VirtualPower: Coordinated Power Management in Virtualized Enterprise Systems

Ripal Nathuji and Karsten Schwan
CERCS Research Center
Georgia Institute of Technology

SOSP 2007
October 16, 2007

Need for Datacenter Power Management

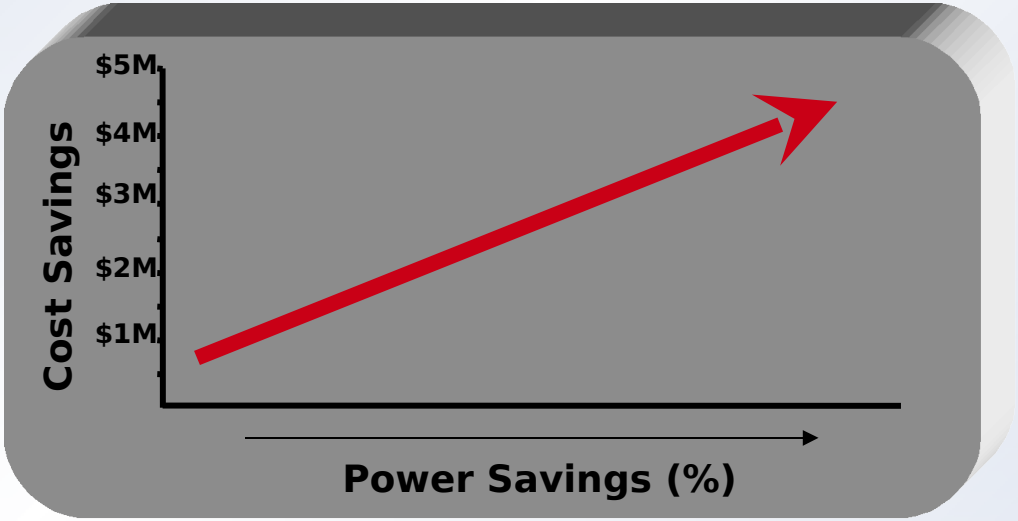

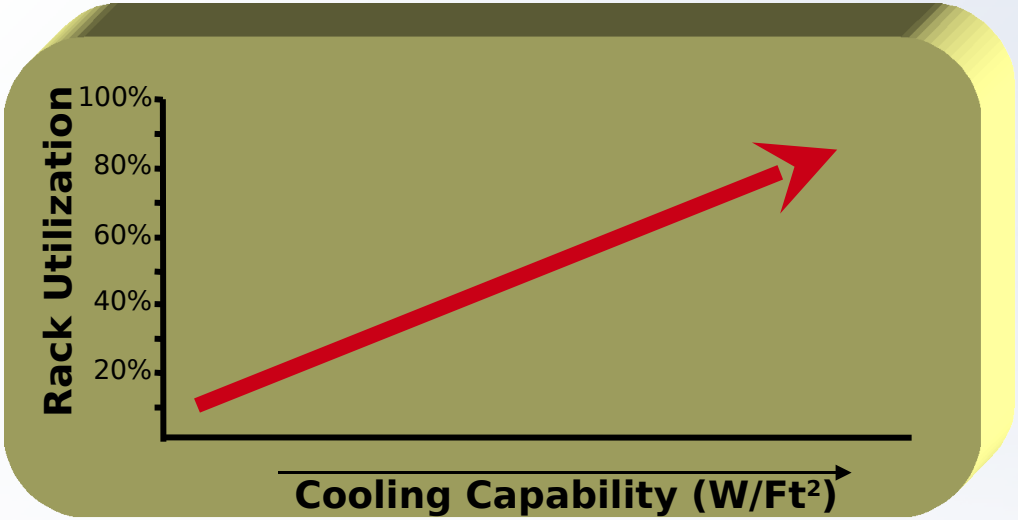

