Quantifying Load Imbalance on Virtualized Enterprise Servers

Emmanuel Arzuaga and David Kaeli

Department of Electrical and Computer Engineering
Northeastern University
Boston MA

NUCAR

Traditional Data Centers

- Physical Enterprise Servers
  - application dedicated
    - multi-tiered: webserver, database, etc.
  - low resource efficiency
  - high cost of operation

2

Virtualized Data Centers

- Virtualization & Server Consolidation:
  - increase system efficiency
  - reduce the number of physical servers
  - reduce cost of operation

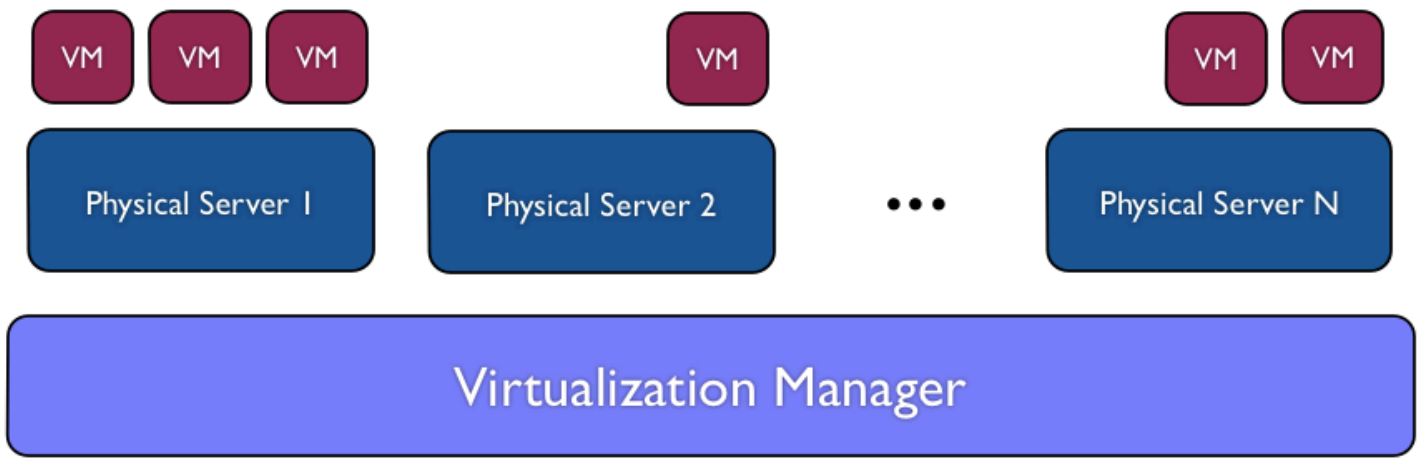

NUCAR

3

Live VM Migration

- VM migration can be used to:
  - maintain a balanced load on the system
  - enable service level agreements (SLA)
  - enhance application performance
  - dynamically tune system to specific workload mix

