PROM in Clouds: Exploring PeRformance OptiMization in Clouds

Chita R. Das

Computer Science & Engineering Department
The Pennsylvania State University
University Park, PA 16802
Email: das@cse.psu.edu

NAS Keynote: July18, 2103

Pope's Inauguration

Then...

"When smartphones and tablets

light up the sky,

load up the clouds."

Now...

Source : http://www.alternet.org/speakeasy/alyssa-figueroa/recording-memories-why-must-we-capture-our-every-moment

Era of Internet and Cloud

*Source: Intel 2012

Cloud Computing

"Cloud" refers to both the *services* delivered over the Internet and the *infrastructure*.

- Pay-as-you-go model of utility computing
- Economy of scale
- Elasticity
- On-demand remote access to shared infrastructure
- Virtualized
- Popularity and usage is growing…

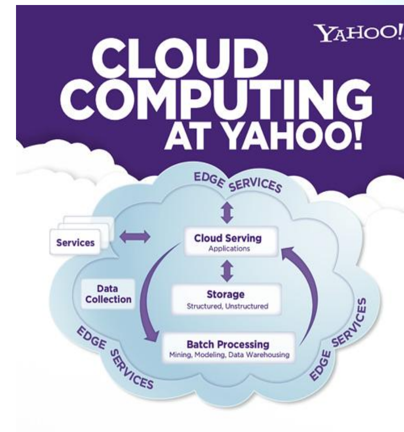

d Stack

Increased provider automation

Gmail

IBM SmartCloud

Salesforce

Software as a Service (SaaS)

Windows Azure

Google App Engine

Platform as a Service (PaaS)

amazon web services™

Eucalyptus Systems

Infrastructure as a Service (IaaS)

Increased end-user control

Cloud Computing Growth

Figure 3 Forecast: Global Public Cloud Market Size, 2011 To 2020

The spreadsheet detailing this forecast is available online.

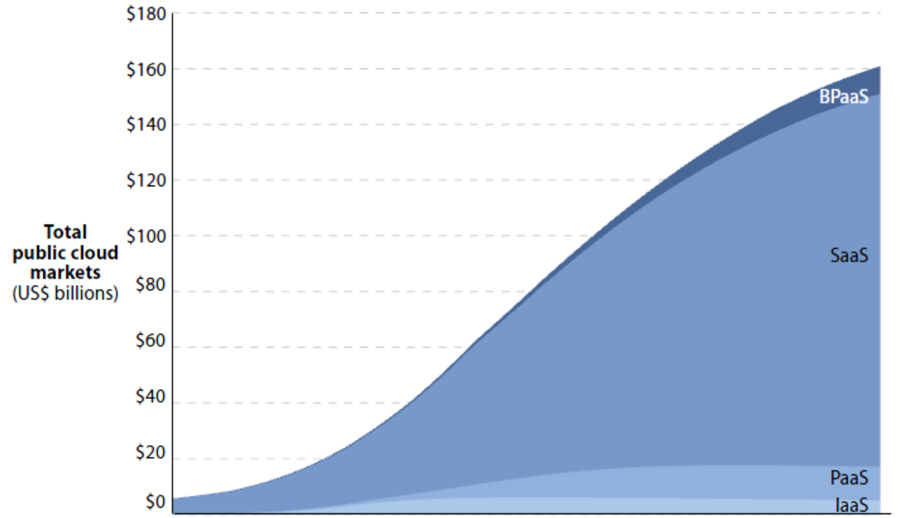

	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
BPaaS ($)	0.15	0.23	0.35	0.53	0.80	1.26	1.95	2.93	4.28	6.00	7.66	9.08	10.02
SaaS ($)	5.56	8.09	13.40	21.21	33.09	47.22	63.19	78.43	92.75	105.49	116.39	125.52	132.57
PaaS ($)	0.05	0.12	0.31	0.82	2.08	4.38	7.39	9.80	11.26	11.94	12.15	12.10	11.91
IaaS ($)	0.06	0.24	1.02	2.94	4.99	5.75	5.89	5.82	5.65	5.45	5.23	5.01	4.78

58161 Source: Forrester Research, Inc.

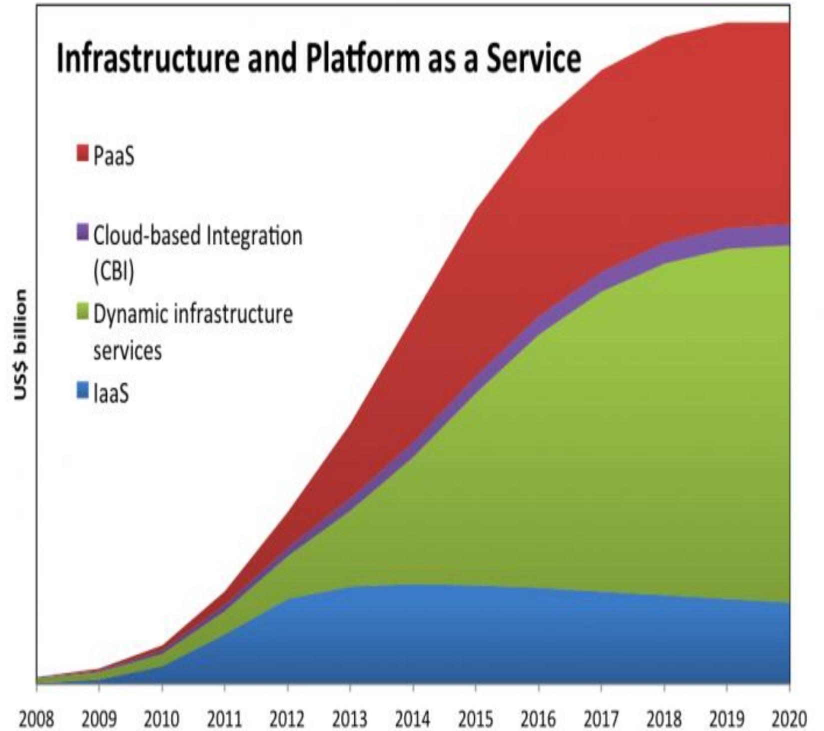

Source: Forrester Research Inc.

Cloud Computing for HPC

- HPC applications require high bandwidth, low latency and very high compute capabilities
- Clouds present a natural choice to meet HPC demands through
  - Flexible performance at scale
  - Time and cost optimization
- E.g., AWS runs variety of HPC applications including CAD, molecular modeling, genome analysis, weather simulation

Amazon's HPC Cloud
www.wired.com

Why Cloud is becoming popular?

- **Most desirable form of IT**
  - Cap-Ex free computing
  - Pay-as-you-go on-demand scaling
  - Resiliency and Redundancy
  - Fast projects deployment with cheap costs
- **Challenges**
  - Unpredictable performance
  - Lack of QoS support
  - Fault-tolarance
  - Security/Privacy issues

Cloud in the context of NAS

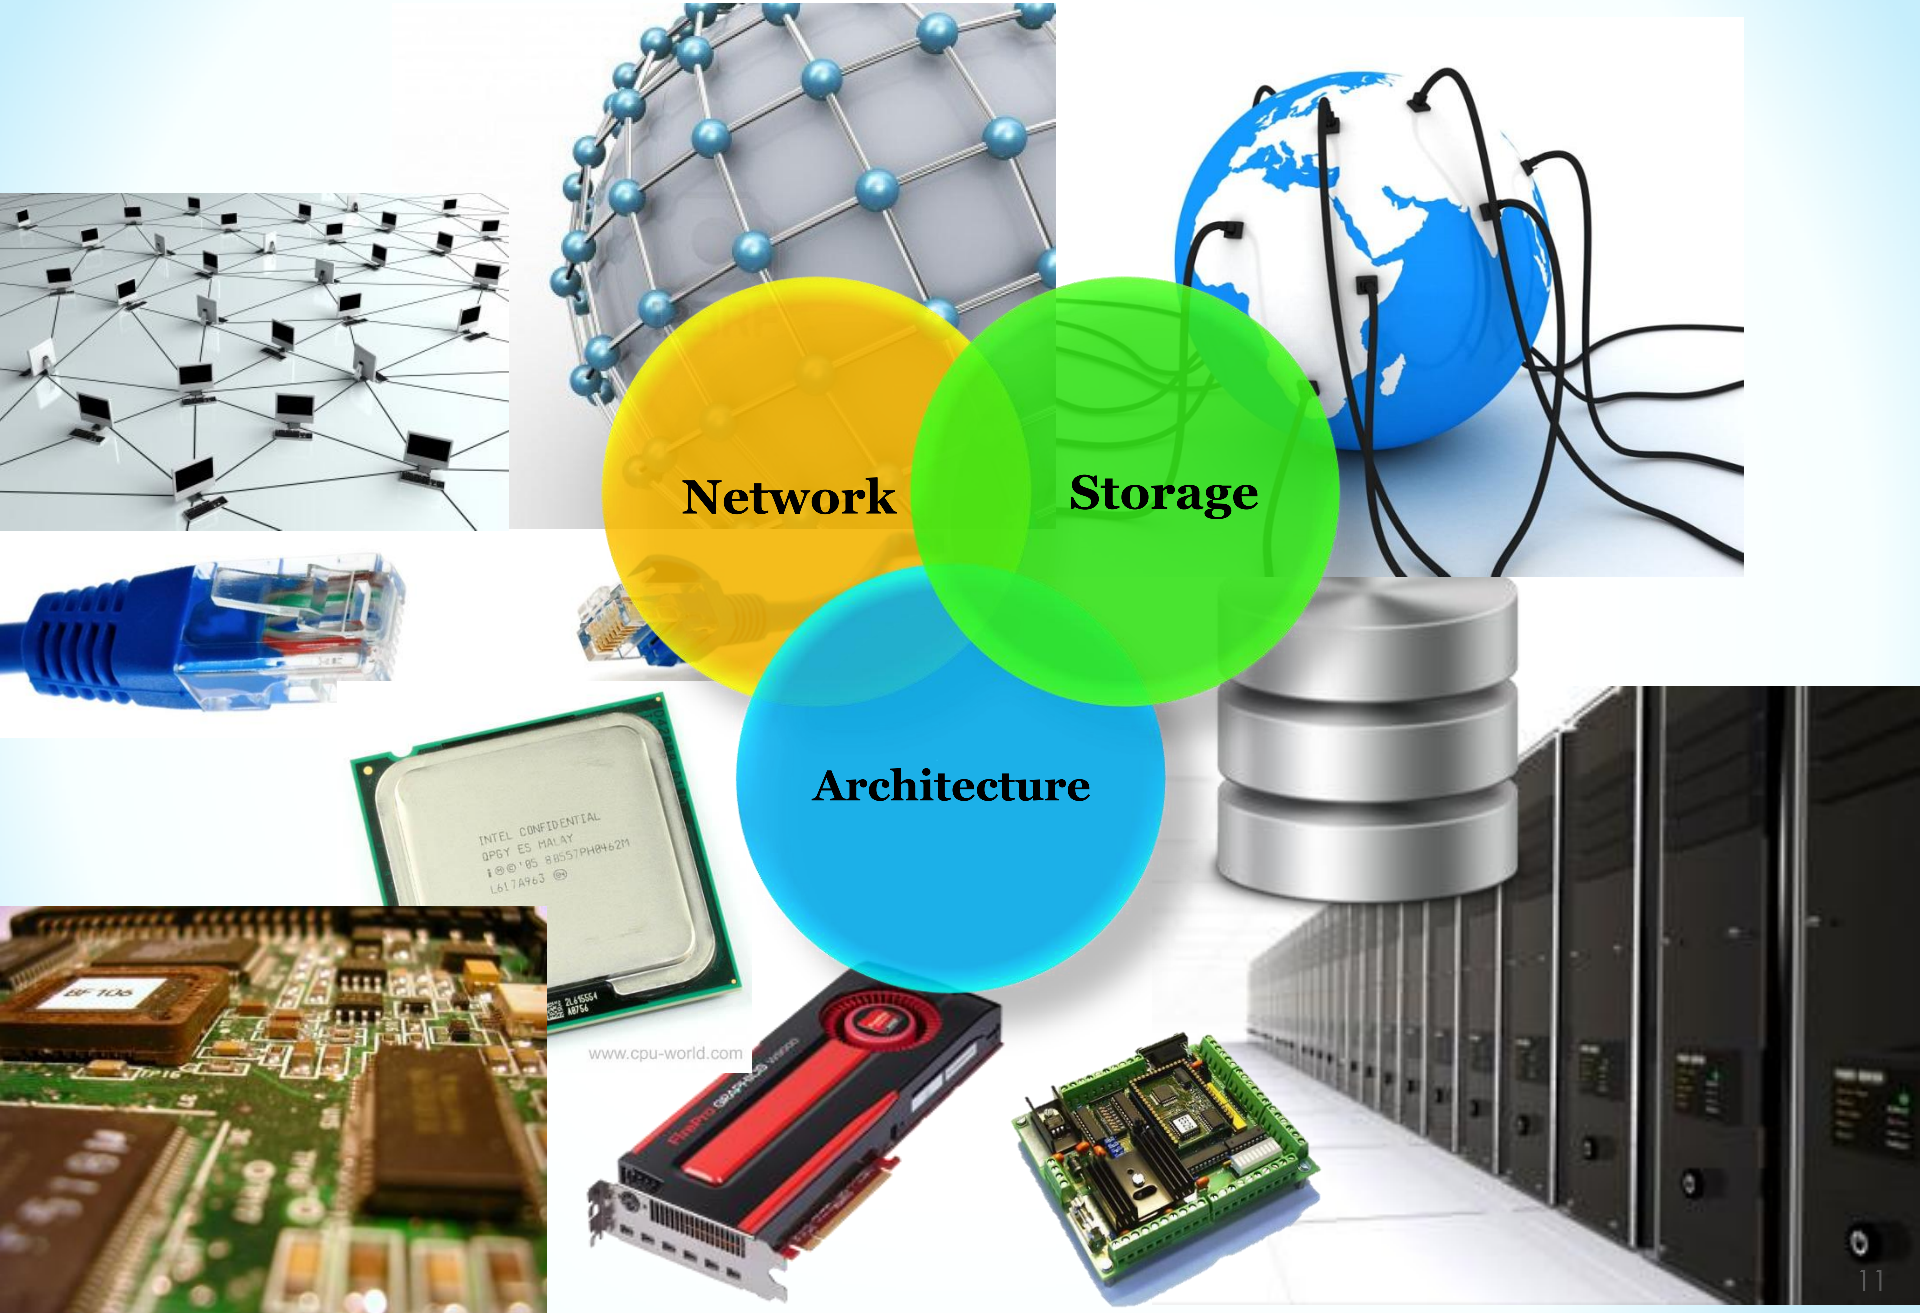

User Surveys - Cloud Challenges

Q: Rate the challenges/issues of the 'cloud'/on-demand model?

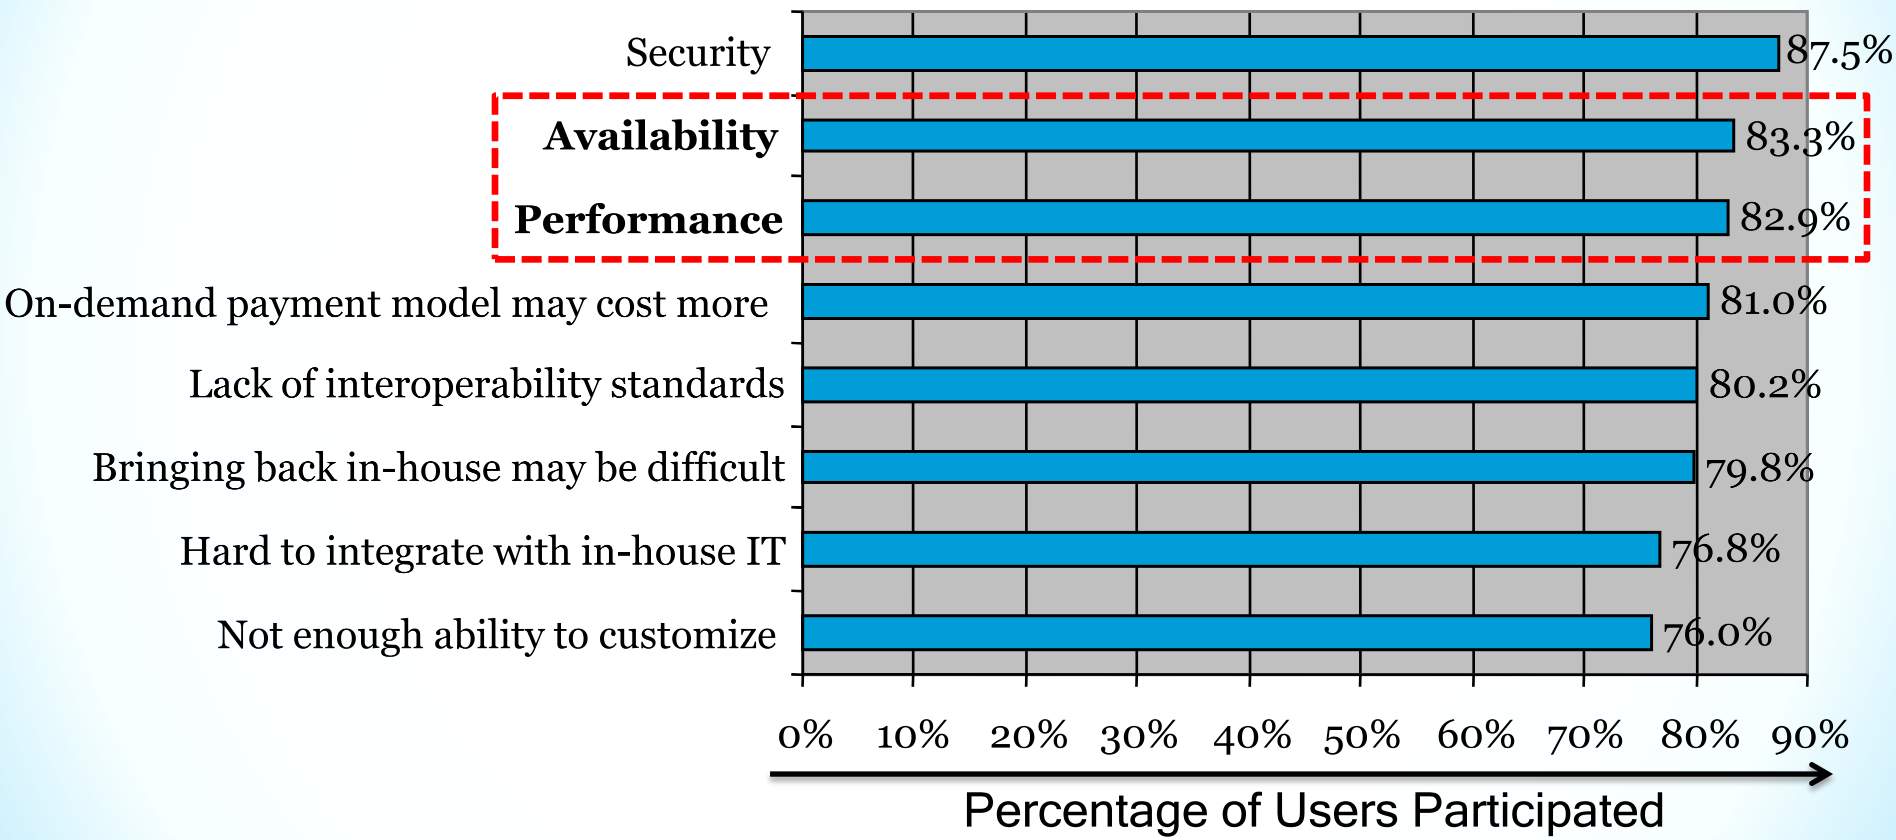

*Source: IDC Enterprise, September 2009. Numbers are representative showing relative importance