Power Management Ecosystem

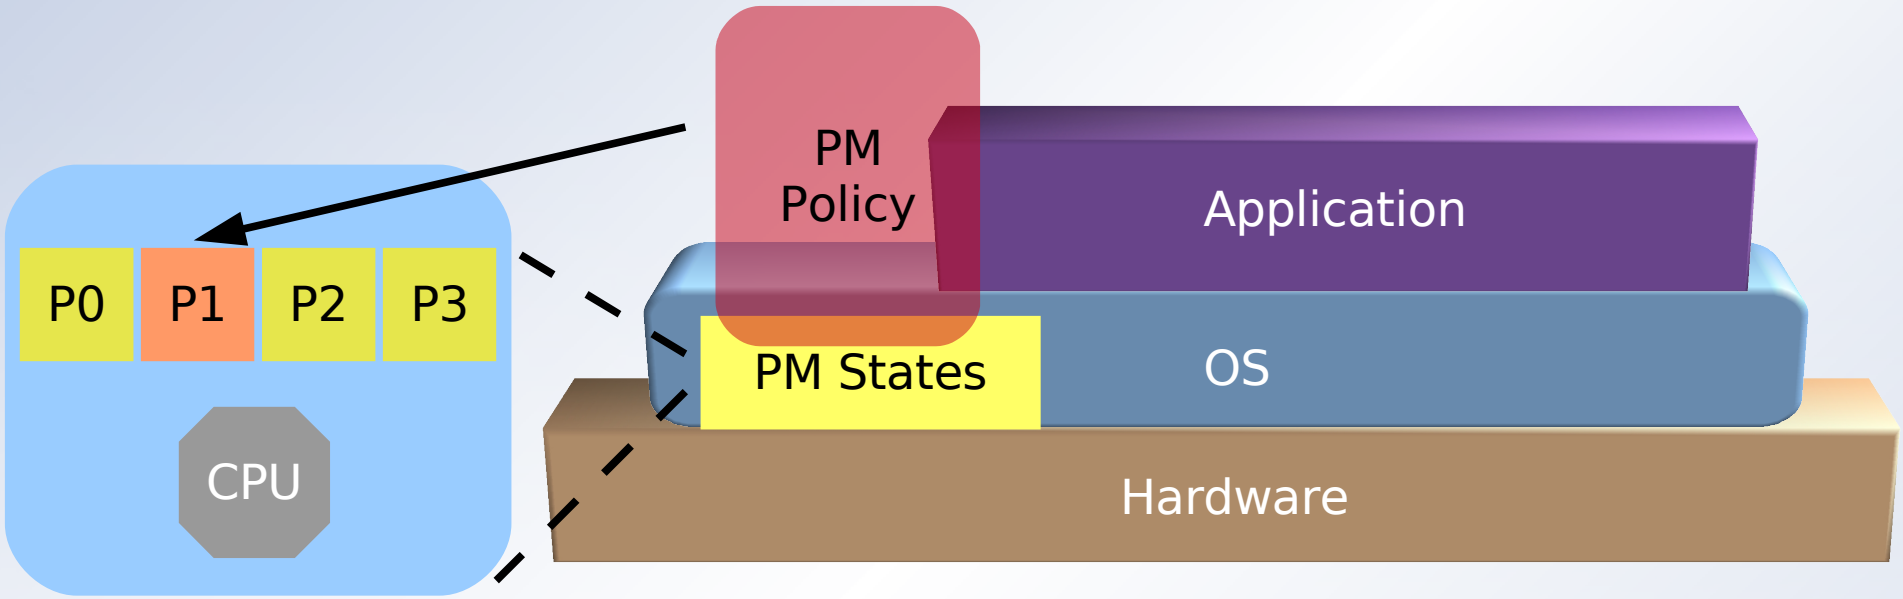

- ACPI exports hardware states (e.g. Px states), with increasingly manageable components beyond CPU
- Investment into application specific power management (PM) policies
- Explicit awareness/modification of hardware states directly impacts platform power consumption