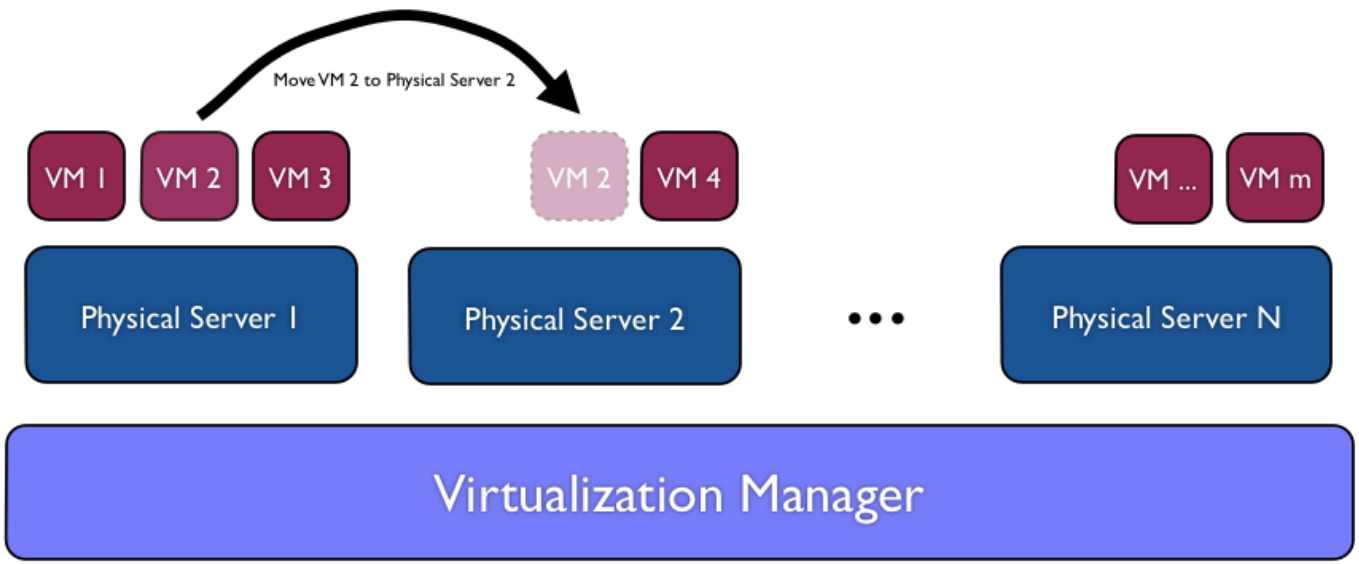

4

Open Questions

- The ability to do live VM migrations is not enough, we still need to decide:
  - what to migrate?
  - where to migrate?
  - when to migrate?
- These decisions often depends on system load

NUCAR

5

Presentation Outline

- Load of a Virtualized Enterprise Server
- Imbalance in Cluster of Virtualized Servers
- Using VM Migration for Load Balancing
- Workload Characteristics
- Experimental Results
- Conclusions and Future Work

6

Virtualized Server Load

- Capturing the load of a Virtualized Enterprise Server:
  - let S be the set of physical servers and VM_{Host} the set of VMs currently assigned to server *Host*
  - the overall server load metric is the VM-to-host usage rate:

$$VSL_{Host} = \sum_{resource} W_{resource} \times \frac{\sum_{v \in VM_{Host}} v_{resource}usage}{Host_{resource}capacity}$$

where *resource* is {*CPU, Memory, Disk*}

- VSL_{Host} varies dynamically depending on the current load of the system
- We can use it to balance the loads of multiple servers

NUCAR

7

Imbalance in Virtualized Cluster

- Based on VSL_{Host}, we can define a load set L containing VSL_{Host} values $\forall$ $servers \in S$
- We want to measure how balanced L is at a particular time
  - our approach is to use the coefficient of variation:

$$I_{Metric} = C_L = \frac{\sigma_L}{\mu_L}$$

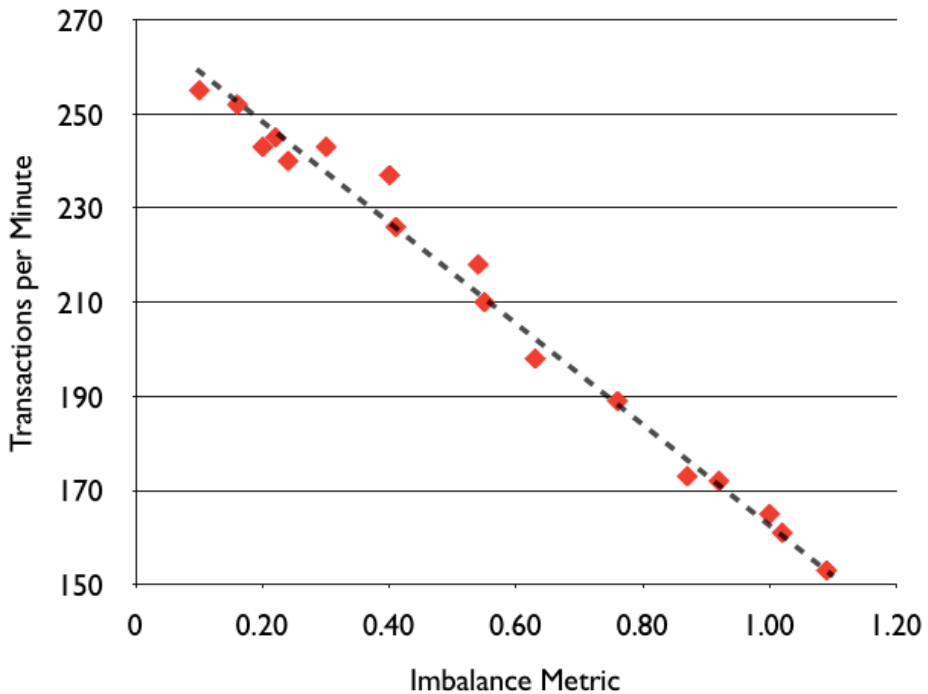

Virtualized Server Load Balancing

- We are interested in the use of live VM migration for load balancing
- The migration criteria is to migrate a VM to a different host if the system is imbalanced according to the imbalance metric (I_{Metric})
- Our problem can be stated as:
  - migrate VM *v* from Host Source *src* to Host Target *target* such that I_{Metric} is reduced