Cloud Performance: State-of-the-art

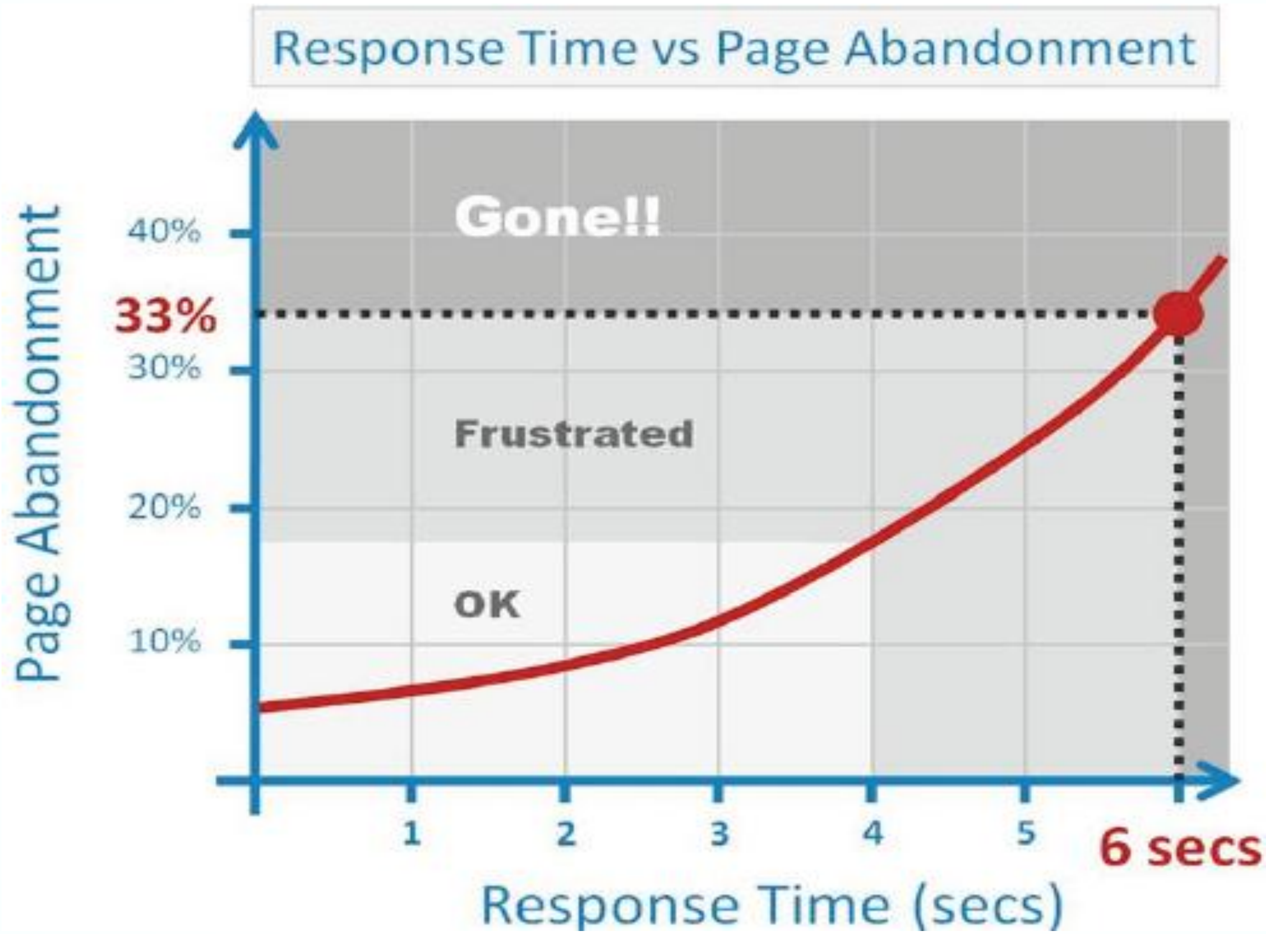

*Source: Performance in the Cloud, Survey report from Compuware.com

Real-life Cloud Failures

- Enterprise-level failures ➡ importance/need of reliability in clouds

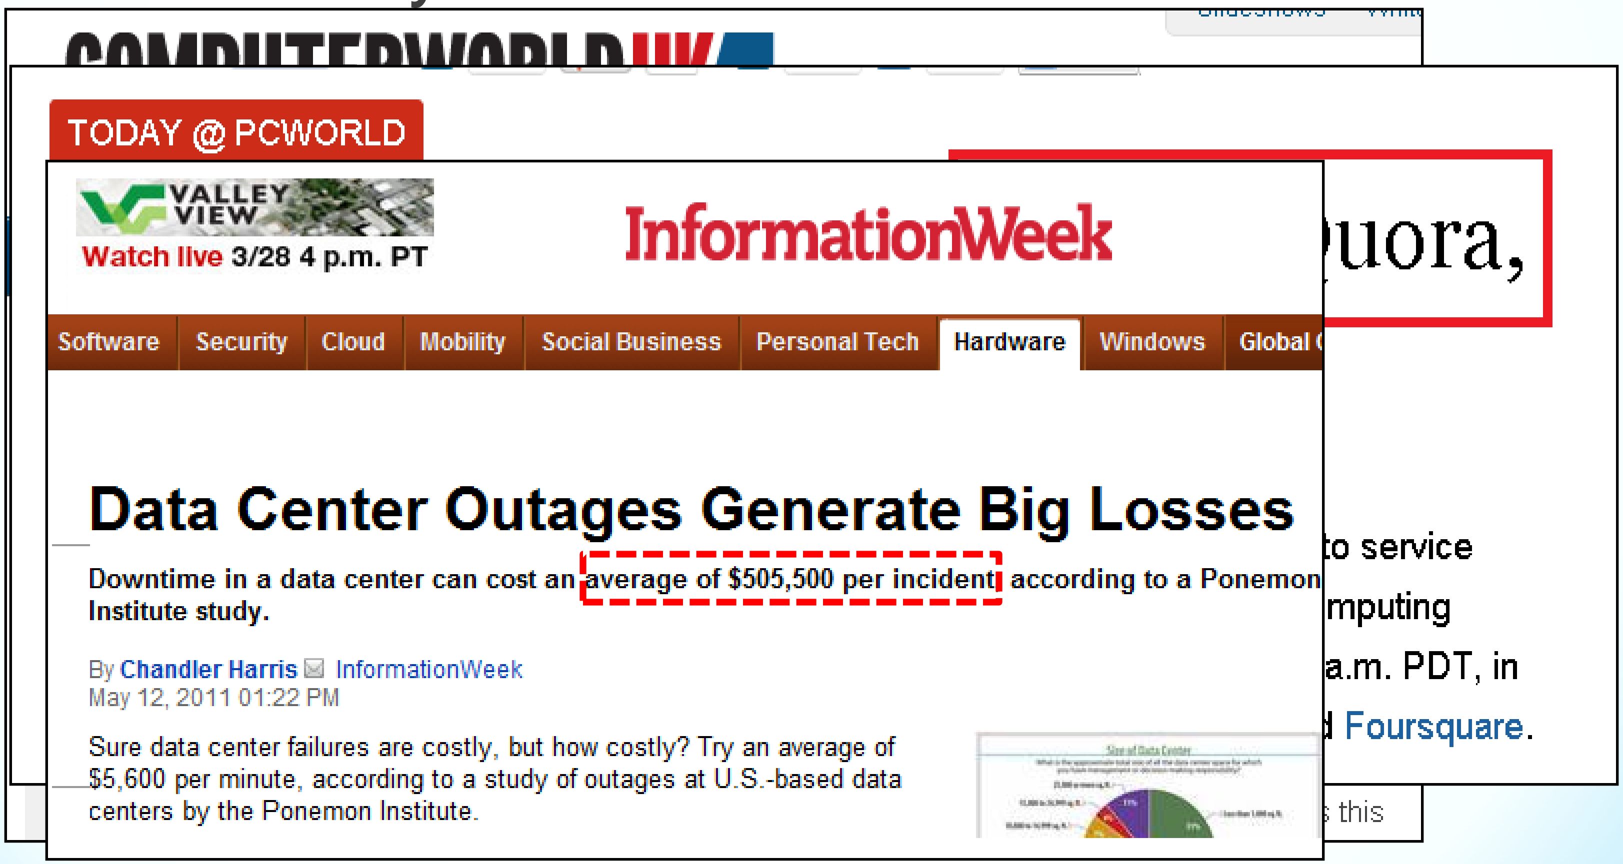

TODAY @ PCWORLD

VALLEY VIEW
Watch live 3/28 4 p.m. PT

InformationWeek

Software | Security | Cloud | Mobility | Social Business | Personal Tech | Hardware | Windows | Global

Data Center Outages Generate Big Losses

Downtime in a data center can cost an average of $5,505,500 per incident according to a Ponemon Institute study.

By Chandler Harris InformationWeek
May 12, 2011 01:22 PM

Sure data center failures are costly, but how costly? Try an average of $5,600 per minute, according to a study of outages at U.S.-based data centers by the Ponemon Institute.

Performance Challenges in Clouds

Cloud Workloads	Resource Management	Performance Model	Reliable services
Lack of representative cloud workloads	Efficient scheduling of resources	Lack of robust performance models	Fault diagnosis in virtualized clouds

Ongoing Research in HPCL

Cloud Performance

- **Workload Characterization**
  - **Task placement constraints**
  - **Modeling and Synthesizing**
    SOCC 2011
- **Resource Management**
  - **Fine-grained resource scheduling**
  - ***MROrchestrator and HybridMR***
    CLOUD 2012, ICDCS 2013
- **Performance Modeling**
  - **Performance quantification in Clouds**
  - ***D-Factor Algorithm***
    SIGMETRICS 2012
- **Fault Diagnosis**
  - **Problem diagnosis in clouds**
  - ***CloudPD*: Fault Management Framework**
    DSN 2013

Modeling and Synthesizing Task Placement Constraints in Google Compute Clusters

Joint work with Google
SOCC, 2011

Cloud Workloads

- Current workload only addresses *how much* resource tasks use
- But machine heterogeneity sometimes requires tasks to specify *which* resources they can use
- Constraints are predicates on machine properties that limit the set of machines a task can run
  - E.g., “kernel_version = x”
- Why constraints?
  - Machine heterogeneity
  - Application optimization
  - Problem avoidance

Thesis: Task placement constraint is an important workload property

Impact of Task Constraints

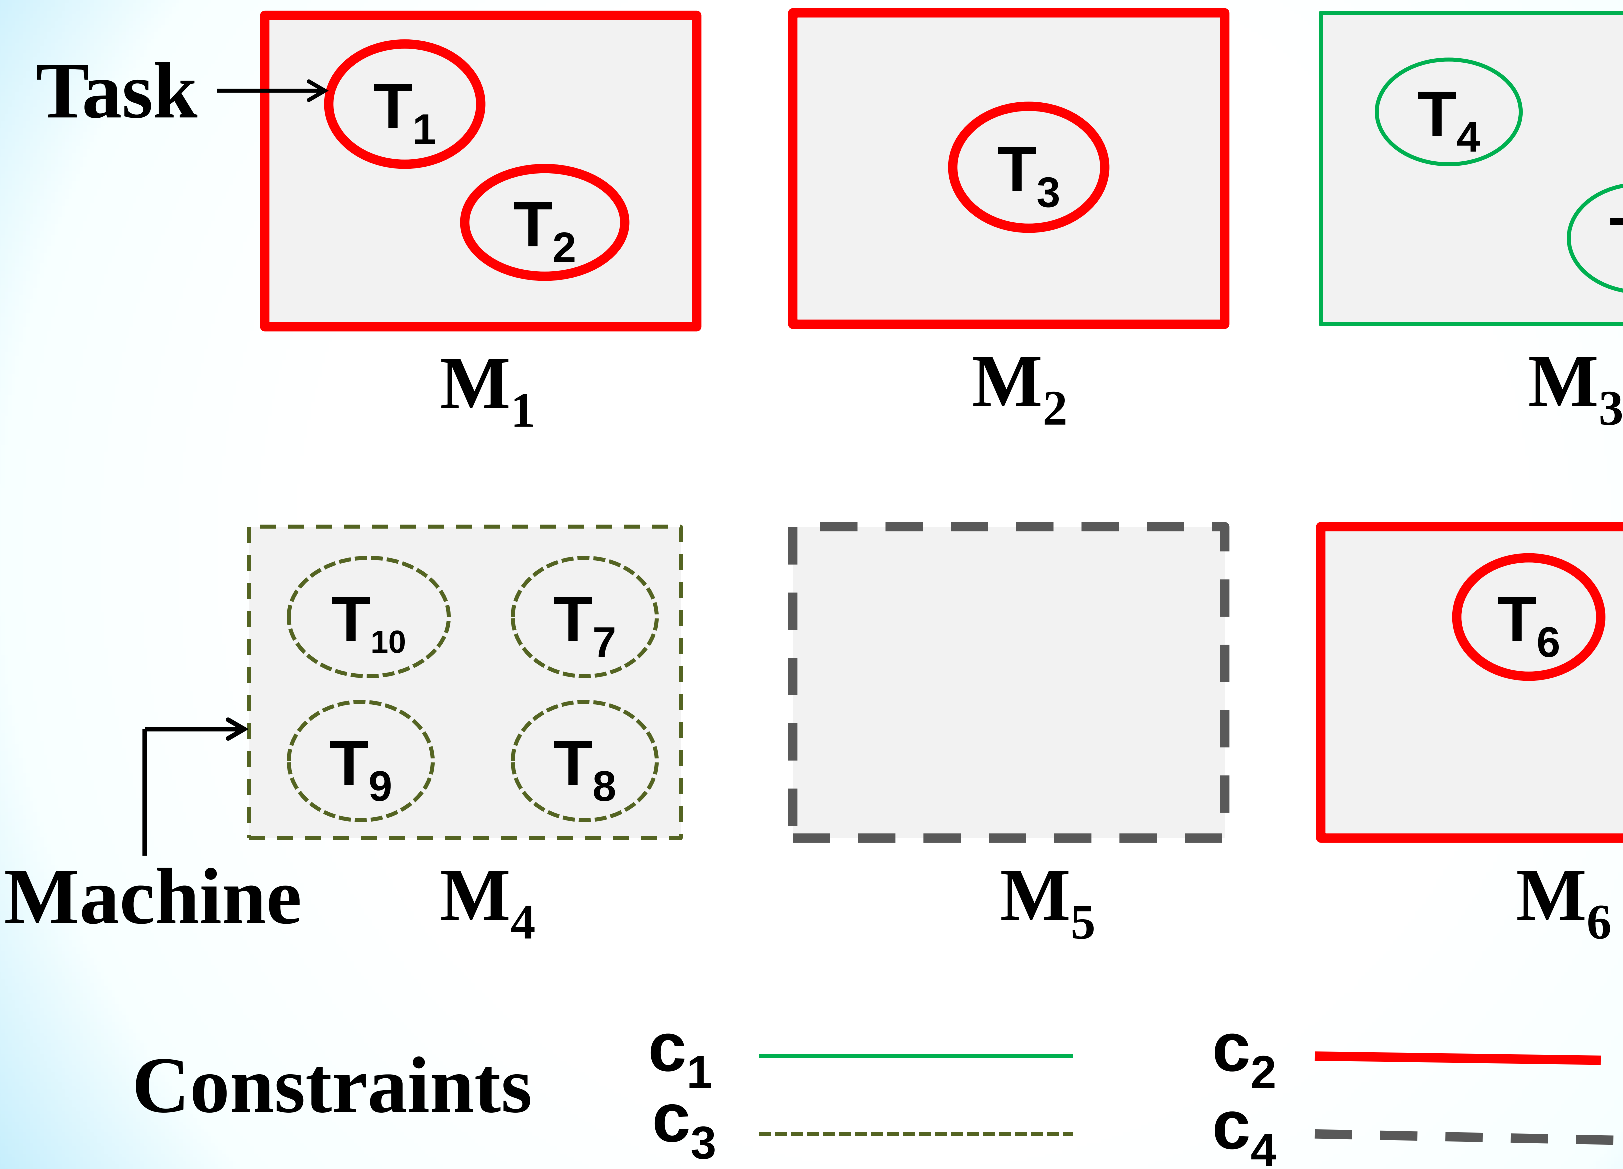

Questions

Q1: Do task placement constraints have a significant impact on task scheduling delays?

Q2: How do we predict the impact of constraints on task scheduling delays?

Q3: How do we extend existing performance benchmarks to include representative task constraints and machine properties?

Q1: Do Constraints Impact Task Scheduling Delays?