Power Management with Virtual Machines

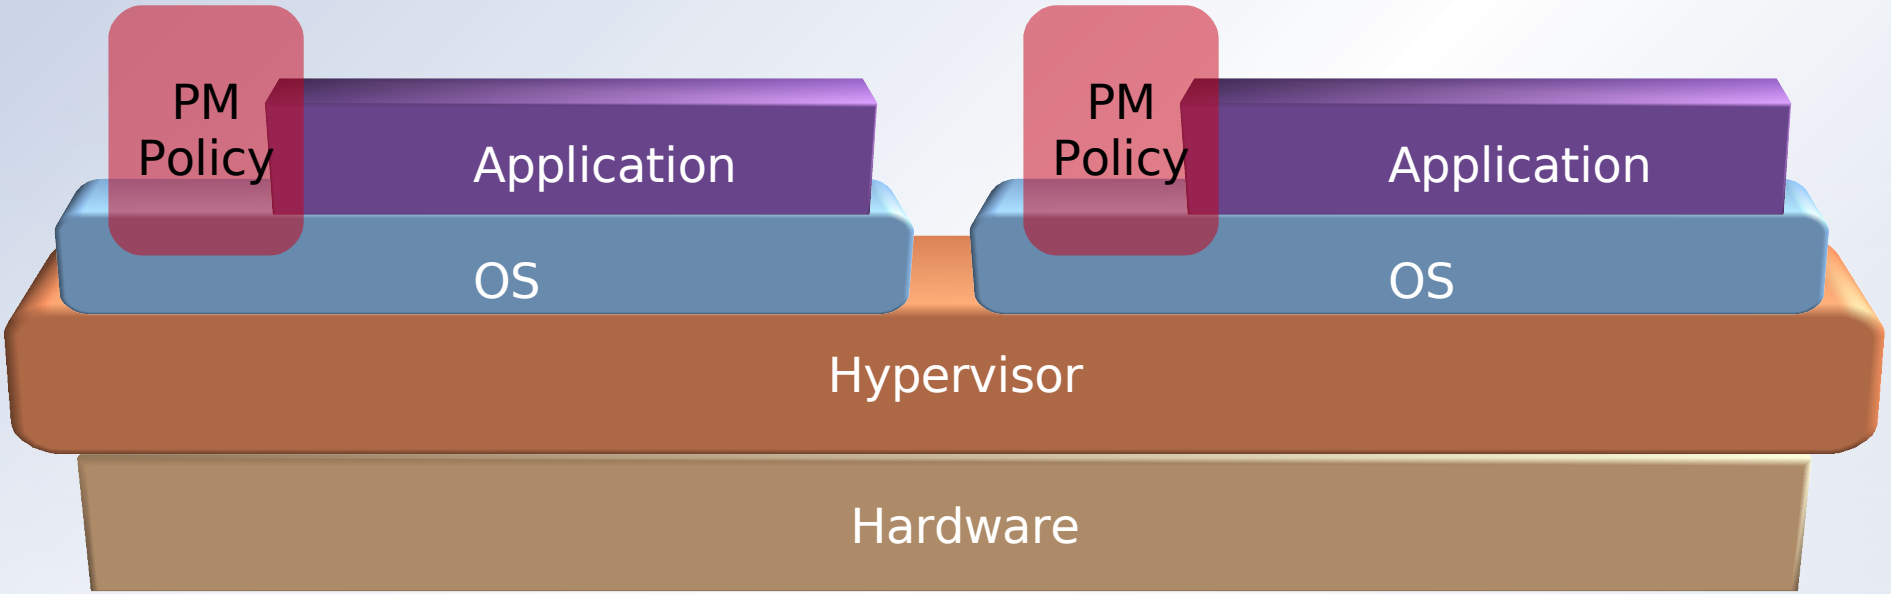

Goal: Continue leveraging existing ecosystem/PM policies

Problem
What manageability to expose?
How to use hardware states without violating isolation?
How to obtain power benefits with VM resource sharing?

VirtualPower Solutions and Opportunities

Problem	Solution
What manageability to expose?	VPM states
How to use hardware states without violating isolation?	VPM channels
How to obtain power benefits with VM resource sharing?	VPM mechanisms

