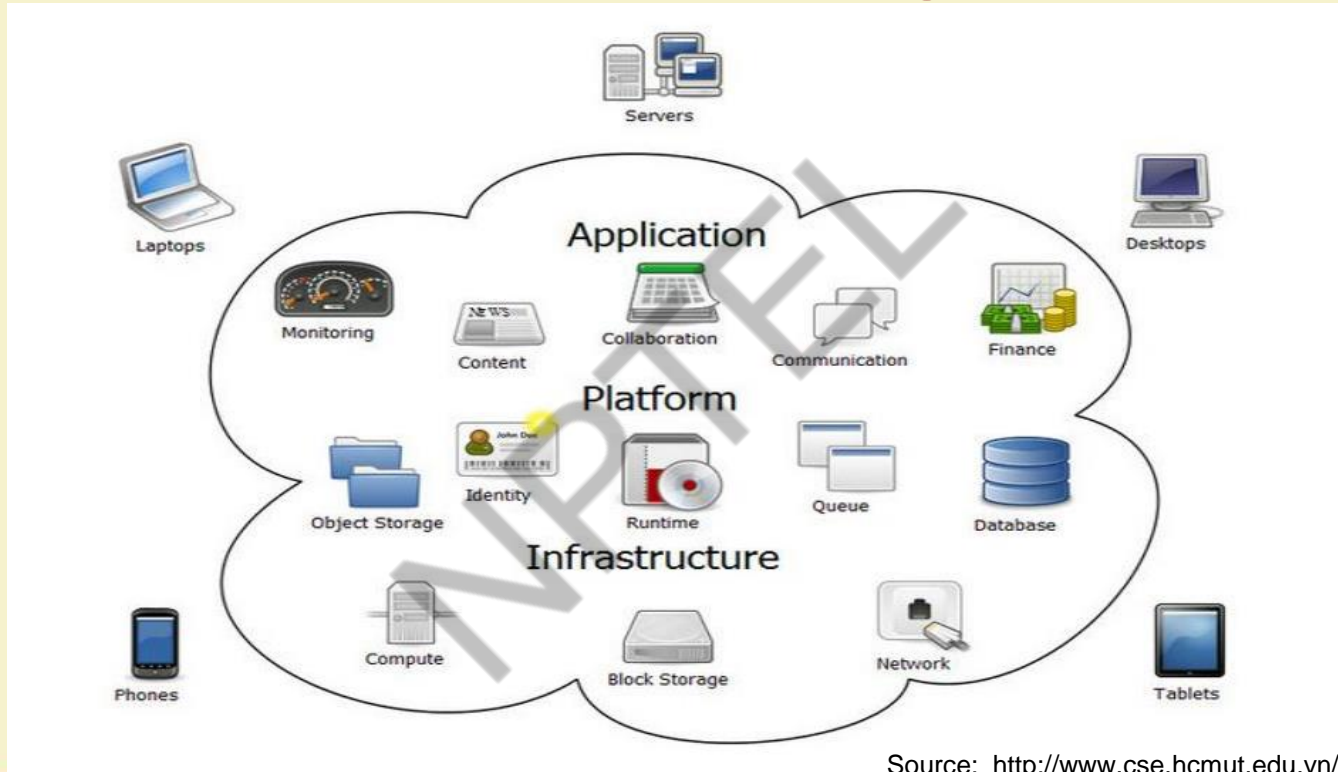


Source: http://www.cse.hcmut.edu.vn/~ptvu/gc/2012/GC-pp.pdf

# Resources types

- **Physical resource**
  - ❑ Computer, disk, database, network, scientific instruments.

- **Logical resource**
  - ❑ Execution, monitoring, communicate application .

IIT KHARAGPUR | NPTEL ONLINE CERTIFICATION COURSES

# Resources Management

- The term ***resource management*** refers to the operations used to control how capabilities provided by Cloud resources and services cane be made available to other entities, whether users, applications, services in an *efficient* manner.

# Resource Management for IaaS

- Infrastructure-as-a-Service (IaaS) is most popular cloud service

- In IaaS, cloud providers offer resources that include computers as virtual machines, raw (block) storage, firewalls, load balancers, and network devices.

- One of the major challenges in IaaS is resource management.