- Methodology
  - Run trace driven benchmarks to obtain task scheduling delays
  - Compare results with and without task constraints
- Evaluation metric – (normalized) task scheduling delay
  - Ratio of delay with constraints to delay without constraints

Effect of Task Constraints

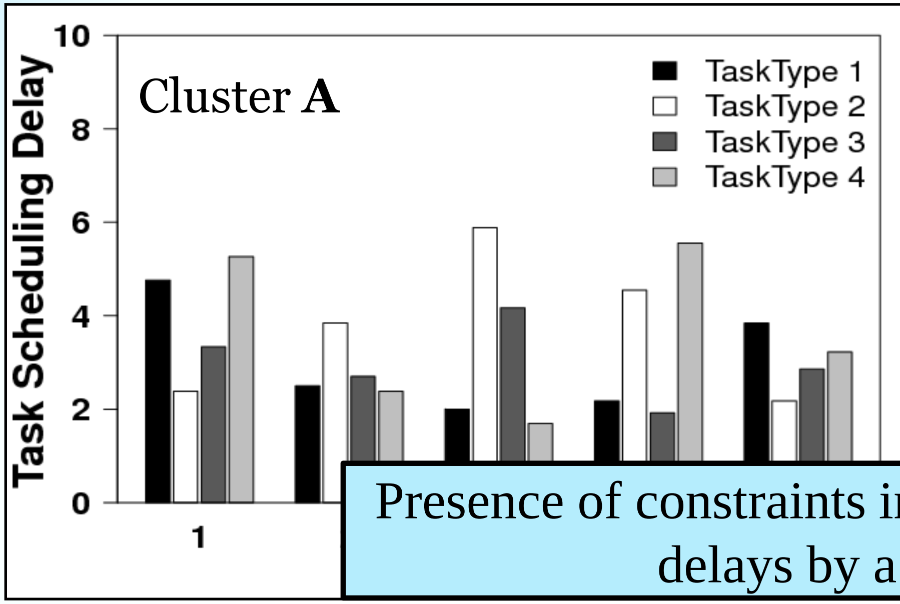

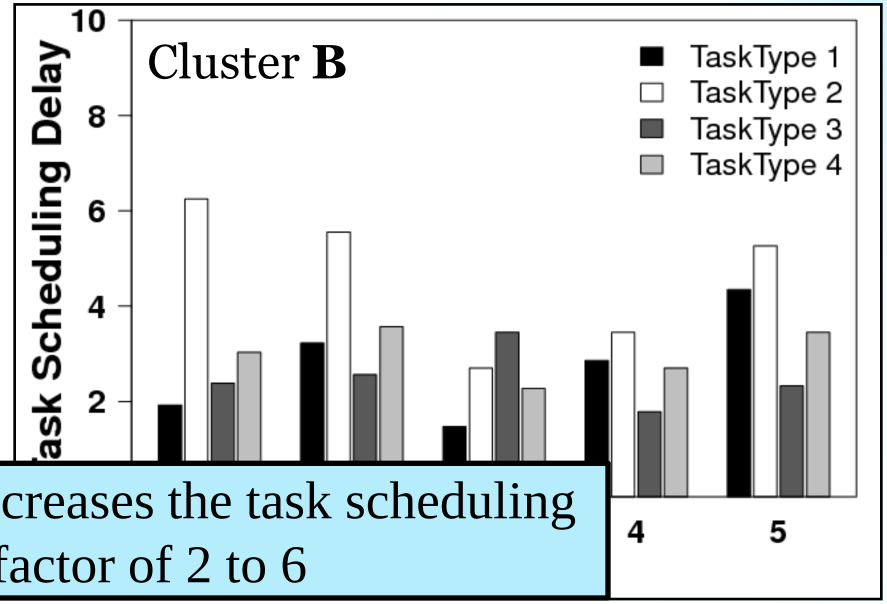

Presence of constraints increases the task scheduling delays by a factor of 2 to 6

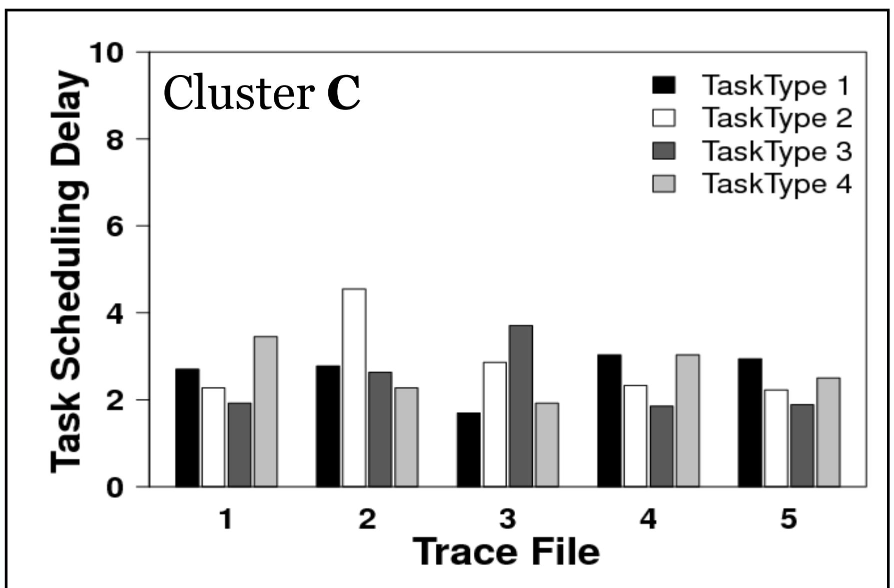

Q2: How to Predict the Impact of Constraints on Task Scheduling Delays?

- Motivation – extend resource utilization to constraints
- **Utilization Multiplier (UM)** is the ratio of resource utilization seen by tasks with a constraint to the average utilization of the resource
  - $u_{r,c}$ = UM metric for resource *r* and constraint *c*
- Maximum utilization multiplier (u_c*)
  - $u_c^* = \max_r(u_{r,c})$

u_c* Predicts Task Scheduling Delays

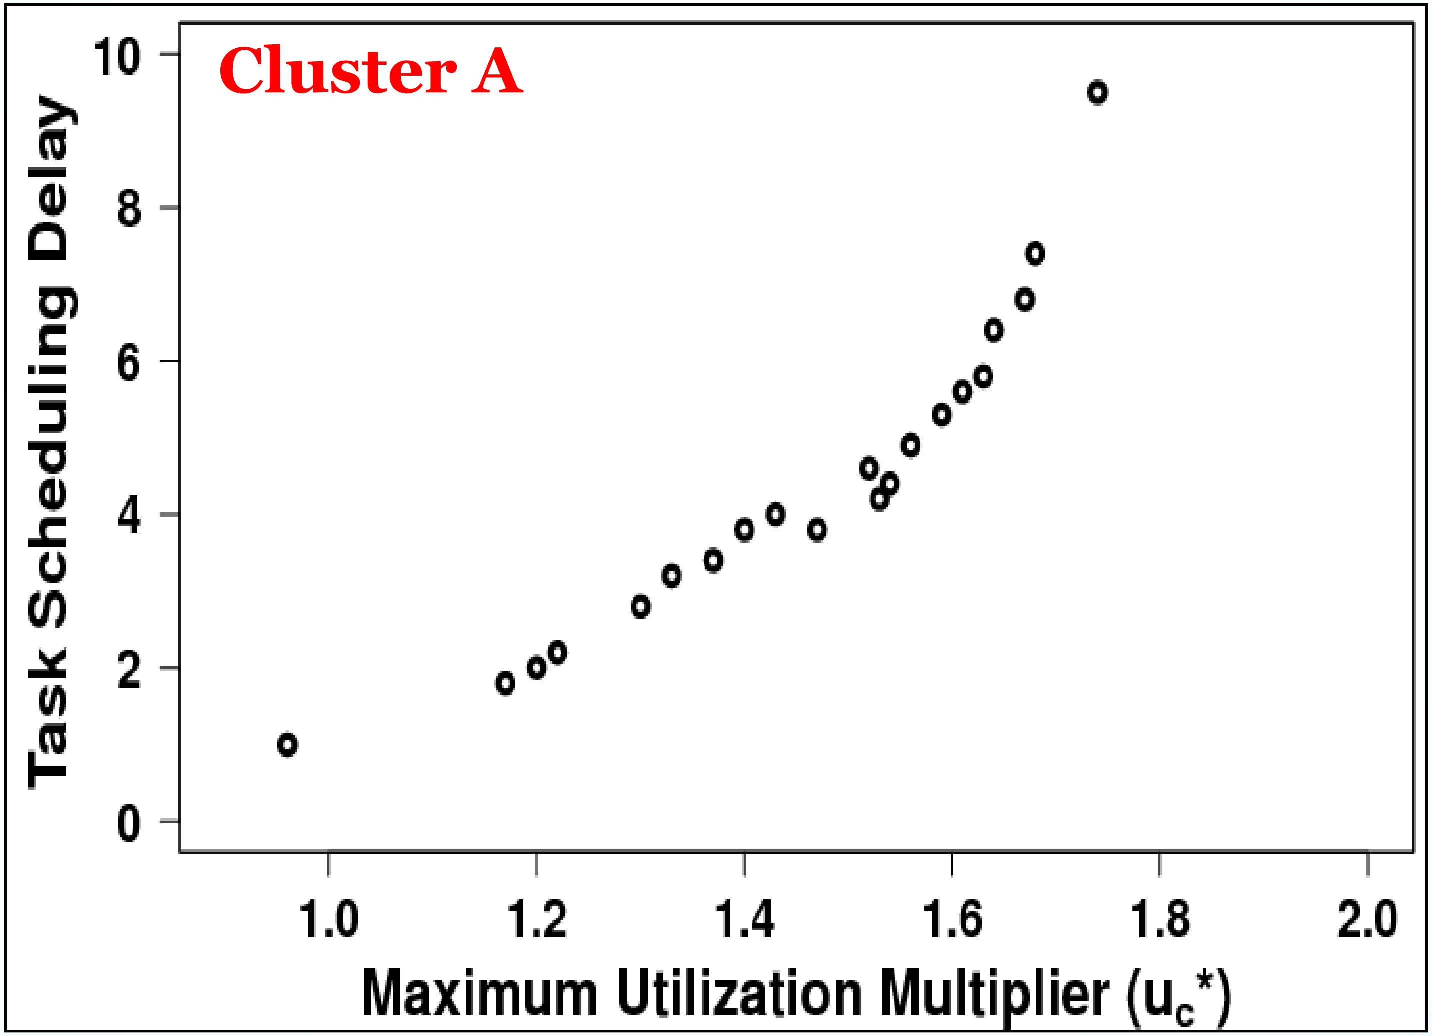

Q3: How to Extend Performance Benchmarks to Include Task Placement Constraints?

Change benchmarking algorithms

- When benchmark initializes
  - Augment machines with representative (synthetic) properties
- When tasks arrive
  - Augment tasks with representative (synthetic) constraints
- When a task is scheduled
  - Only consider machines whose properties are compatible with the task's constraints

Validation of Synthetic Characterization

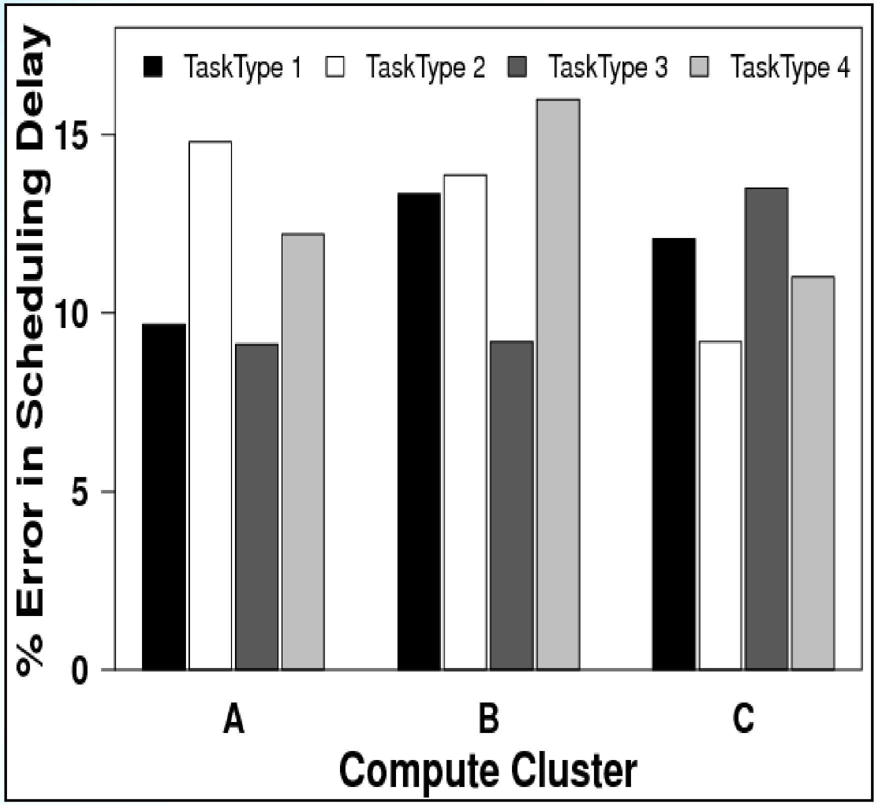

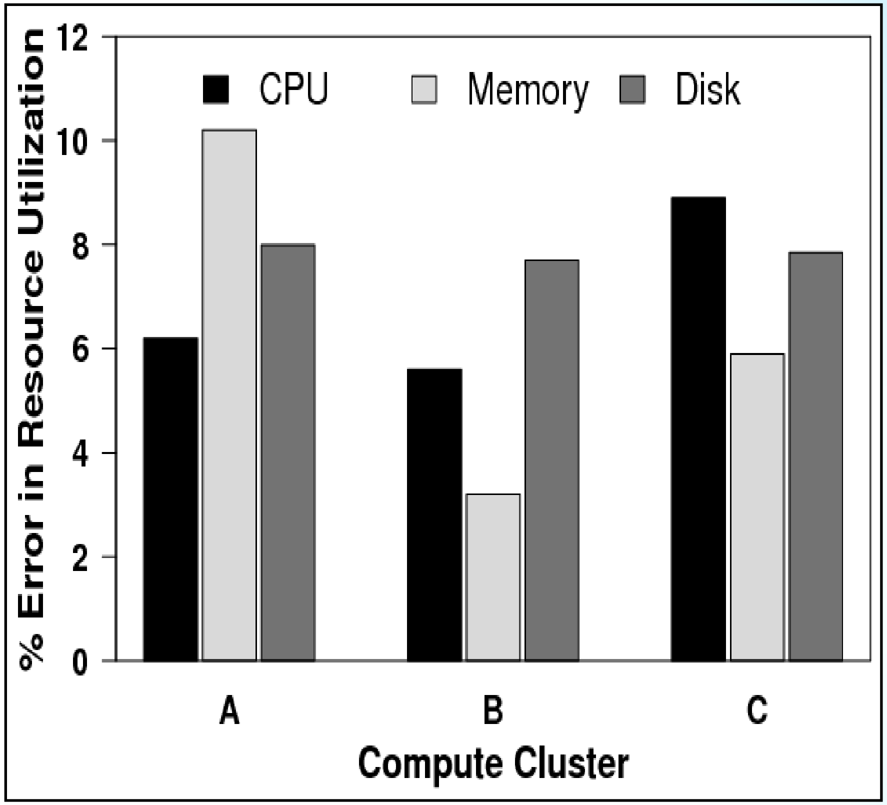

Characterization reproduces observed performance with

~ 13% error in task scheduling delay;

~ 5% error in resource utilization

Resource Management in Hadoop MapReduce Clusters

Cloud 2012

Who is using MapReduce/Hadoop?

Bank of America
digg
Booz | Allen | Hamilton
explorys
YAHOO!
pentaho
facebook
sematext
twitter
Aol.
Karmasphere
mozilla Firefox
cloudera
hp
GE
NTT Data
cme
Chicago Mercantile Exchange
TexelTek
BBN TECHNOLOGIES
ebay
intel
ORBITZ
Greenplum
Visible MEASURES
Datameer
Infochimps
OpenLogic
StumbleUpon