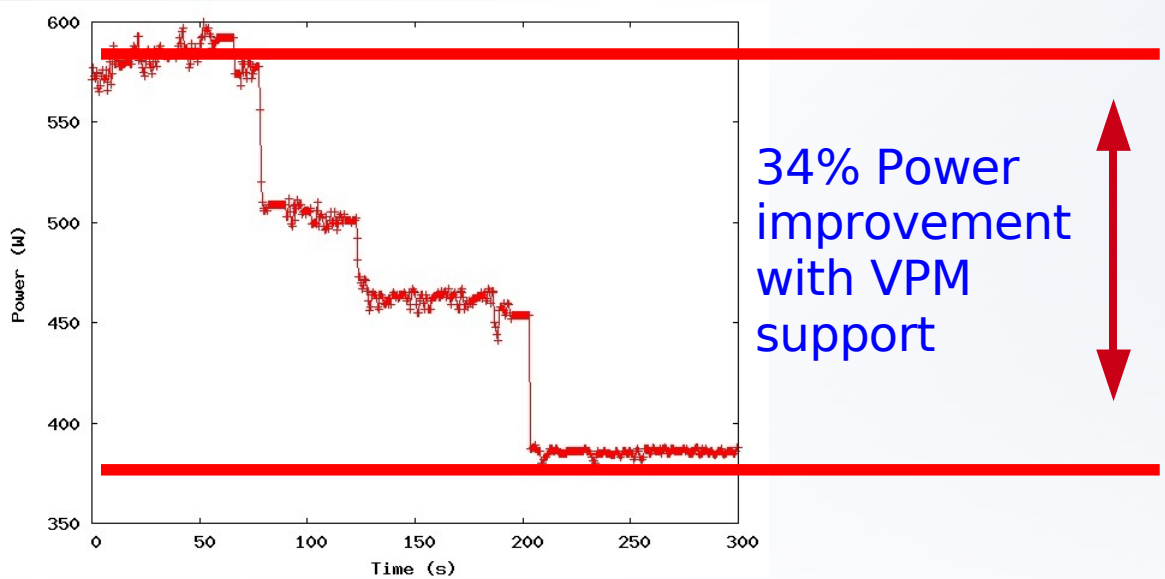

Heterogeneity in Modern Datacenters

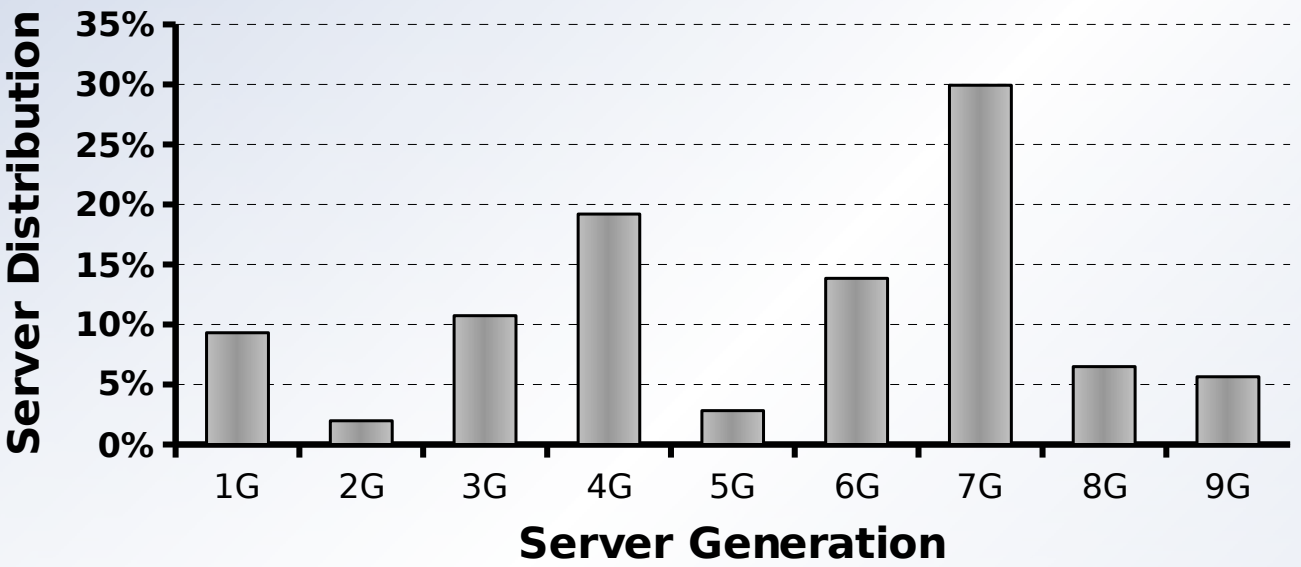

[ICAC2007]

Platform heterogeneity

- Caused by upgrade cycles/failures
- Variations in power, performance, and manageability