Source:
http://www.zearon.com/down/Resource%20management%20for%20Infrastructure%20as%20a%20Service%20%28IaaS%29%20in%20cloud%20computing%20A%20survey.pdf

# Resource Management - Objectives

- Scalability
- Quality of service
- Optimal utility
- Reduced overheads
- Improved throughput
- Reduced latency
- Specialized environment
- Cost effectiveness
- Simplified interface

# Resource Management – Challenges (Hardware)

- CPU (central processing unit)

- Memory

- Storage

- Workstations

- Network elements

- Sensors/actuators

# Resource Management – Challenges (Logical resources)

- Operating system

- Energy

- Network throughput/bandwidth

- Load balancing mechanisms

- Information security

- Delays

- APIs/(Applications Programming Interfaces)

- Protocols

# Resource Management Aspects

- Resource provisioning
- Resource allocation
- Resource requirement mapping
- Resource adaptation
- Resource discovery
- Resource brokering
- Resource estimation
- Resource modeling

# Resource Management

| Type | Details |
|---|---|
| Resource provisioning | Allocation of a service provider's resources to a customer |
| Resource allocation | Distribution of resources economically among competing groups of people or programs |
| Resource adaptation | Ability or capacity of that system to adjust the resources dynamically to fulfill the requirements of the user |
| Resource mapping | Correspondence between resources required by the users and resources available with the provider |
| Resource modeling | Resource modeling is based on detailed information of transmission network elements, resources and entities participating in the network. Attributes of resource management: states, transitions, inputs and outputs within a given environment. Resource modeling helps to predict the resource requirements in subsequent time intervals |
| Resource estimation | A close guess of the actual resources required for an application, usually with some thought or calculation involved |
| Resource discovery and selection | Identification of list of authenticated resources that are available for job submission and to choose the best among them |
| Resource brokering | It is the negotiation of the resources through an agent to ensure that the necessary resources are available at the right time to complete the objectives |
| Resource scheduling | A resource schedule is a timetable of events and resources. Shared resources are available at certain times and events are planned during these times. In other words, It is determining when an activity should start or end, depending on its (1) duration, (2) predecessor activities, (3) predecessor relationships, and (4) resources allocated |

# Resource Provisioning Approaches

| Approach | Description |
|---|---|
| Nash equilibrium approach using Game theory | Run time management and allocation of IaaS resources considering several criteria such as the heterogeneous distribution of resources, rational exchange behaviors of cloud users, incomplete common information and dynamic successive allocation |
| Network queuing model | Presents a model based on a network of queues, where the queues represent different tiers of the application.The model sufficiently captures the behavior of tiers with significantly different performance characteristics and application idiosyncrasies, such as, session-based workloads, concurrency limits, and caching at intermediate tiers |
| Prototype provisioning | Employs the k-means clustering algorithm to automatically determine the workload mix and a queuing model to predict the server capacity for a given workload mix. |
| Resource (VM) provisioning | Uses virtual machines (VMs) that run on top of the Xen hypervisor. The system provides a Simple Earliest Deadline First (SEDF) scheduler that implements weighted fair sharing of the CPU capacity among all the VMs<br>The share of CPU cycles for a particular VM can be changed at runtime |
| Adaptive resource provisioning | Automatic bottleneck detection and resolution under dynamic resource management which has the potential to enable cloud infrastructure providers to provide SLAs for web applications that guarantee specific response time requirements while minimizing resource utilization. |
| SLA oriented methods | Handling the process of dynamic provisioning to meet user SLAs in autonomic manner. Additional resources are provisioned for applications when required and are removed when they are not necessary |
| Dynamic and automated framework | A dynamic and automated framework which can adapt the adaptive parameters to meet the specific accuracy goal, and then dynamically converge to near-optimal resource allocation to handle unexpected changes |
| Optimal cloud resource provisioning (OCRP) | The demand and price uncertainty is considered using optimal cloud resource provisioning (OCRP) including deterministic equivalent formulation, sample-average approximation, etc. |