MapReduce Overview

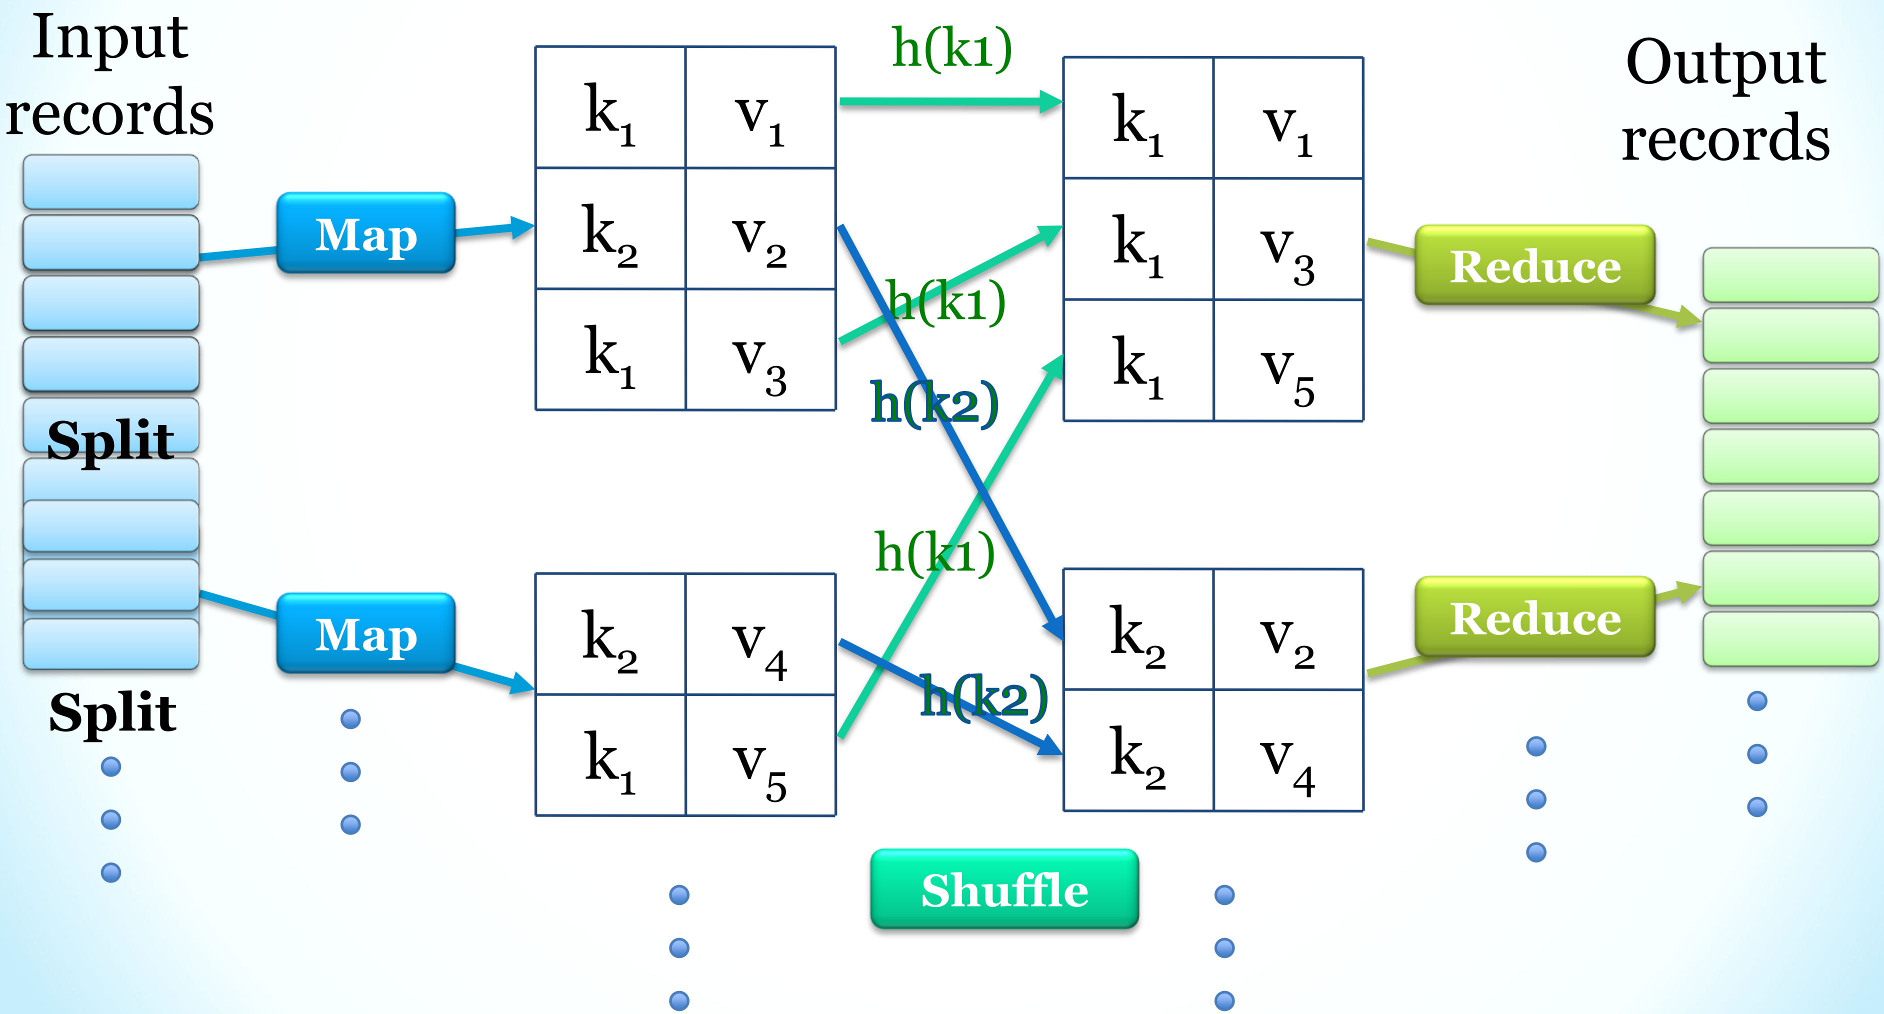

Hadoop MapReduce Framework

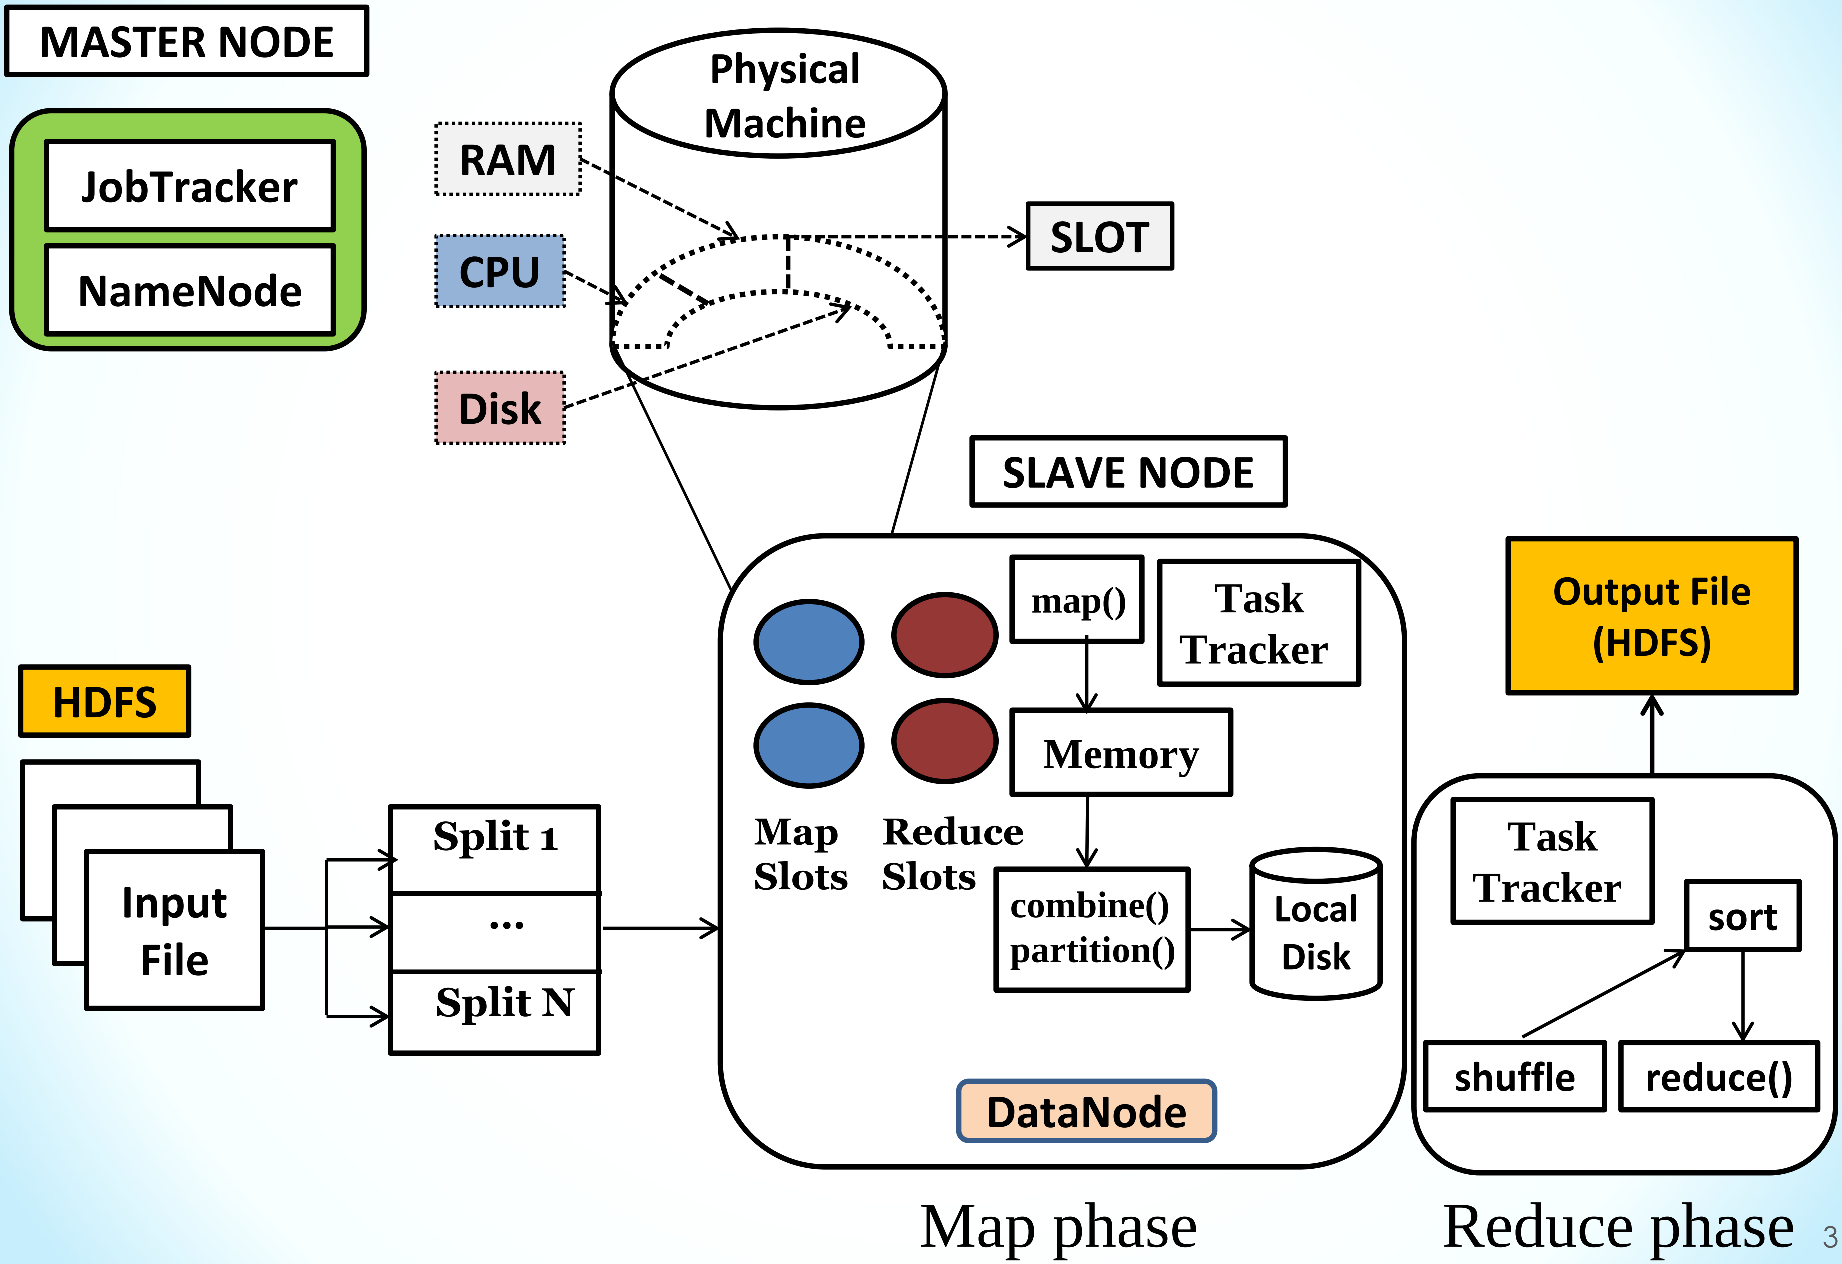

Motivation

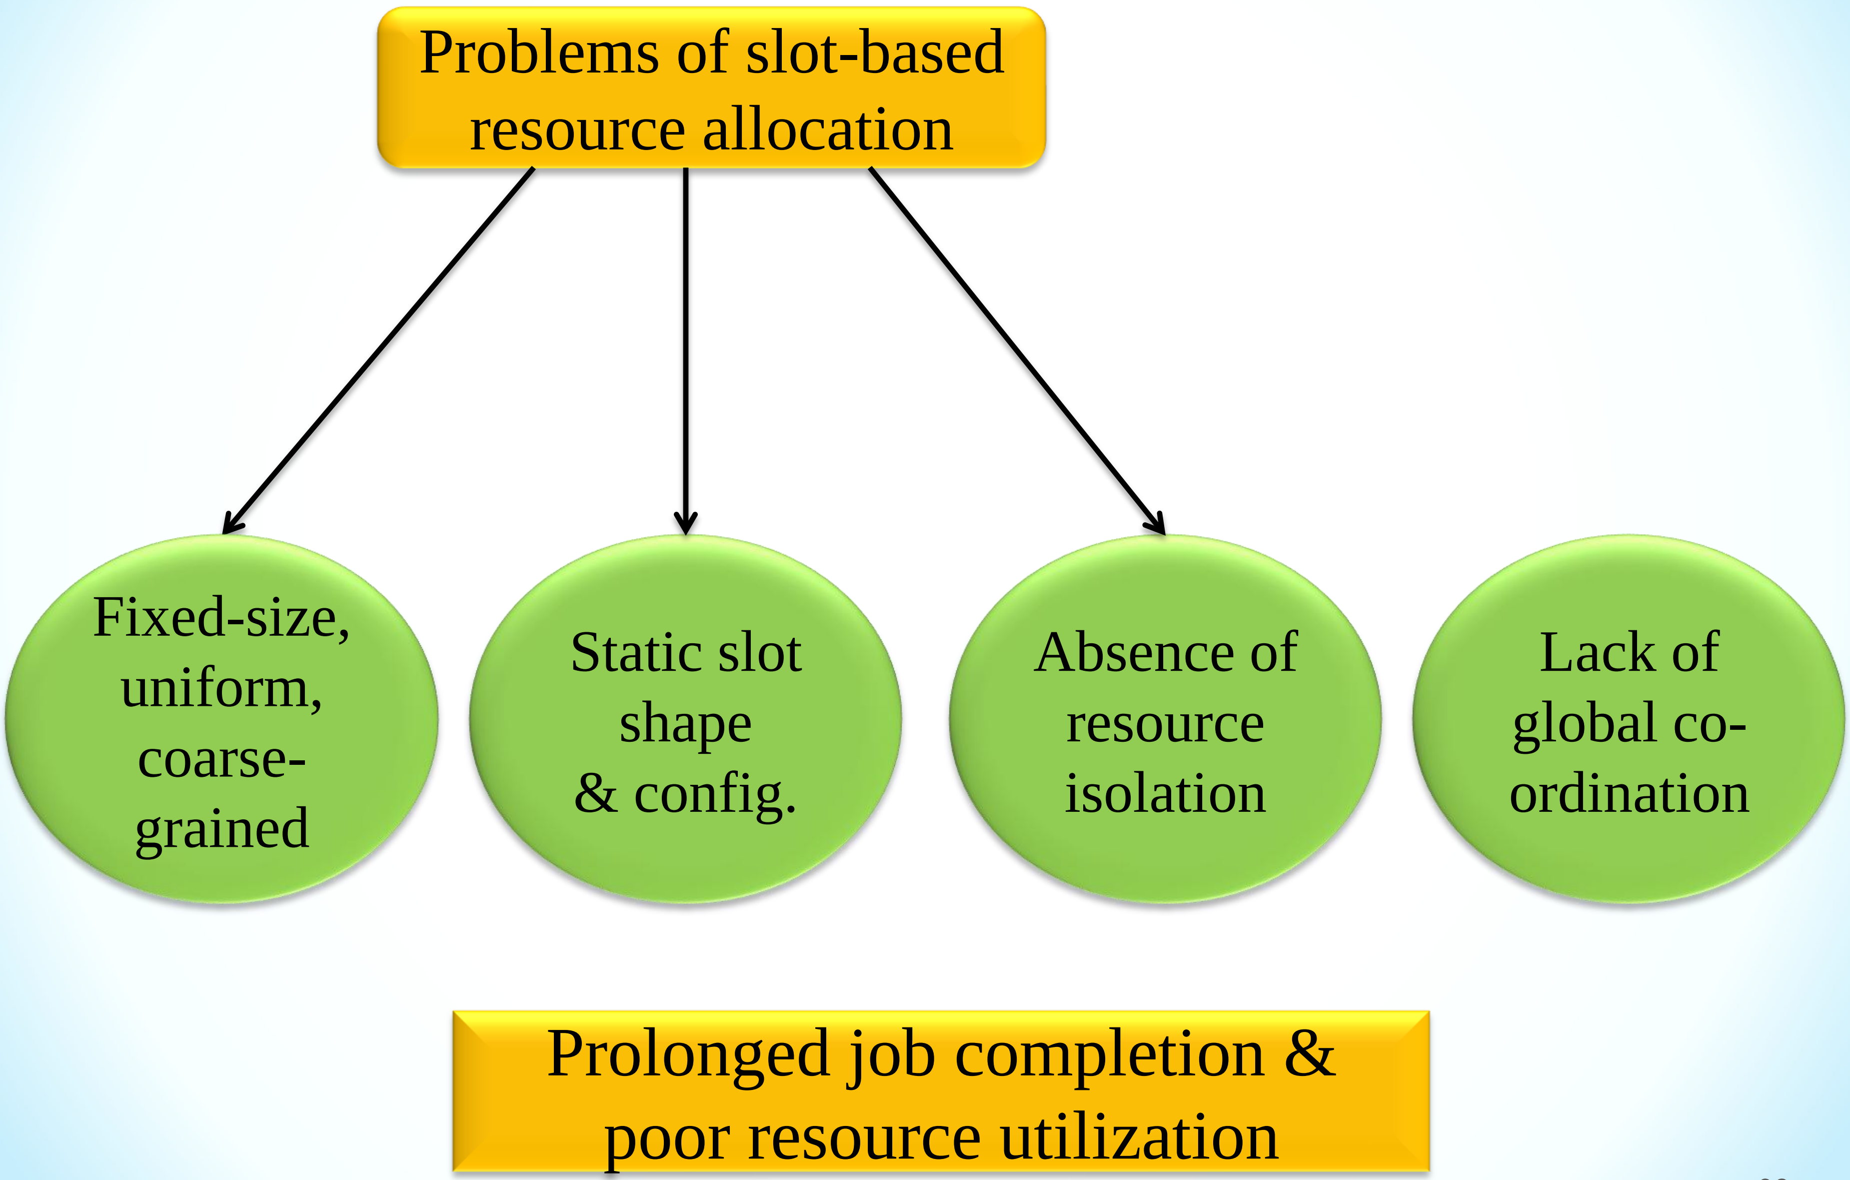

Problems of slot-based resource allocation
Fixed-size, uniform, coarse-grained
Static slot shape & config.
Absence of resource isolation
Lack of global co-ordination
Prolonged job completion & poor resource utilization