Problem: VM Management View with Heterogeneity

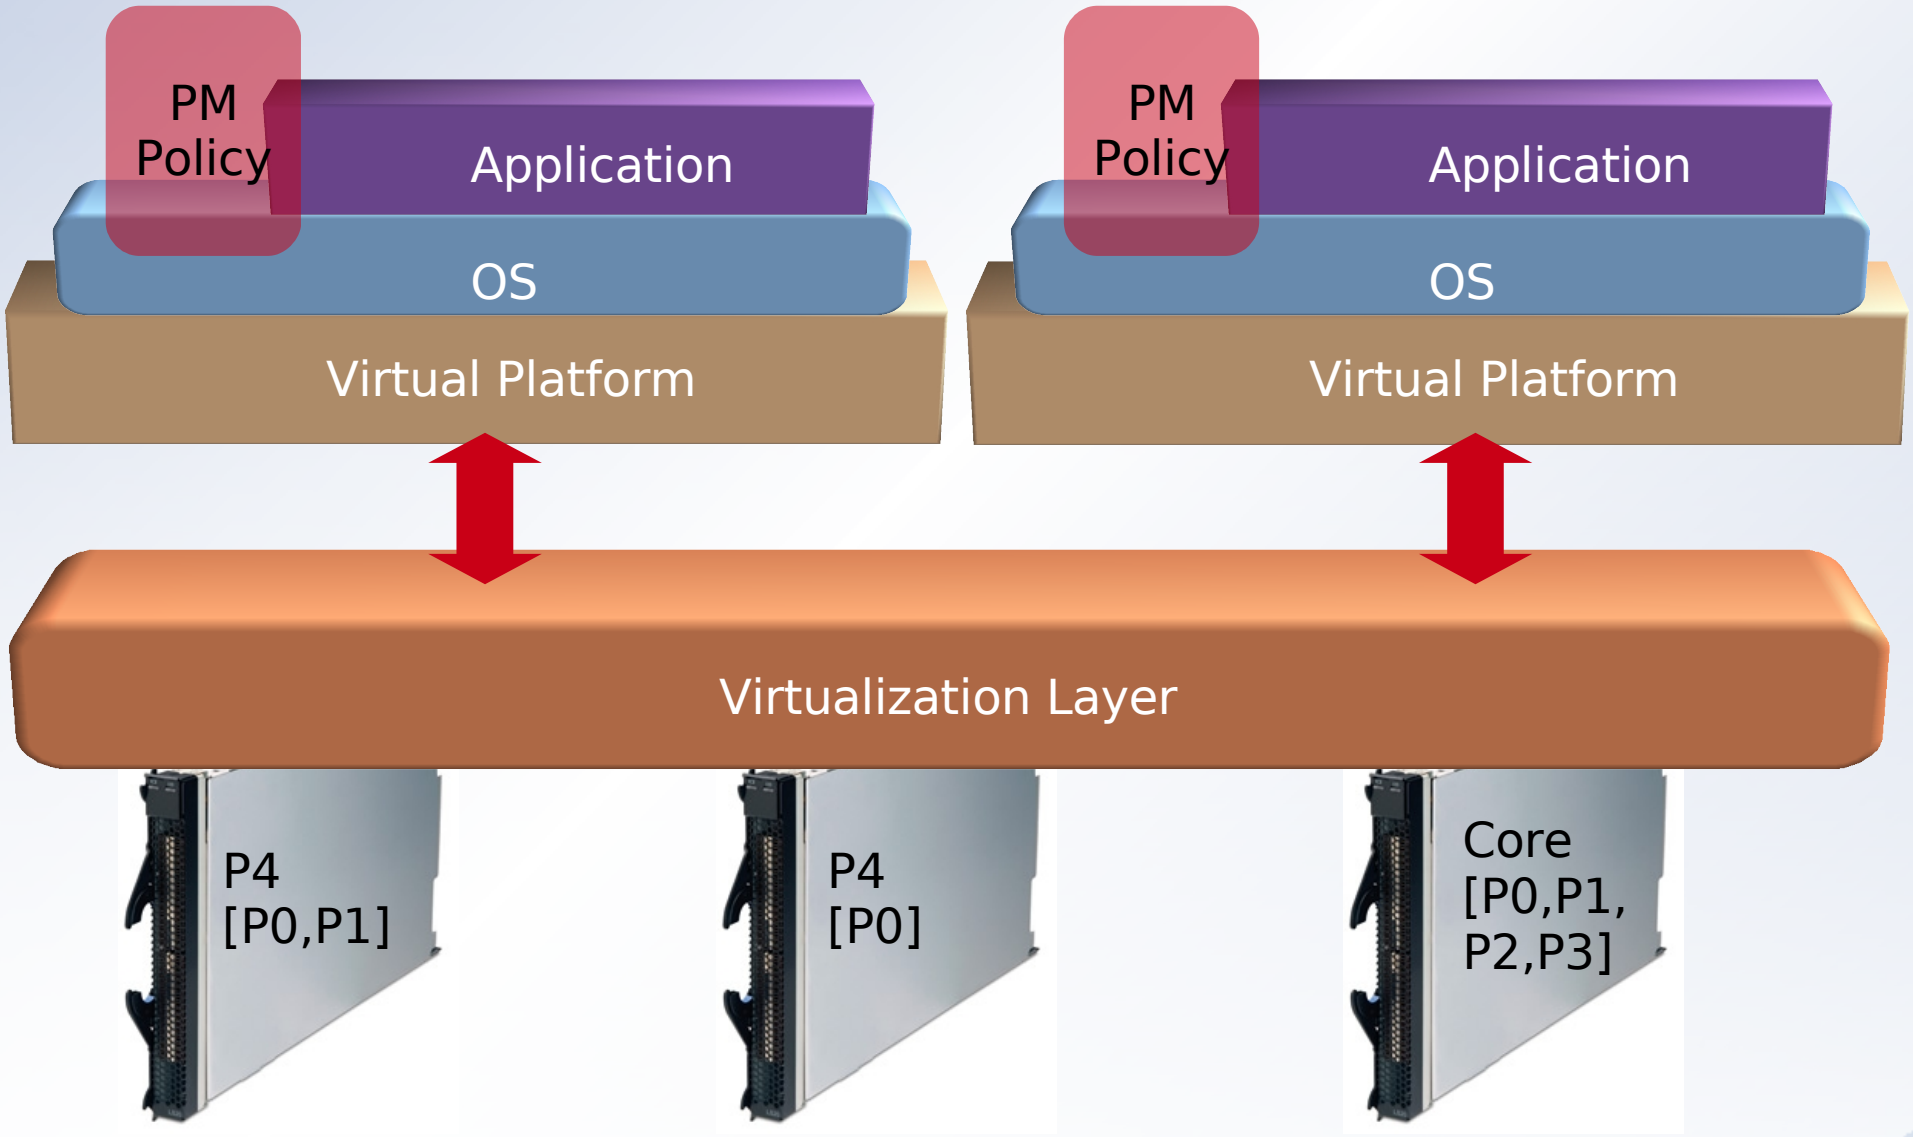