NUCAR

9

VSL Inductive Balancing Method (*VIBM*)

- Inductively predict future state:
  - calculate $I_{Metric\ PREDICTED}$ if we move *v* to *target* and choose the move that provides the lowest value for I_{Metric} ,

$$VSL_{target} \longleftarrow VSL_{target} + v_{candidate}$$
$$VSL_{src} \longleftarrow VSL_{src} - v_{candidate}$$

- Follow a greedy approach
  - improve future state given current state

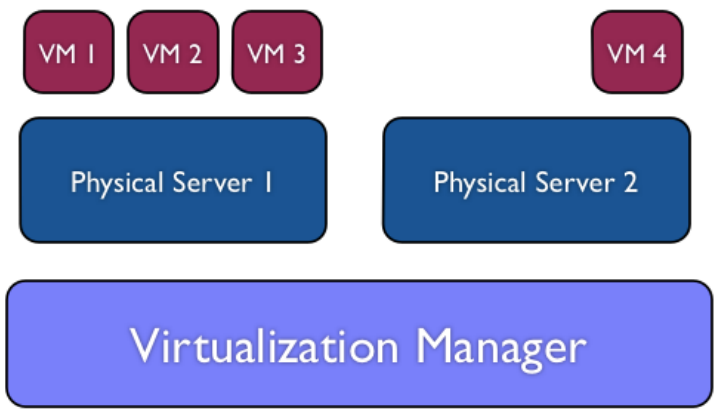

I_{Metric}	0.83
VM	$I_{Metric\ PREDICTED}$
VM 1	0.65
VM 2	0.43
VM 3	0.55

10

Workload Characteristics

- Enterprise servers run varied types of applications
  - database, webserver, application server
- Applications behave differently in terms of resource usage
  - even same application may change resource consumption over time (burstiness)
- I_{Metric} can account for such changes

NUCAR

11

Experiments

- Understand I_{Metric} in terms of:
  - accuracy of predicting future system state
    - simple & complex workloads
  - how it relates to other resource management solutions (VMware DRS)
- Evaluate possible performance improvement using *VIBM*

NUCAR

12

System configuration

- ESX 3.5 Servers
  - dual processor, dual core Intel Xeon 2.33 Ghz
  - 4 GB main memory
- shared iSCSI SAN VM
  - 700 GB capacity
- VMware VCenter Server
  - *VIBM* Migration Handler