Architecture of MROrchestrator

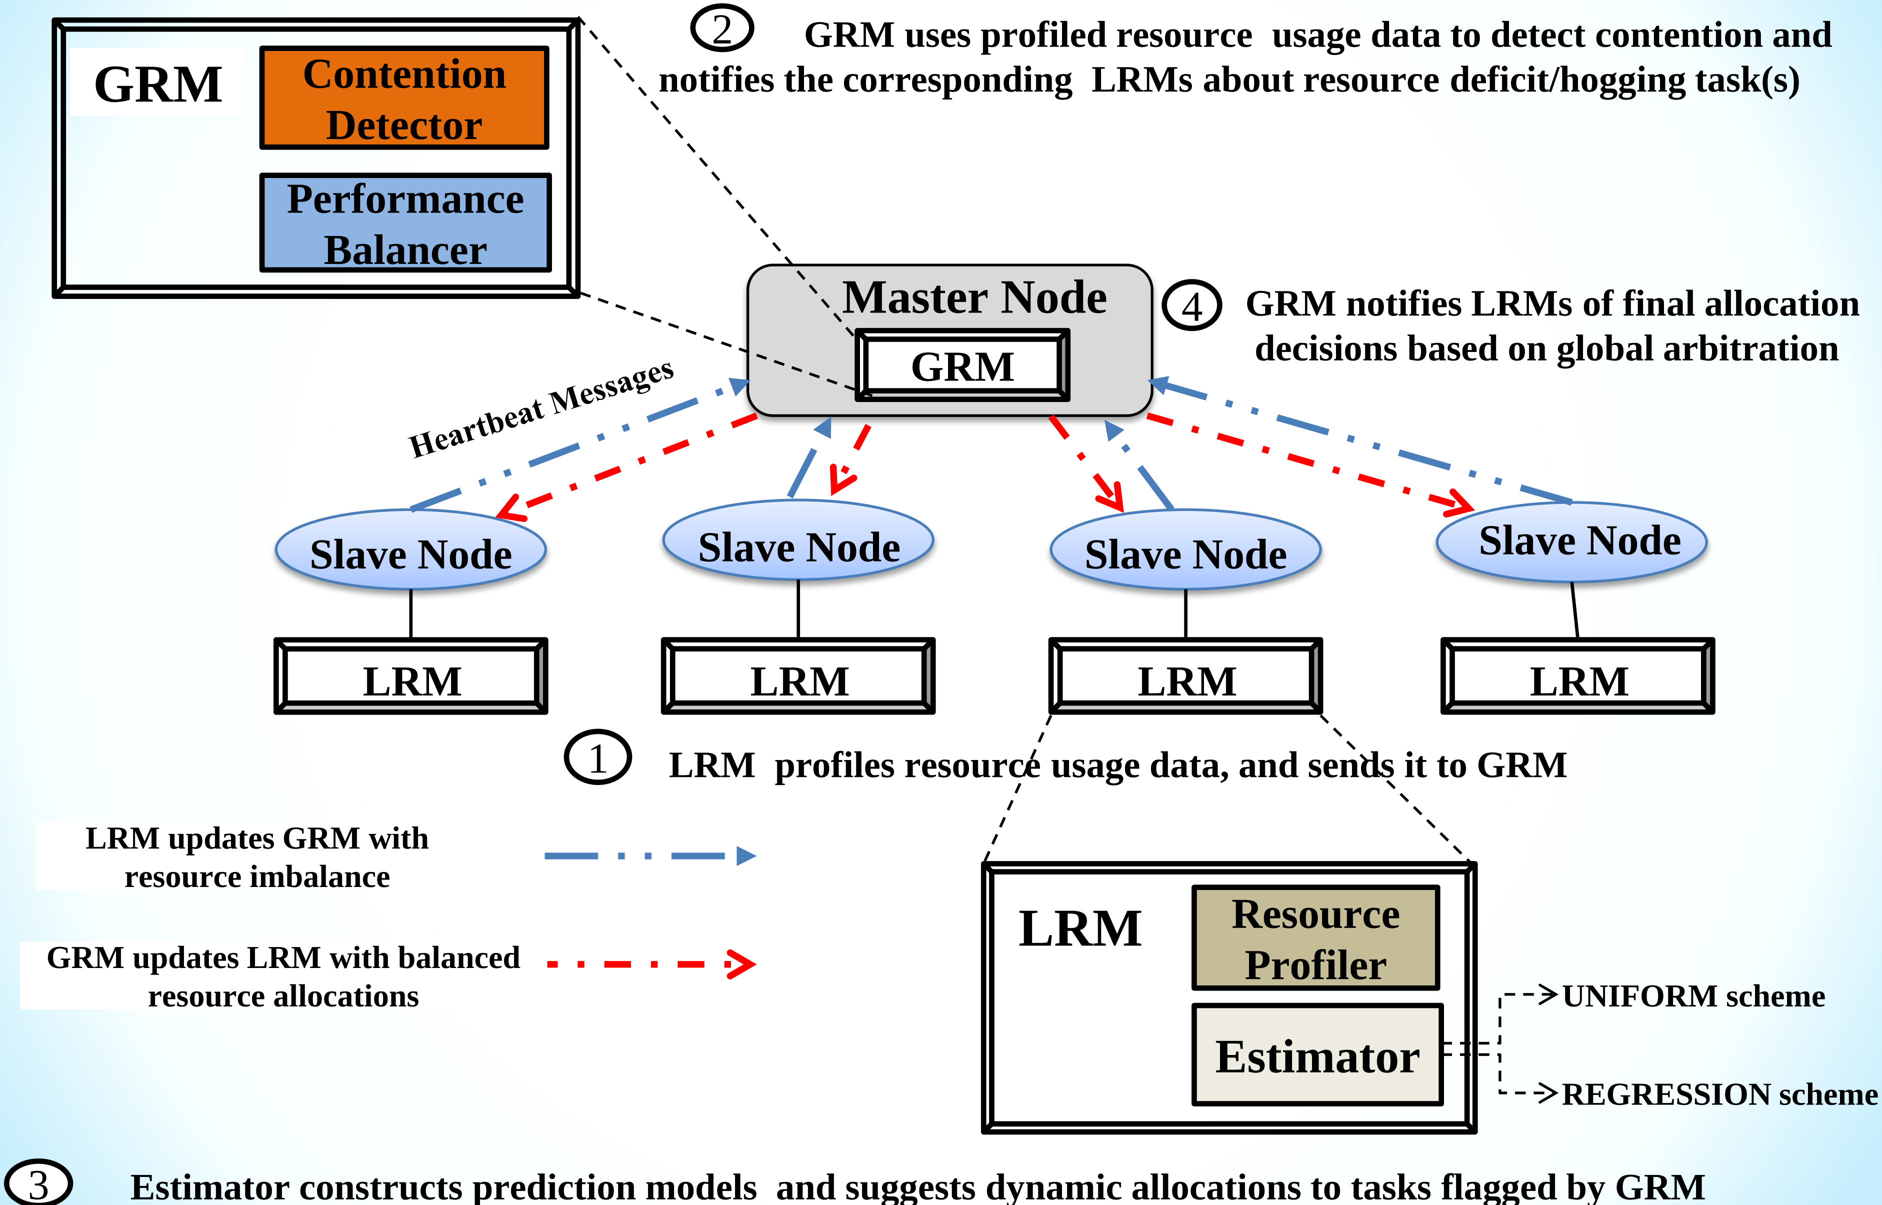

Evaluation

Experimental Platform

Environment	# of nodes	Machine specs.	Tool	Software
Native Hadoop Cluster	24 physical hosts	64-bit, 2.4 GHz AMD Opteron CPU, 4GB RAM, 1GB Ethernet	Linux Containers	Hadoop v0.20.203.0
Virtualized Hadoop Cluster	24 virtualized hosts on 12 physical hosts	Xen Hypervisor with same machine specs. as native Hadoop	Xen-xm	Hadoop v0.20.203.0

Benchmarks

Applications	Data set	Resource sensitiveness
Sort	20 GB text data	CPU + I/O
Wcount	20 GB text data	CPU + Memory
PiEst	10 million points	CPU
DistGrep	20 GB text data	CPU + I/O
Twitter	25 GB Twitter graph data	CPU + Memory
Kmeans	10 GB numeric data	CPU + I/O

Virtualized Infrastructure

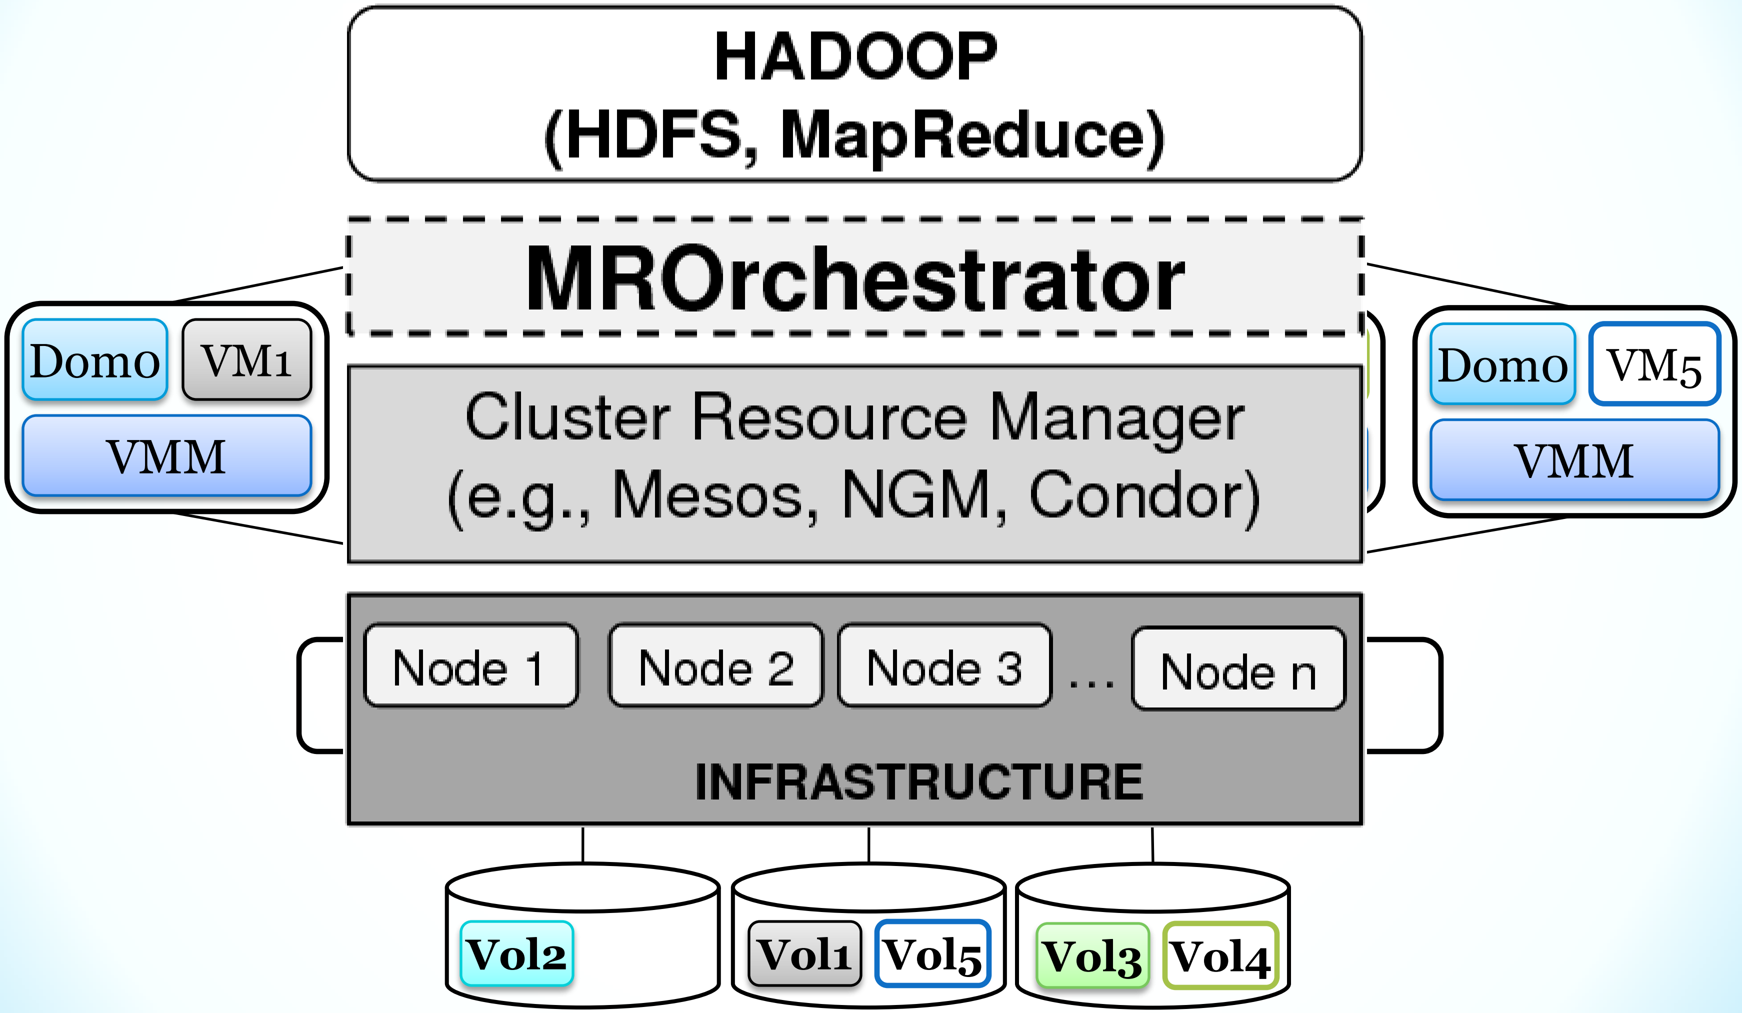

Native Hadoop Cluster

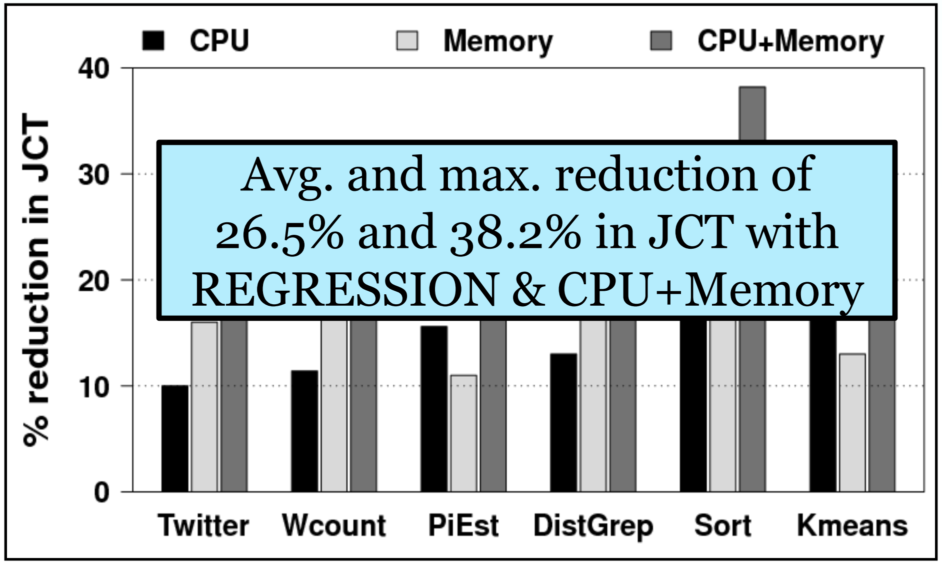

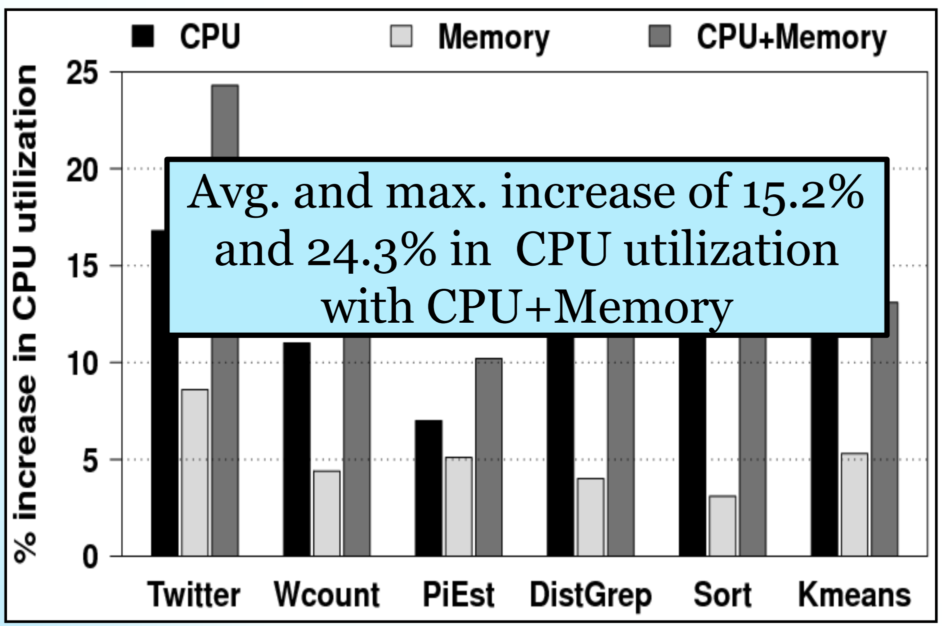

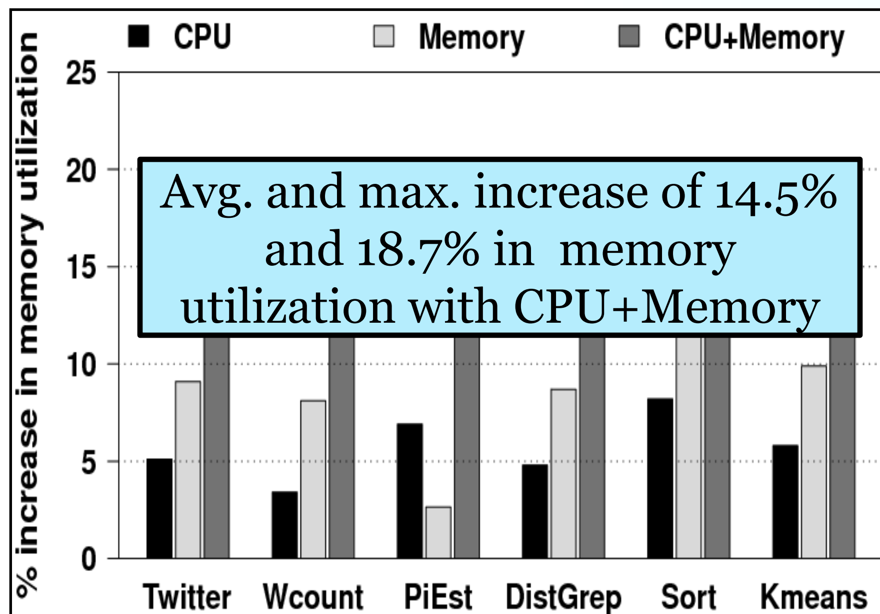

MROrchestrator with Mesos and NGM

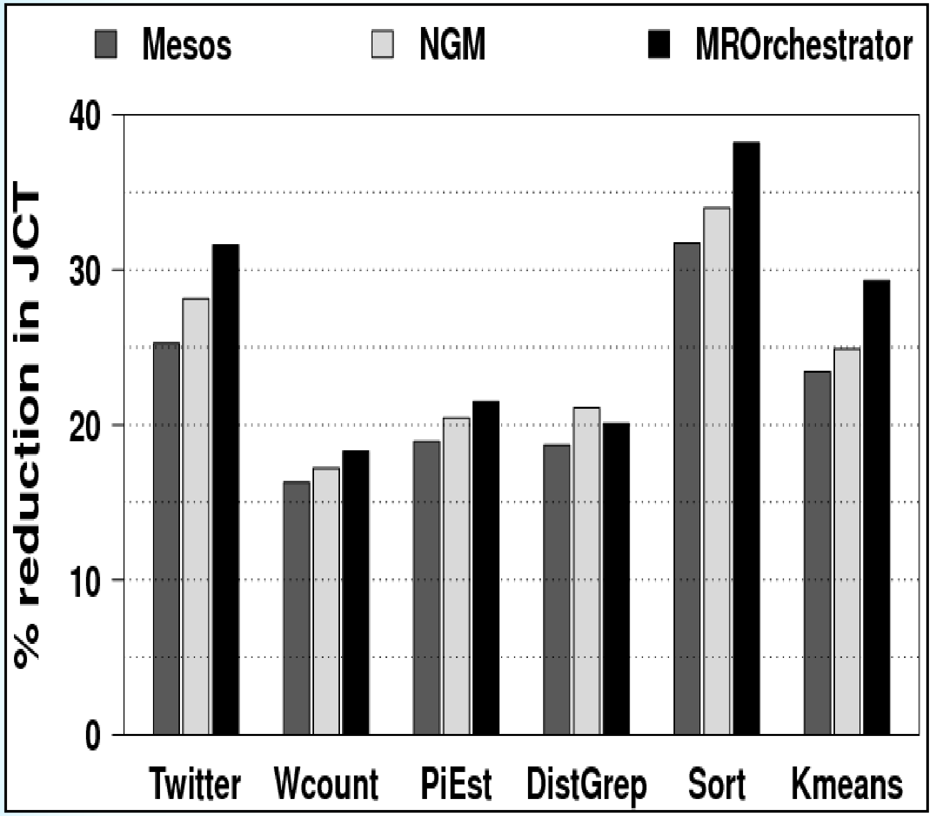

Performance comparison of Mesos, NGM and MROrchestrator.