# Resource Allocation Approaches

| Approach | Description |
|---|---|
| Market-oriented resource allocation | Considers the case of a single cloud provider and address the question how to best match customer demand in terms of both supply and price in order to maximize the providers revenue and customer satisfactions while minimizing energy cost. In particular, it models the problem as a constrained discrete-time optimal control problem and uses Model Predictive Control(MPC) to find its solution |
| Intelligent multi-agent model | An intelligent multi-agent model based on virtualization rules for resource virtualization to automatically allocate service resources suitable for mobile devices. It infers user demand by analyzing and learning user context information. |
| Energy-Aware Resource allocation | Resource allocation is carried out by mimicking the behavior of ants, that the ants are likely to choose the path identified as a shortest path, which is indicated by a relatively higher density of pheromone left on the path compared to other possible paths |
| Measurement based analysis on performance | Focuses on measurement based analysis on performance impact of co-locating applications in a virtualized cloud in terms of throughput and resource sharing effectiveness, including the impact of idle instances on applications that are running concurrently on the same physical host |
| Dynamic resource allocation method | Dynamic resource allocation method based on the load of VMs on IaaS, which enables users to dynamically add and/or delete one or more instances on the basis of the load and the conditions specified by the user |
| Real time resource allocation mechanism | Designed for helping small and medium sized IaaS cloud providers to better utilize their hardware resources with minimum operational cost by a well-designed underlying hardware infrastructure, an efficient resource scheduling algorithm and a set of migrating operations of VMs |
| Dynamic scheduling and consolidation mechanism | Presents the architecture and algorithmic blueprints of a framework for workload co-location, which provides customers with the ability to formally express workload scheduling flexibilities using Directed Acyclic Graphs (DAGs), and optimizes the use of cloud resources to collocate client's workloads |

# Resource Mapping Approaches

| Approach | Description |
|---|---|
| Symmetric mapping pattern | Symmetric mapping pattern for the design of resource supply systems. It divides resource supply in three functions: (1) users and providers match and engage in resource supply agreements, (2) users place tasks on subscribed resource containers, and (3) providers place supplied resource containers on physical resources |
| Load-aware mapping | Explores how to simplify VM image management and reduce image preparation overhead by the multicast file transferring and image caching/reusing. Load-Aware Mapping to further reduce deploying overhead and make efficient use of resources. |
| Minimum congestion mapping | Framework for solving a natural graph mapping problem arising in cloud computing. Applying this framework to obtain offline and online approximation algorithms for workloads given by depth-d trees and complete graphs |
| Iterated local search based request partitioning | Request partitioning approach based on iterated local search is introduced that facilitates the cost- efficient and on-line splitting of user requests among eligible Cloud Service Providers (CSPs) within a networked cloud environment |
| SOA API | Designed to accept different resource usage prediction models and map QoS constraints to resources from various IaaS providers |
| Impatient task mapping | Batch mapping via genetic algorithms with throughput as a fitness function that can be used to map jobs to cloud resources |
| Distributed ensembles of virtual appliances (DEVAs) | Requirements are inferred by observing the behavior of the system under different conditions and creating a model that can be later used to obtain approximate parameters to provide the resources. |
| Mapping a virtual network onto a substrate network | An effective method (using backbone mapping) for computing high quality mappings of virtual networks onto substrate networks. The computed virtual networks are constructed to have sufficient capacity to accommodate any traffic pattern allowed by user-specified traffic constraints. |

# Resource Adaptation Approaches

| Approach | Description |
|---|---|
| Reinforcement learning guided control policy | A multi-input multi-output feedback control model-based dynamic resource provisioning algorithm which adopts reinforcement learning to adjust adaptive parameters to guarantee the optimal application benefit within the time constraint |
| Web-service based prototype | A web-service based prototype framework, and used it for performance evaluation of various resource adaptation algorithms under different realistic settings |
| OnTimeMeasure service | Presents an application – adaptation case study that uses OnTimeMeasure-enabled performance intelligence in the context of dynamic resource allocation within thin-client based virtual desktop clouds to increase cloud scalability, while simultaneously delivering satisfactory user quality-of-experience |
| Virtual networks | Proposes virtual networks architecture as a mechanism in cloud computing that can aggregate traffic isolation, improving security and facilitating pricing, also allowing customers to act in cases where the performance is not in accordance with the contract for services |
| DNS-based Load Balancing | Proposes a system that contain the appropriate elements so that applications can be scaled by replicating VMs (or application containers), by reconfiguring them on the fly, and by adding load balancers in front of these replicas that can scale by themselves |
| Hybrid approach | Proposes a mechanism for providing dynamic management in virtualized consolidated server environments that host multiple multi-tier applications using layered queuing models for Xen-based virtual machine environments, which is a novel optimization technique that uses a combination of bin packing and gradient search |

# Performance Metrics for Resource Management

- Reliability

- Ease  of deployment

- QoS

- Delay

- Control overhead

# Thank you!