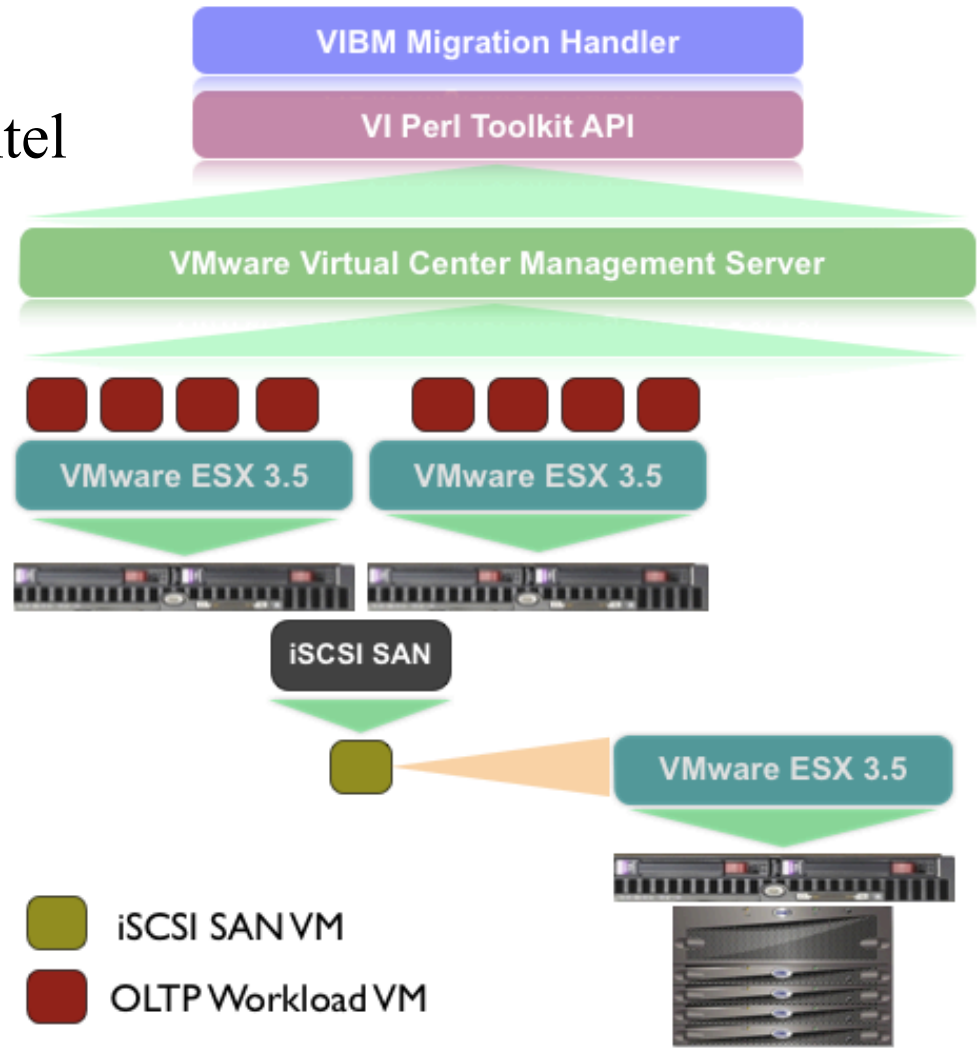

VM Configuration

- Two different VM configurations:
  - 2 VCPU 1GB 50GB (large)
  - 1 VCPU 512MB 50GB (small)
- Experiments with two different VM sets:
  - 6VM (2-large,4-small)
  - 8VM (3-large,5-small)
- Initial VM Placement (all VMs running on one host)

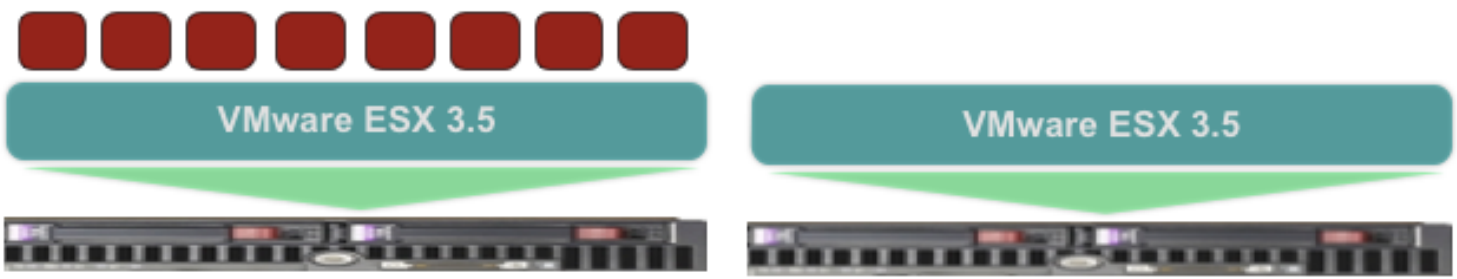

NUCAR

14

Workloads

- Simple CPU intensive workload
  - I_{Metric} prediction tests
- Online Transaction Processing (OLTP)
  - TPC-C based
  - wholesale supplier managing orders
  - transactions show random I/O behavior