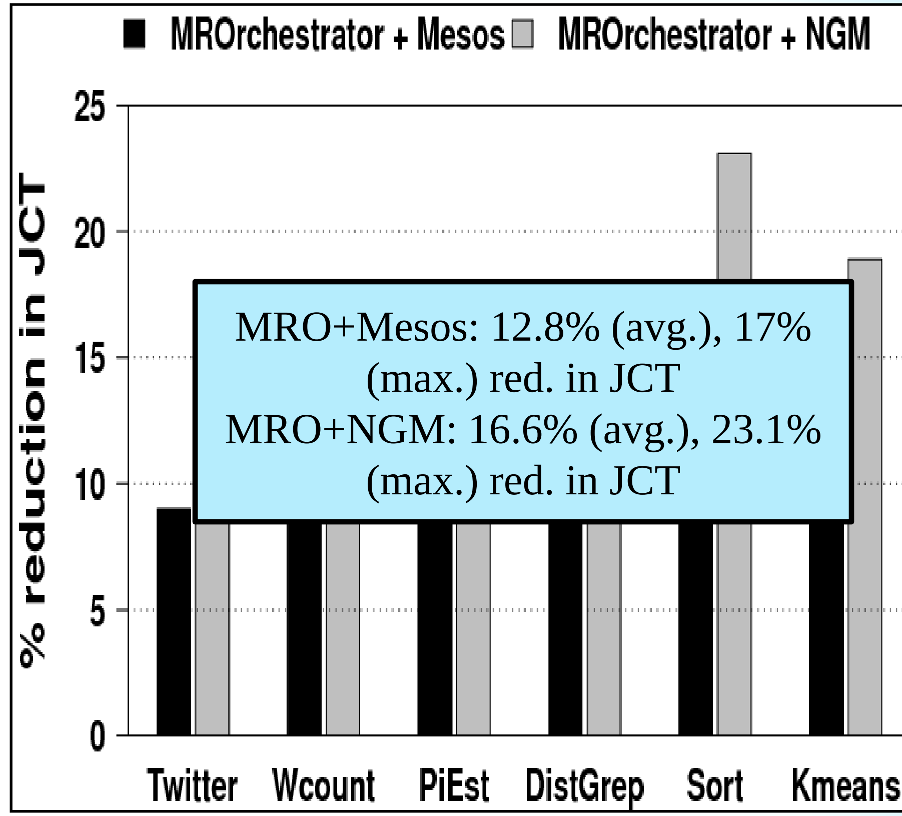

Performance benefits from the integration of MROrchestrator with Mesos and NGM.

HybridMR: A Hierarchical Scheduler for Hybrid Data Centers

ICDCS 2013

Motivation for Hybrid Platform

- Interactive applications – virtual environment
- Batch jobs (MapReduce) – native environment

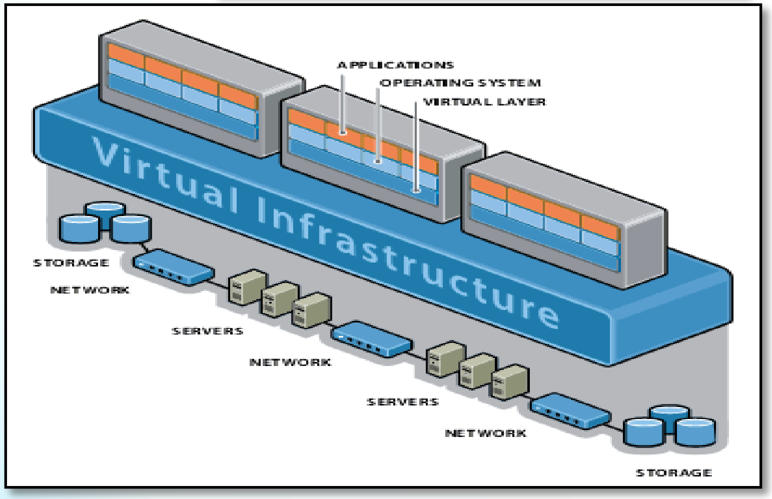

Opportunity: Best of Two Worlds!

Hybrid Compute Cluster

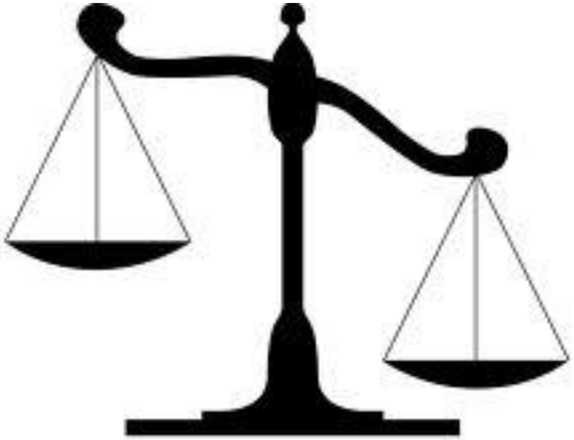

- ➢ Suitable for batch workloads
- ➢ Incurs high cost

- ➢ Attractive for interactive applications
- ➢ Poor I/O performance

HybridMR

- 2-phase hierarchical scheduler for effective resource management in hybrid datacenters
- 1st phase: estimates virtualization overheads to guide placement of MapReduce jobs
- 2nd phase: dynamic resource management of MapReduce jobs co-running interactive applications

HybridMR Architecture

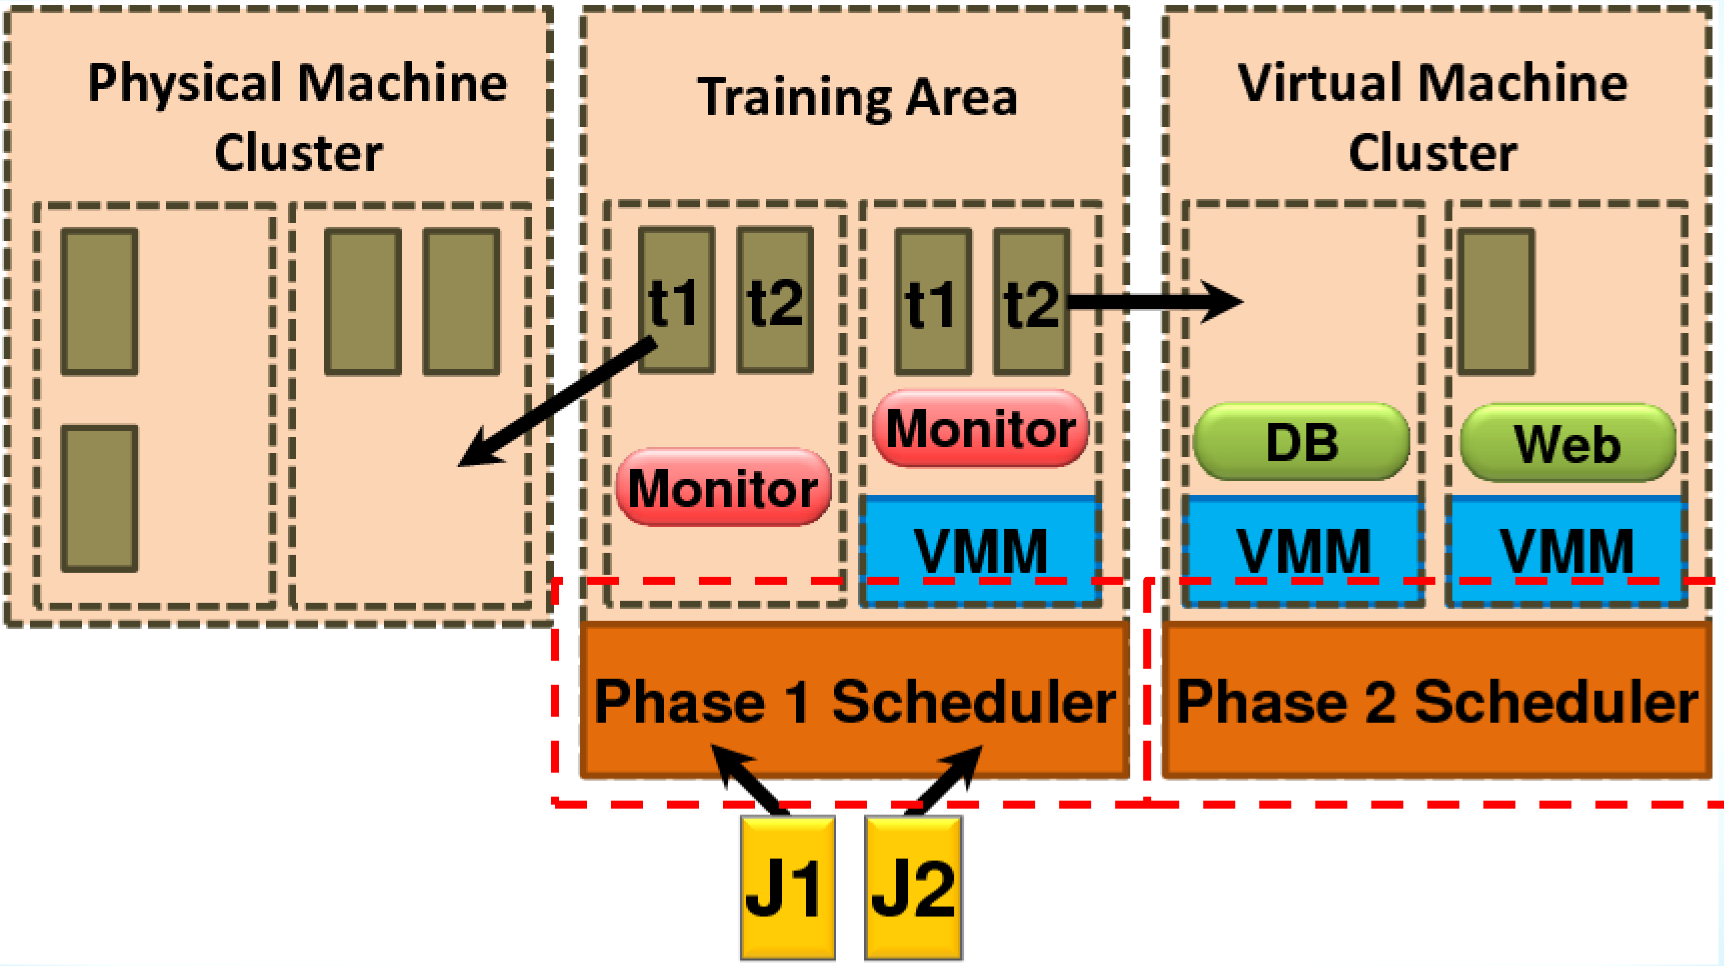

Phase II Scheduler: Architecture

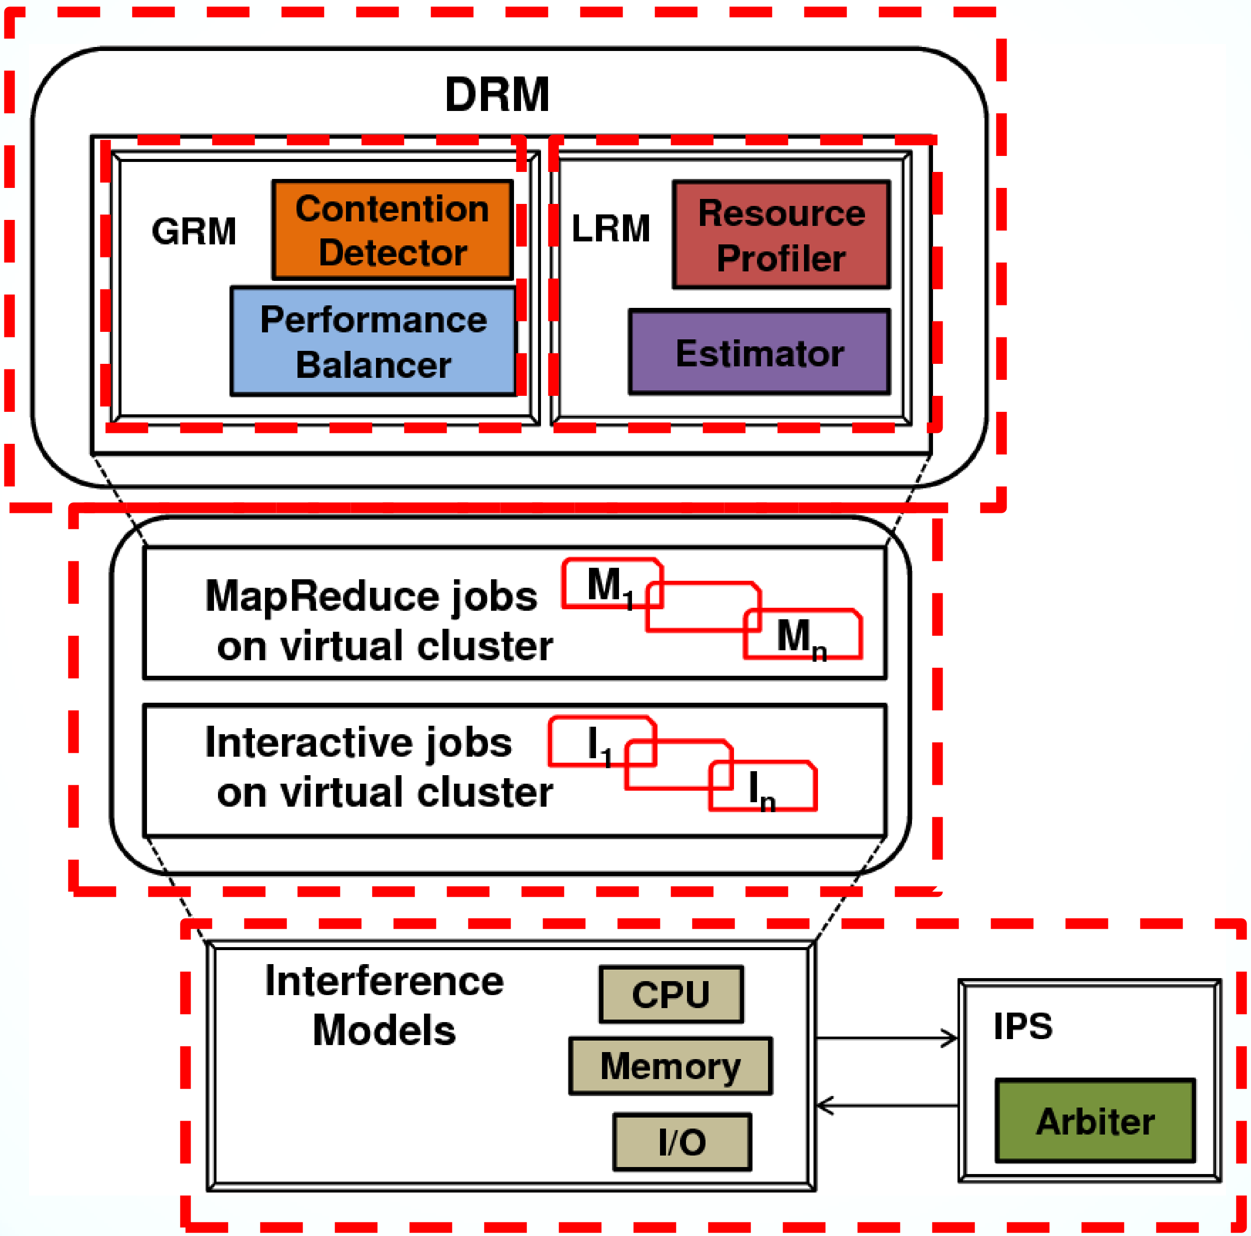

Results

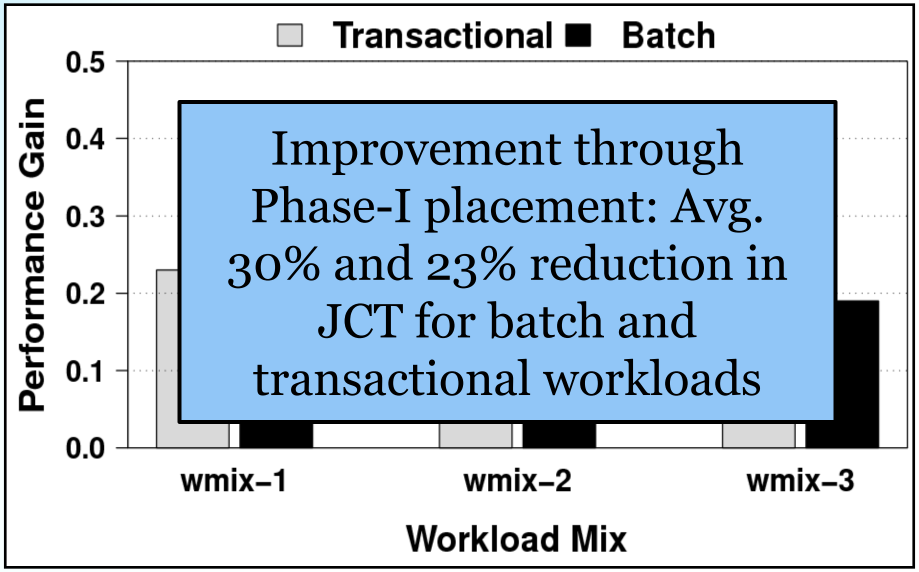

Improvement through Phase-I placement: Avg. 30% and 23% reduction in JCT for batch and transactional workloads