Solution: VPM States

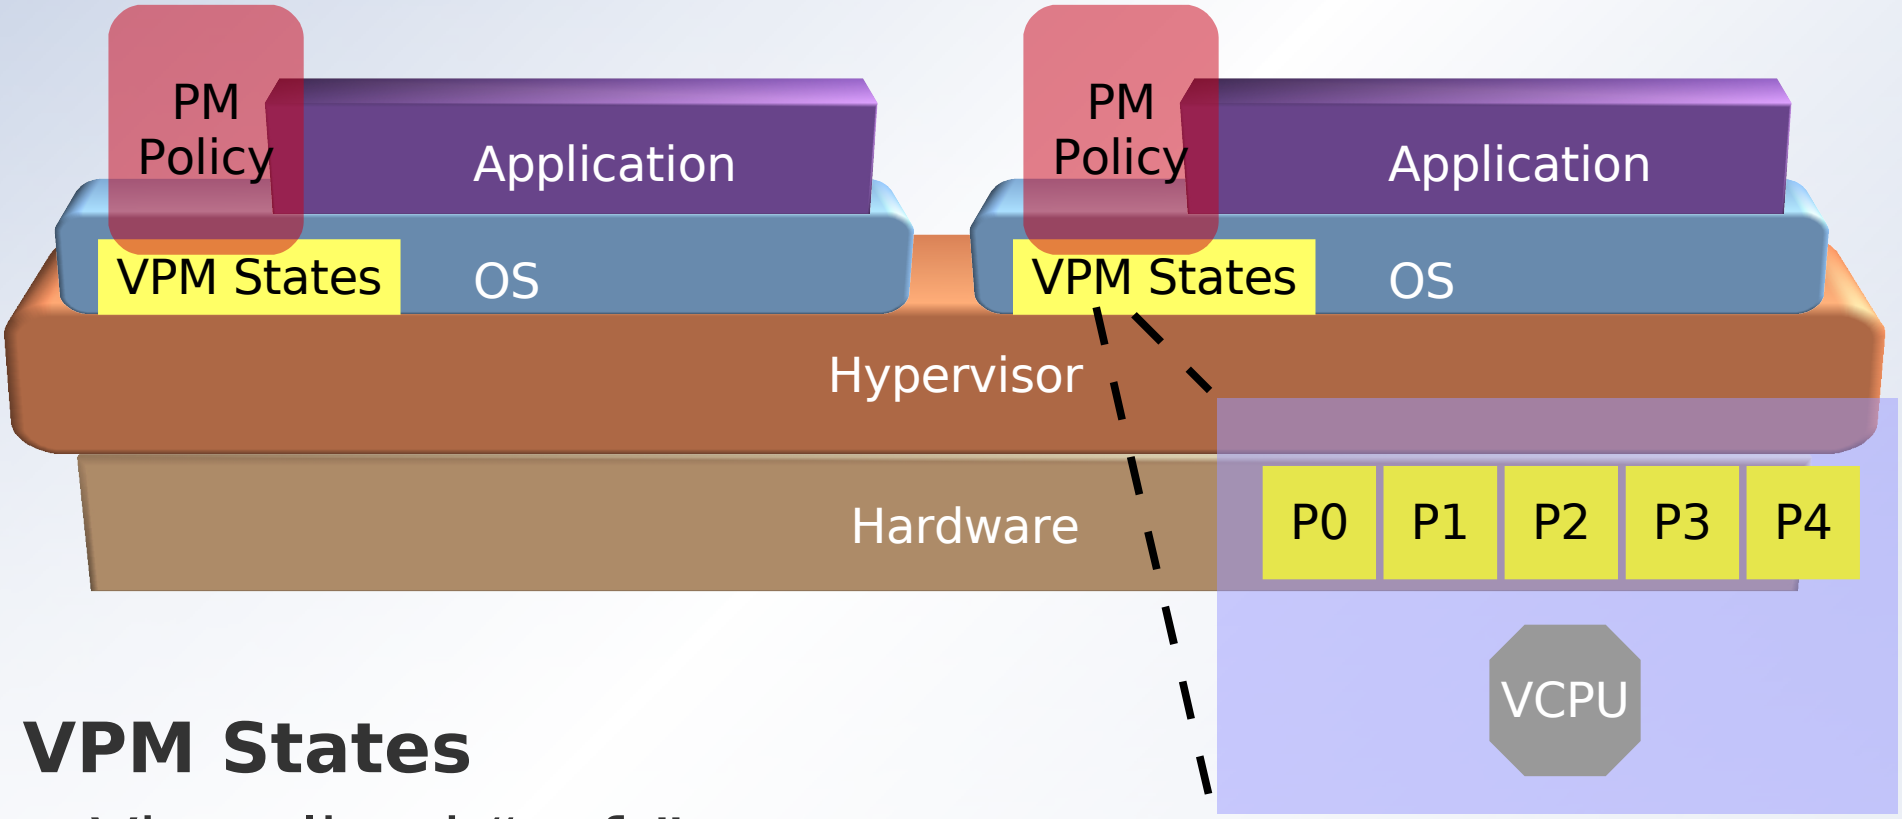

VPM States

- Virtualized “soft” states
- Provide consistent view of manageability across migrations