I_{Metric} Analysis: Prediction

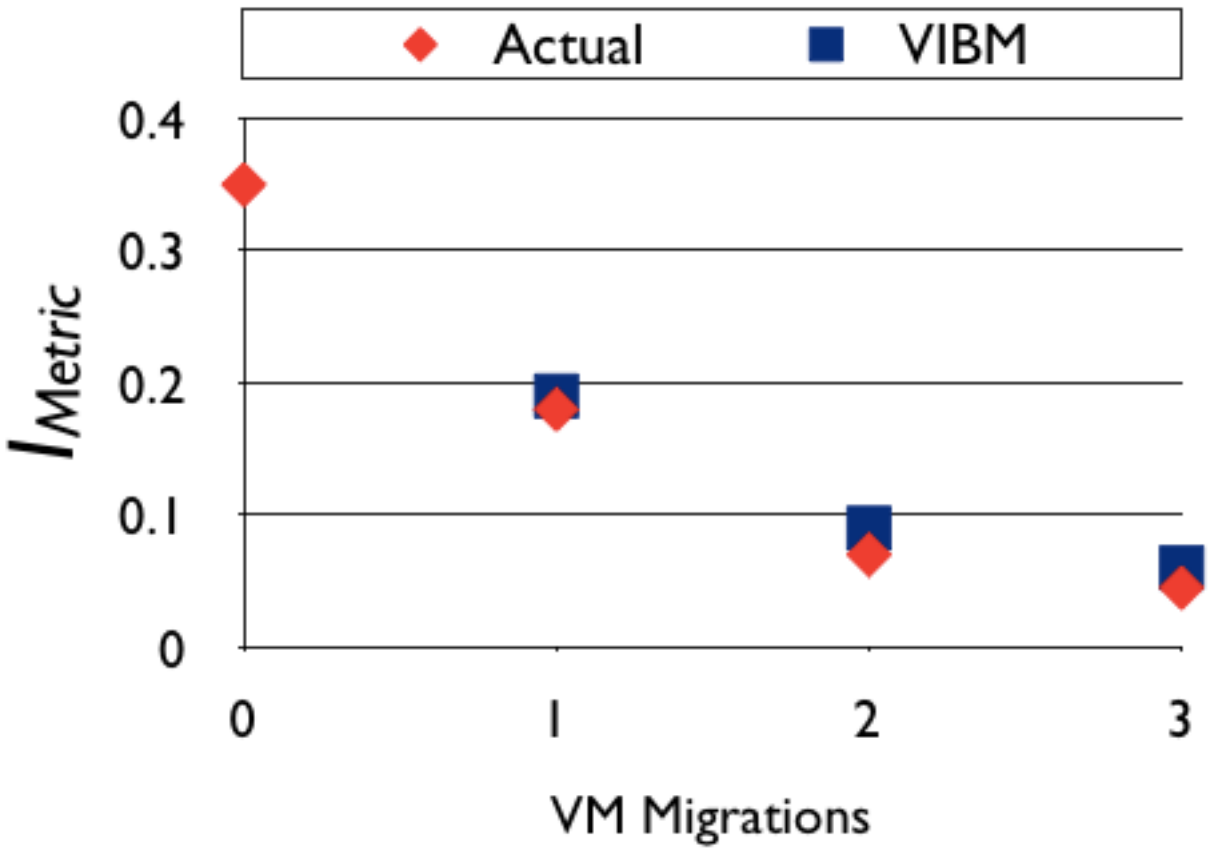

6VM

I_{Metric} Analysis: *VIBM* vs *DRS*

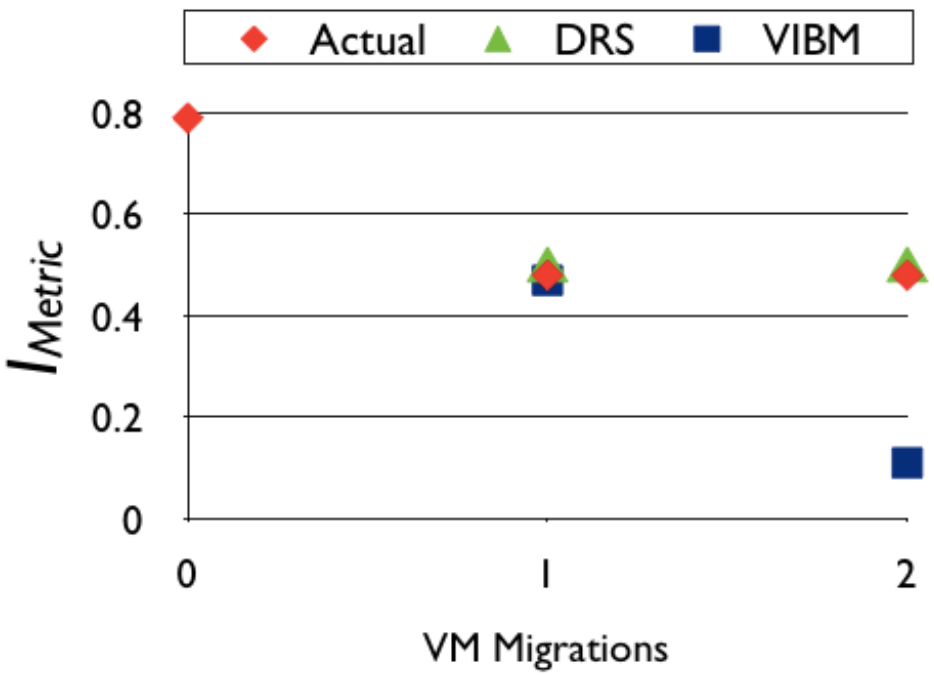

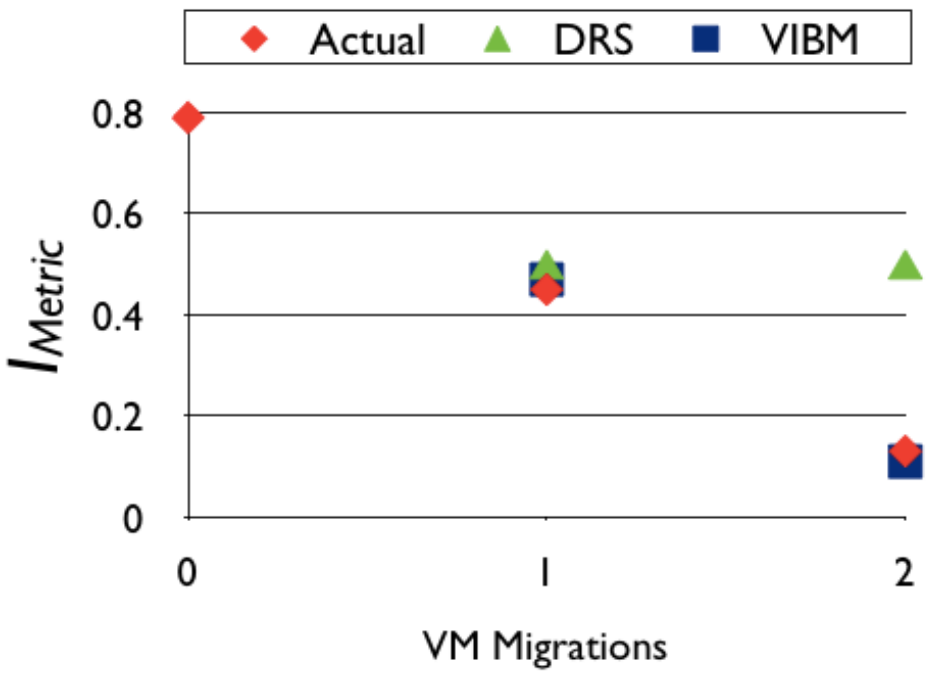

6VM

17

I_{Metric} Analysis: Throughput Test

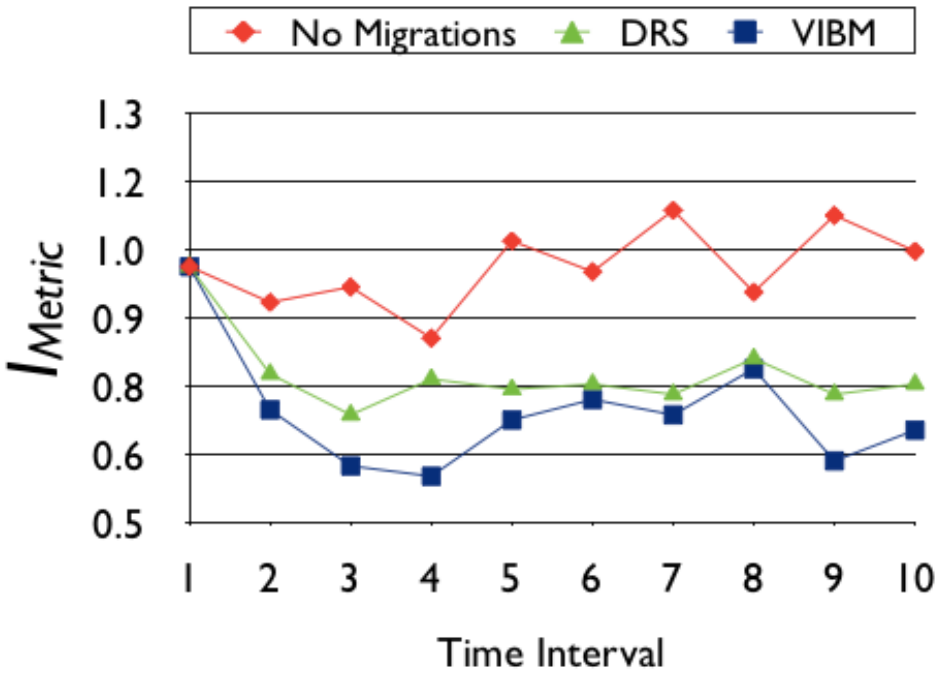

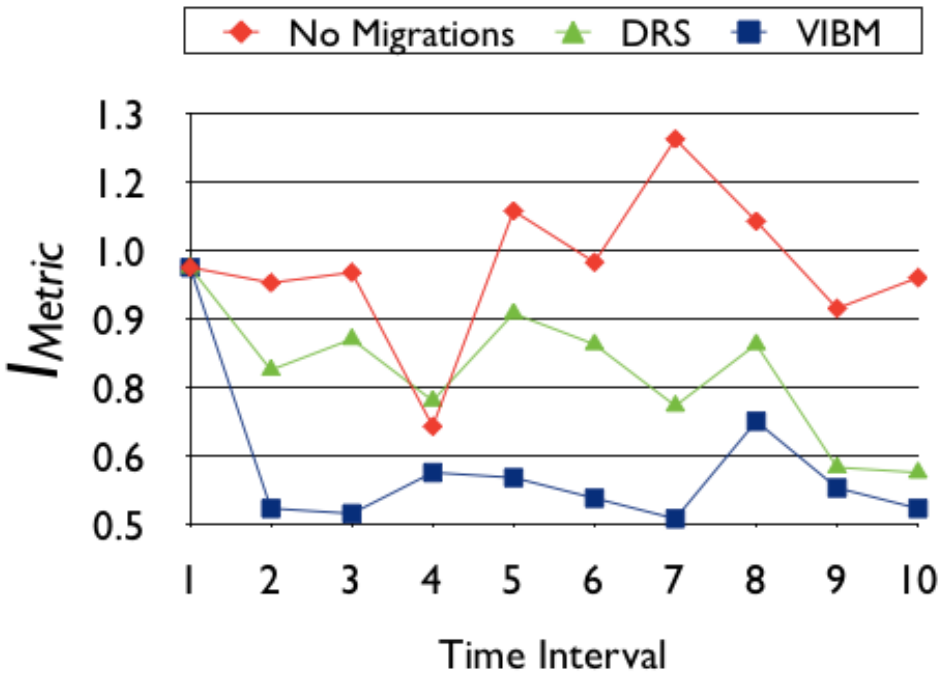

6VM	8VM

18