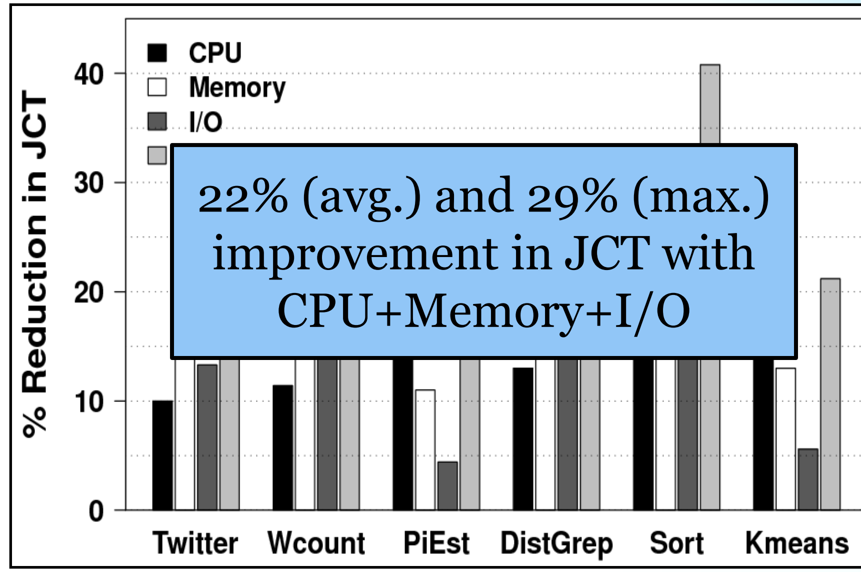

22% (avg.) and 29% (max.) improvement in JCT with CPU+Memory+I/O

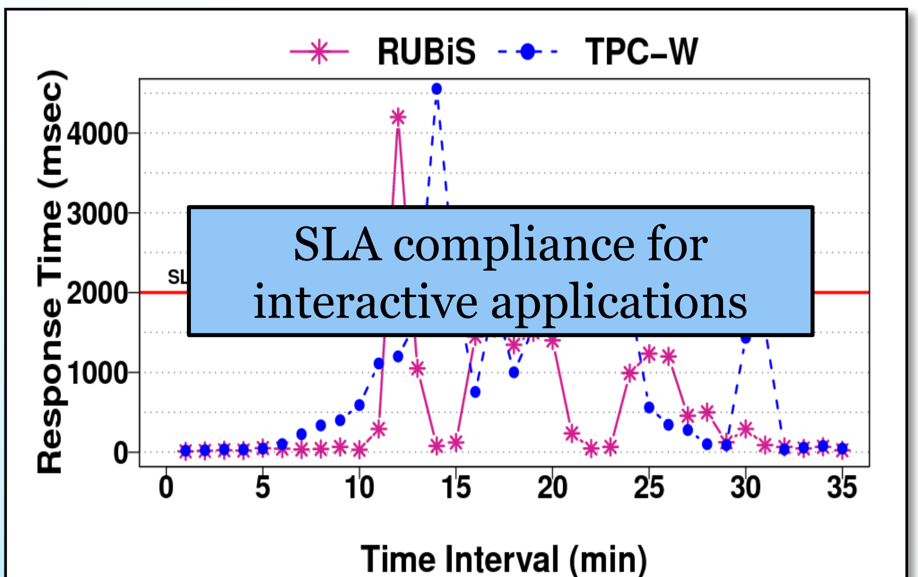

SLA compliance for interactive applications

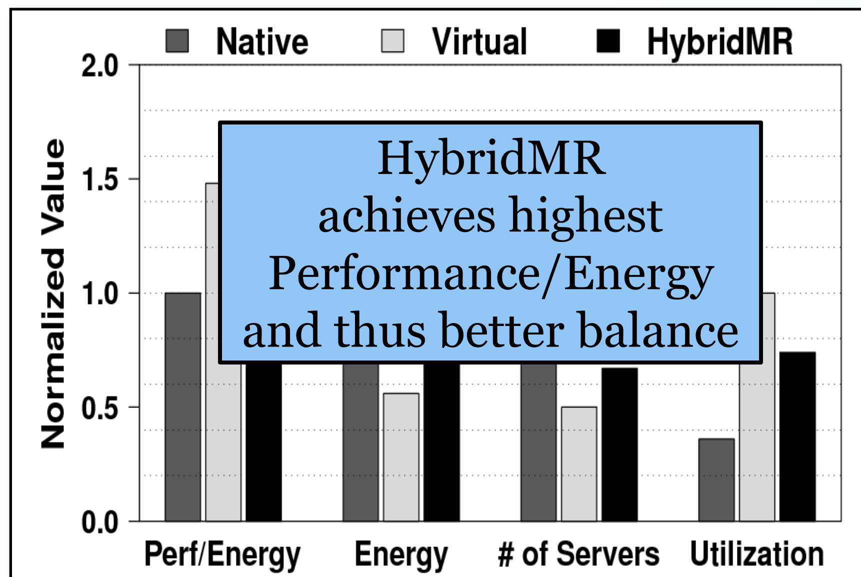

HybridMR achieves highest Performance/Energy and thus better balance

HybridMR Summary

- ✓ Efficient scheduling of workload mix on hybrid compute clusters
- ✓ Investigates Hadoop performance on virtual cluster
- ✓ Dynamic resource management
- ✓ Achieves best of two worlds (native and virtual)

↓ 40%	↑ 45%	↓ 43%
Completion time	**Utilization**	**Energy**

CloudPD : Problem Determination and Diagnosis in Shared Dynamic Clouds

Joint work with IBM Research, India, DSN 2013

Cloud Related Faults

Large data centers and clouds experience frequent faults

Fault diagnosis challenges in virtualized clouds

How to automate problem determination and diagnosis, for a large, shared, dynamic and virtualized clouds?

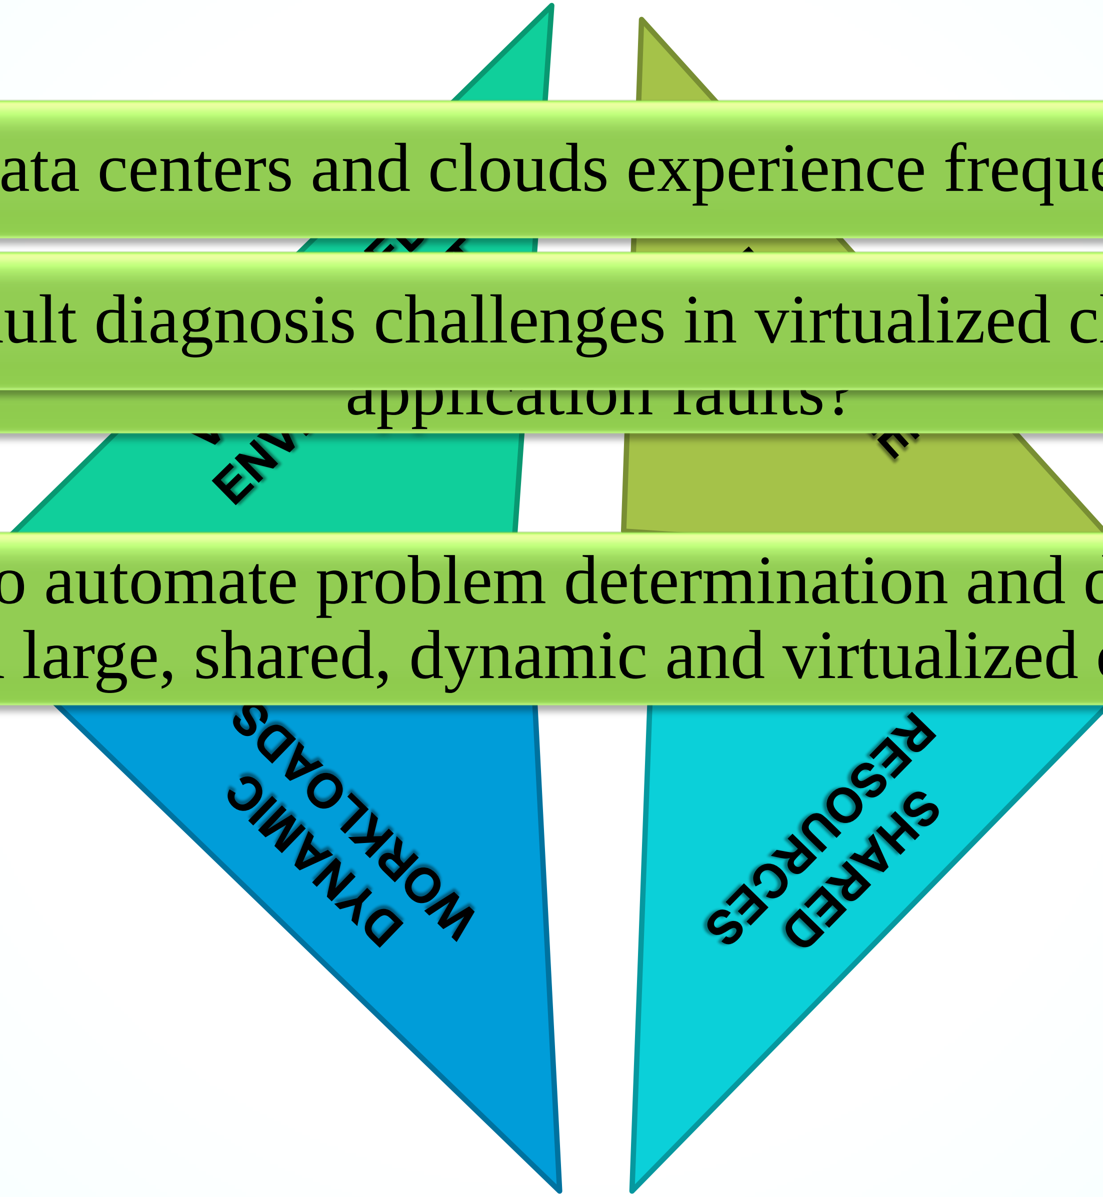

Clouds Usher New Challenges

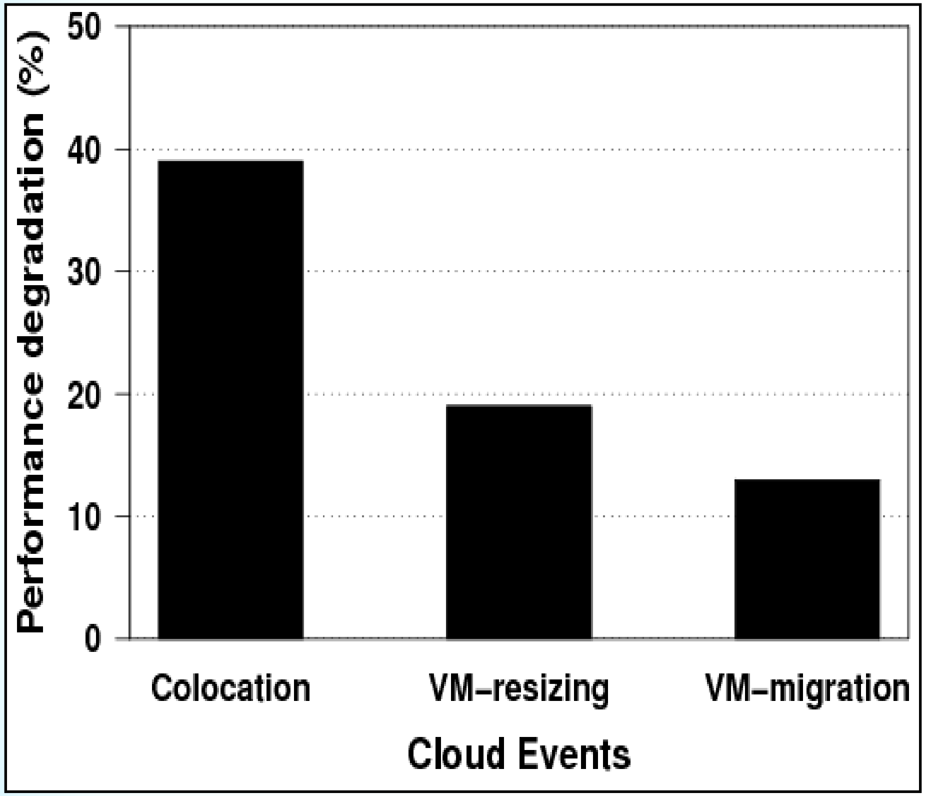

(a) Increase in application latency due to faulty cloud events

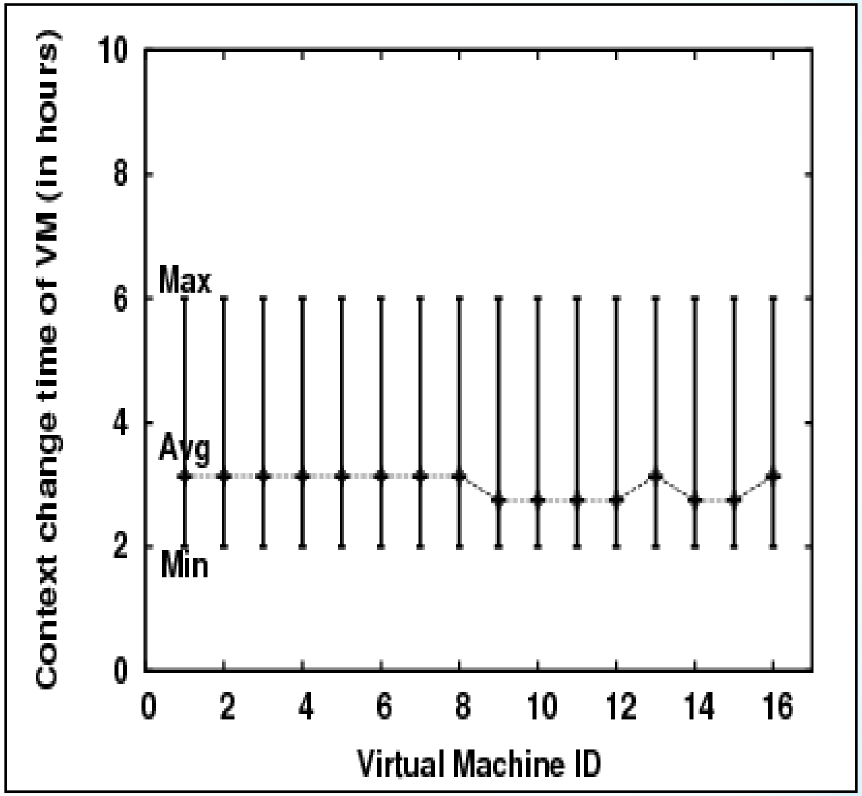

(b) High rate of change of VMs operating context

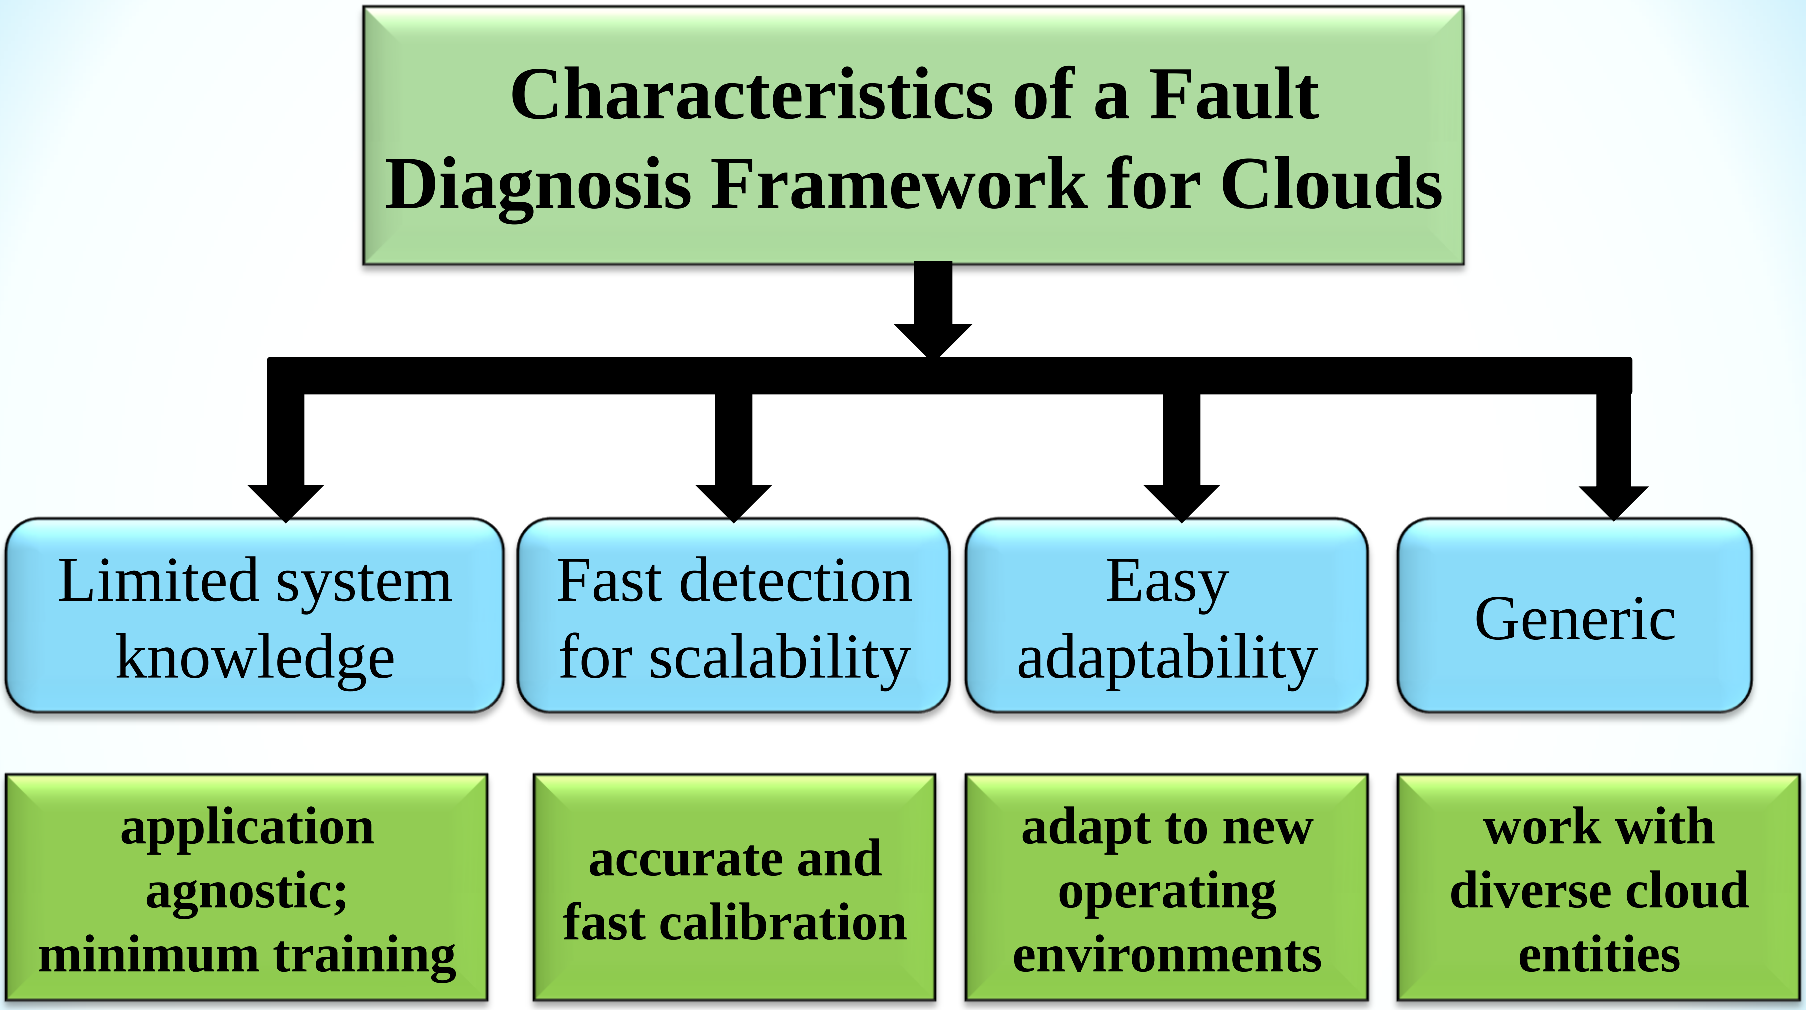
Characteristics of a Fault Diagnosis Framework for Clouds
Limited system knowledge
Fast detection for scalability
Easy adaptability
Generic
application agnostic; minimum training
accurate and fast calibration
adapt to new operating environments
work with diverse cloud entities