Throughput Test Results

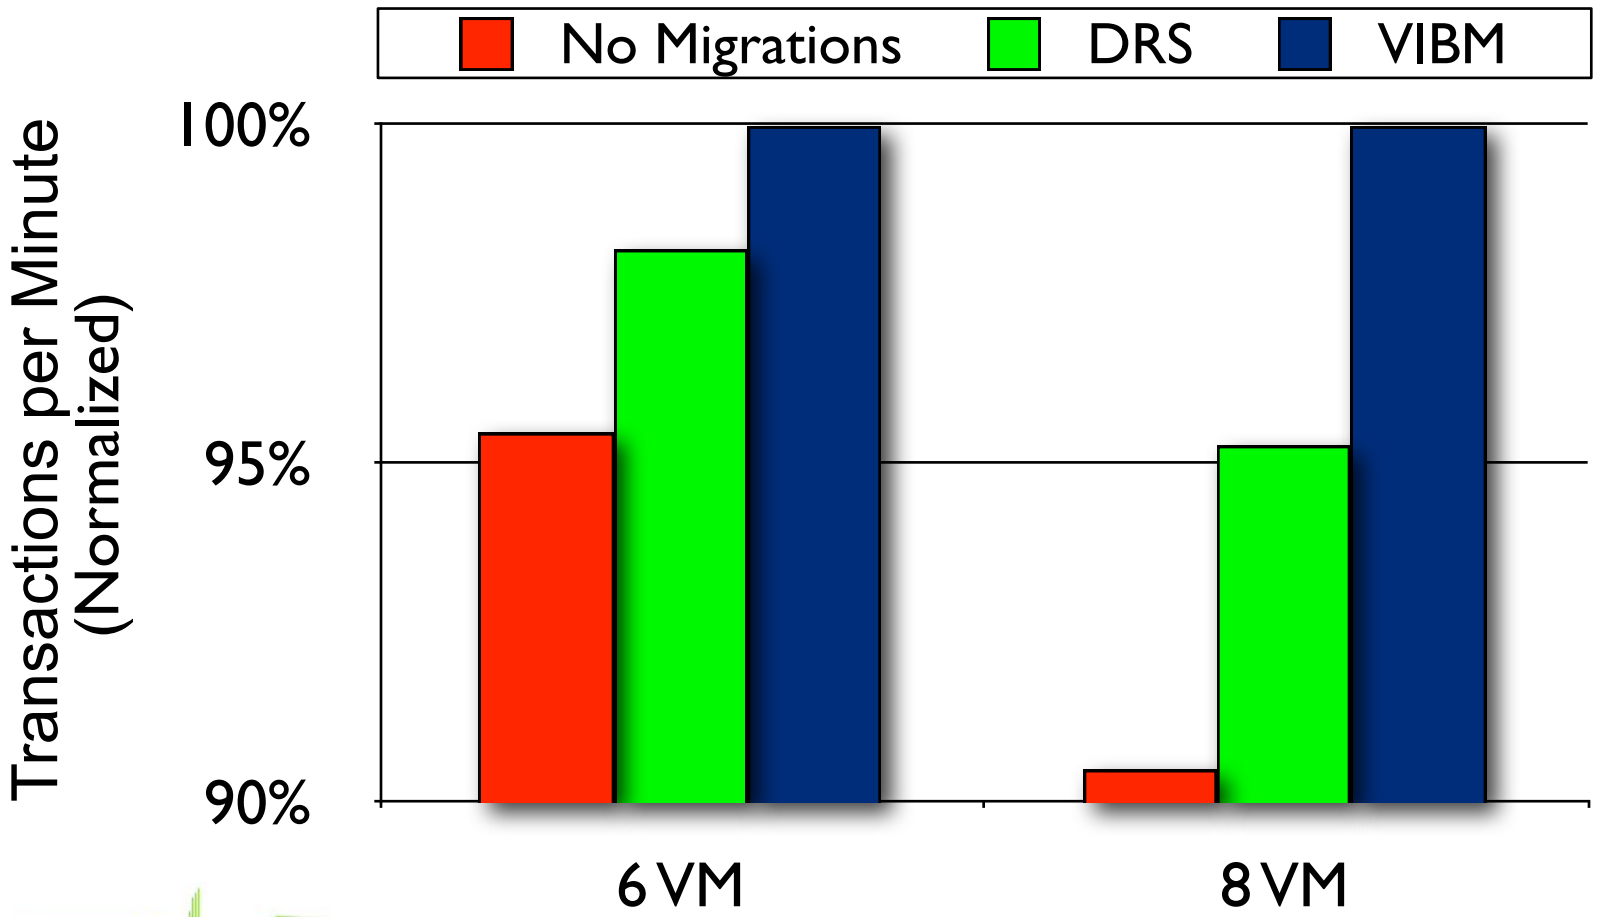

19

Conclusions

- Presented a load metric for virtualized enterprise servers: VSL_{Host}
- Built a Load Balancing Scheme based on VSL_{Host}: *VIBM*
- *VIBM* produced migration patterns that improved system balance and throughput superior to VMware DRS
- Future work includes the extension of VSL_{Host}
  - enable *VIBM* to suggest migrations that reduce power consumption
  - resource weights analysis
  - increase the workload mix inside VMs

NUCAR

20

Acknowledgements

- This work was supported in part by:
  - NSF Major Research Instrumentation Grant (Award Number MRI-0619616)
  - Institute for Complex Scientific Software
  - Gordon-CenSSIS, the Bernard M. Gordon Center for Subsurface Sensing and Imaging Systems, under the Engineering Research Centers Program of the National Science Foundation (Award Number EEC-9986821)
  - VMware through its VMAP program

ICSS
Institute for Complex Scientific Software

VMware Academic Program

NUCAR

21

Thank you!

22