System Context

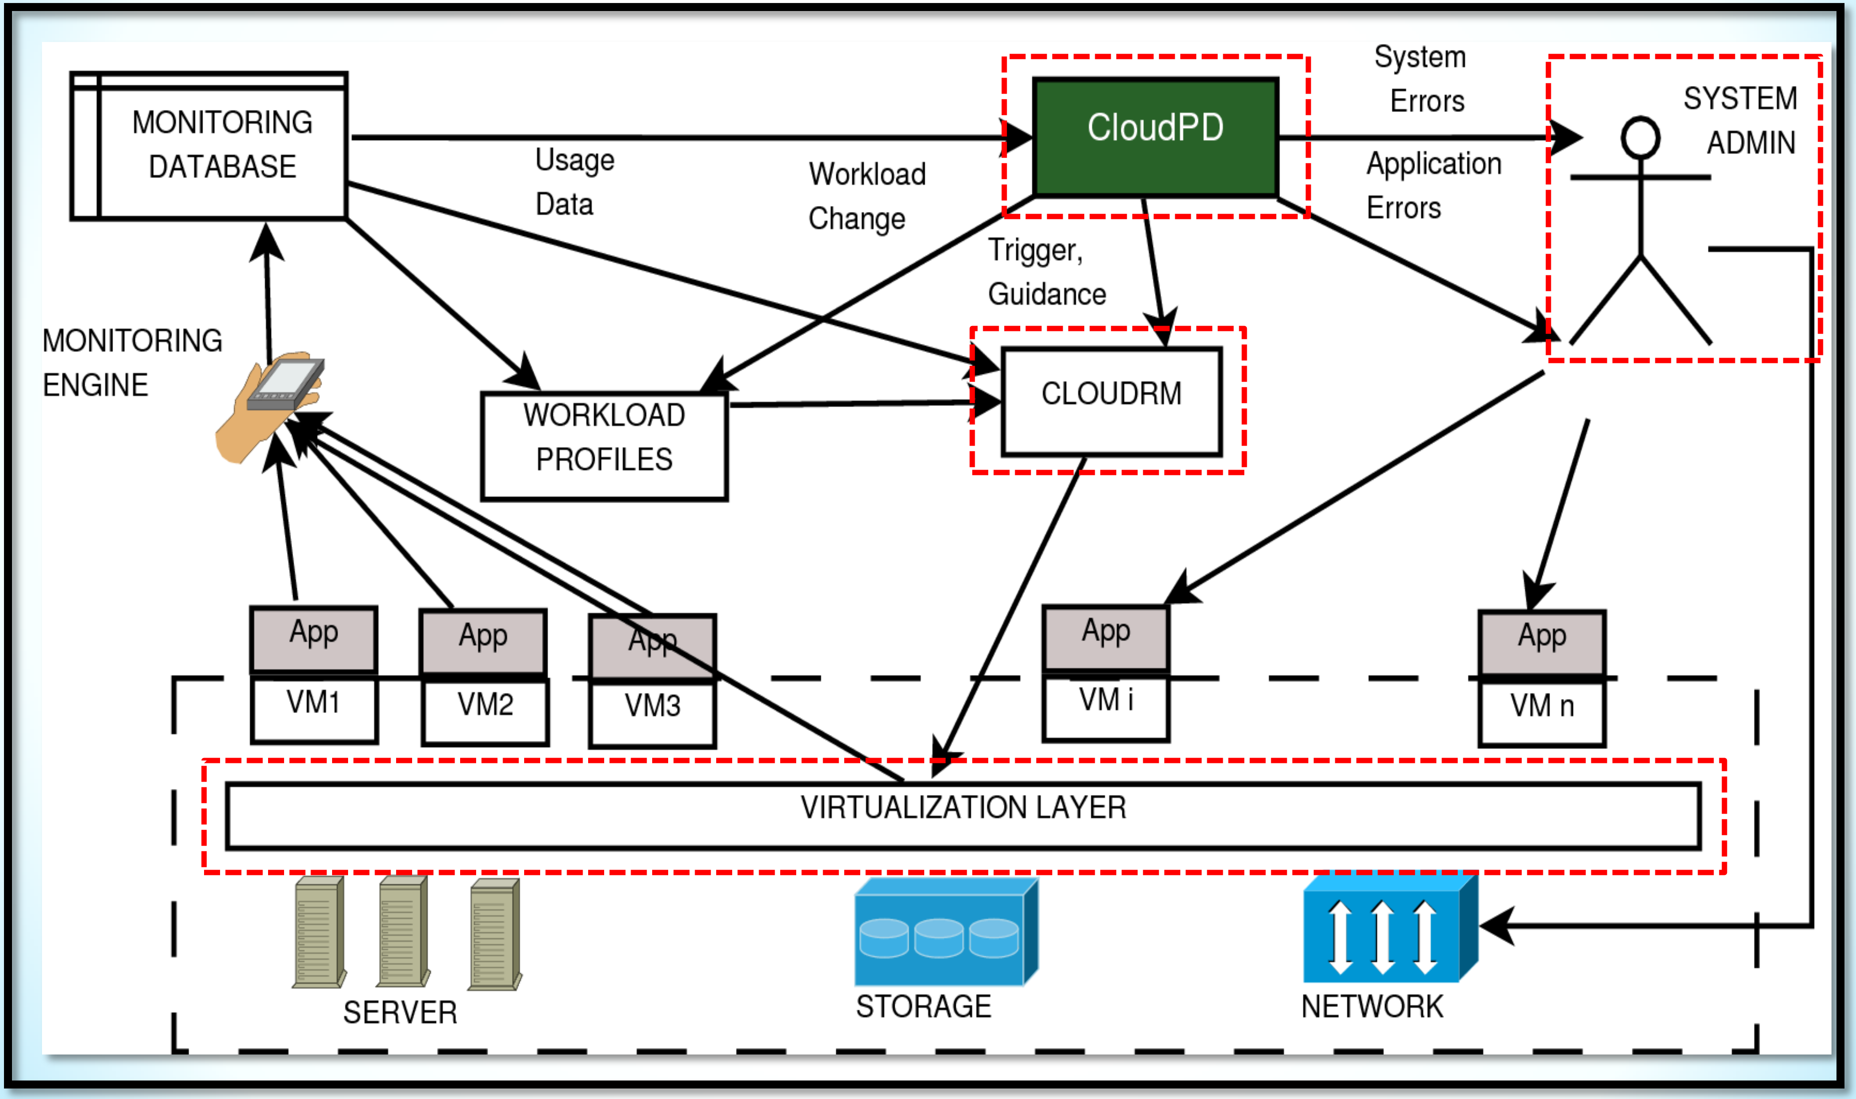

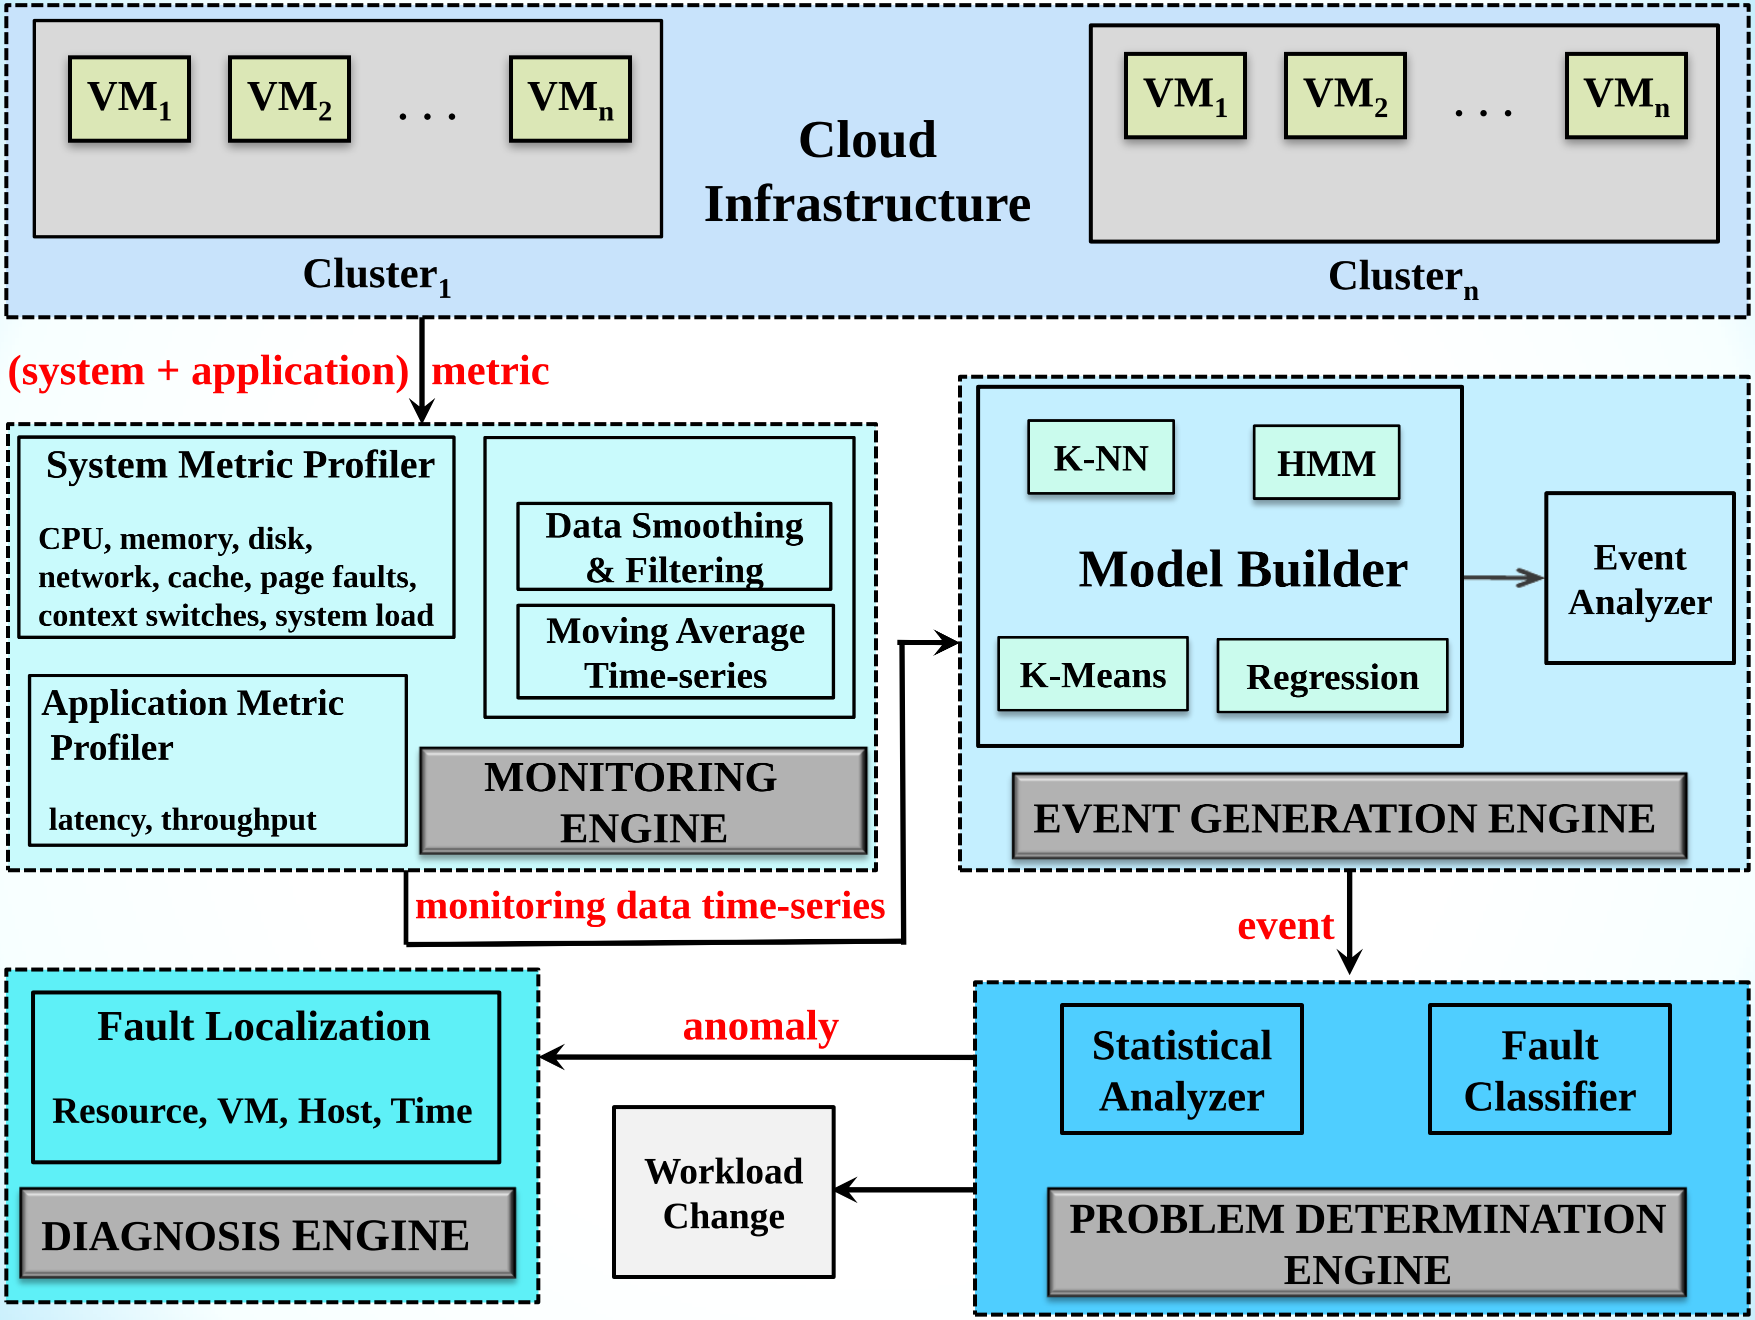

VM$_1$
VM$_2$
. . .
VM$_n$
Cluster$_1$
Cloud Infrastructure
VM$_1$
VM$_2$
. . .
VM$_n$
Cluster$_n$
(system + application) metric
System Metric Profiler
CPU, memory, disk, network, cache, page faults, context switches, system load
Application Metric Profiler
latency, throughput
Data Smoothing & Filtering
Moving Average Time-series
MONITORING ENGINE
K-NN
HMM
Model Builder
K-Means
Regression
Event Analyzer
EVENT GENERATION ENGINE
monitoring data time-series
event
Fault Localization
Resource, VM, Host, Time
DIAGNOSIS ENGINE
anomaly
Workload Change
Statistical Analyzer
Fault Classifier
PROBLEM DETERMINATION ENGINE

Faults Examples

Cloud-related faults	Non-cloud/Application faults
Impact due to resource sharing	Misconfigured application
Wrong VM sizing	Software bugs
Incorrect VM reconfiguration	Application or OS update
Faulty VM migration	Anomalous workload change

Competitive Methodologies

- Baseline B1: no operating context notion; only considers VM CPU and memory + CloudPD's three stages
- Baseline B2: *oracle*; analyzes every interval in detail
- Baseline B3: no correlation across peers
- Baseline B4: uses static thresholds to trigger events

Evaluation Metrics	Definition
Recall	# of successful detections / total # of anomalies
Precision	# of successful detections / total # of alarms
Accuracy	2 * Recall * Precision / Recall + Precision
False Alarm Rate	# of false alarms / total # of alarms

Results

Method	# of correct normal detections	# of correct anomalous detections	# of correct Phase 1 detections	# of total predicted anomalies	Recall	Precision	Accuracy	False Alarm Rate
CloudPD	67	18	21	24	0.78	0.75	0.77	0.25
B1	58	10	14	25	0.43	0.40	0.42	0.60
B2	67	21	23	27	0.91	0.78	0.84	0.22
B3	60	11	21	24	0.48	0.46	0.47	0.54
B4	60	13	15	26	0.57	0.50	0.53	0.50

Comparing end-to-end diagnosis effectiveness of CloudPD for a 24-hour enterprise trace-based case study

CloudPD Summary

- ✓ CloudPD is a problem determination framework for clouds
- ✓ Introduces the notion of *operating context*
  Hierarchical architecture to address massive scale
- ✓ Integrates with cloud manager for remediation actions
- ✓ Comprehensive evaluation with representative Web 2.0
- ✓ Achieves

< 20%	**85%**	**< 30 sec**
False positives	**Accuracy**	**Analysis time**

Conclusions

- **Why is research in clouds important?**
  - Cost-effective and flexible business model
  - Numerous challenges and umpteen research opportunities
- Performance and reliability in clouds are major concerns
  - Characterization of cloud workloads to better understand their performance impact
  - Effective resource management and scheduling for cloud-based MapReduce clusters and hybrid data centers
  - Efficient end-to-end reliability management in clouds
  - A preliminary performance model (D-factor)

Future Research Directions

- Heterogeneity-aware scheduling and resource management in cloud-based clusters
- Analytical modeling of MapReduce performance in hybrid data centers
- Better diagnosis and classification of faults in large-scale virtualized clouds
- Optimizing MapReduce deployment in shared memory systems with focus on network communication (NoCs)
- Many more …

Thank You!

Questions?