Problem: PM Policies and Isolation + Independence

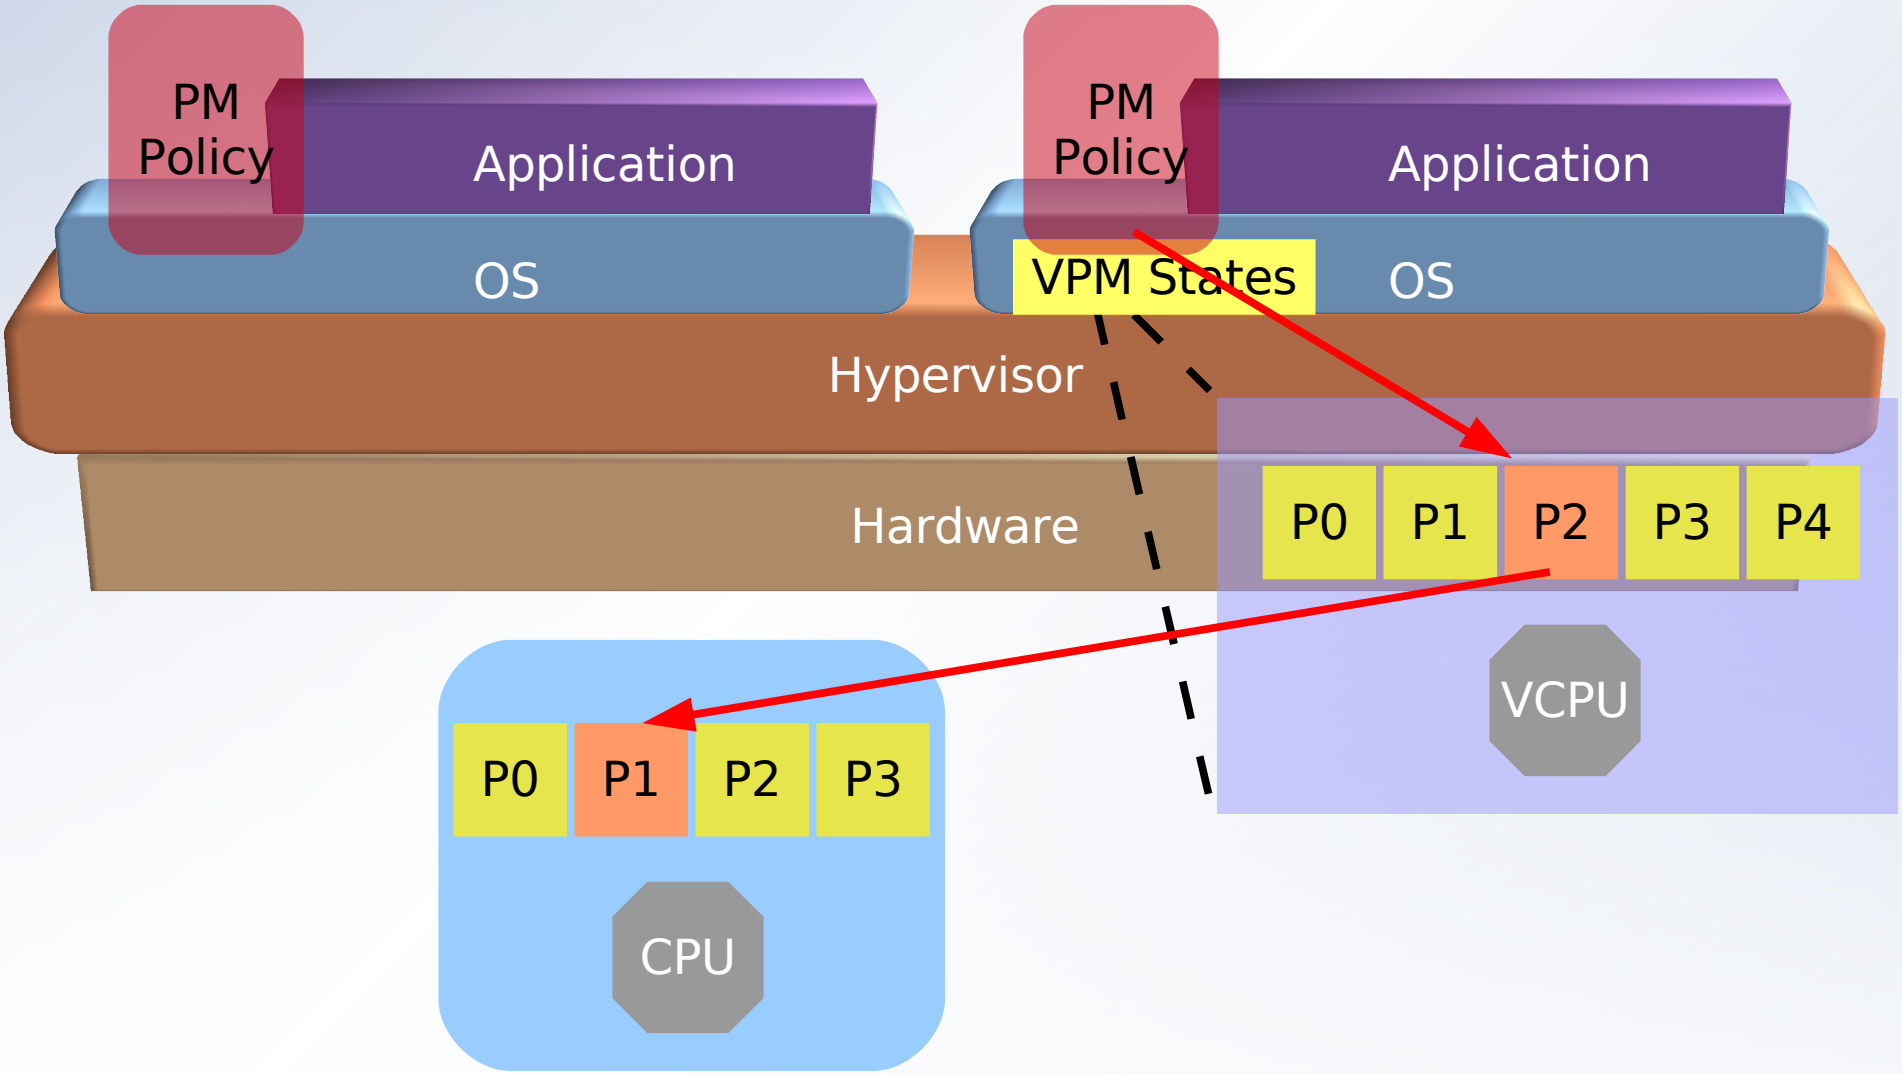

Solution: VPM Channels

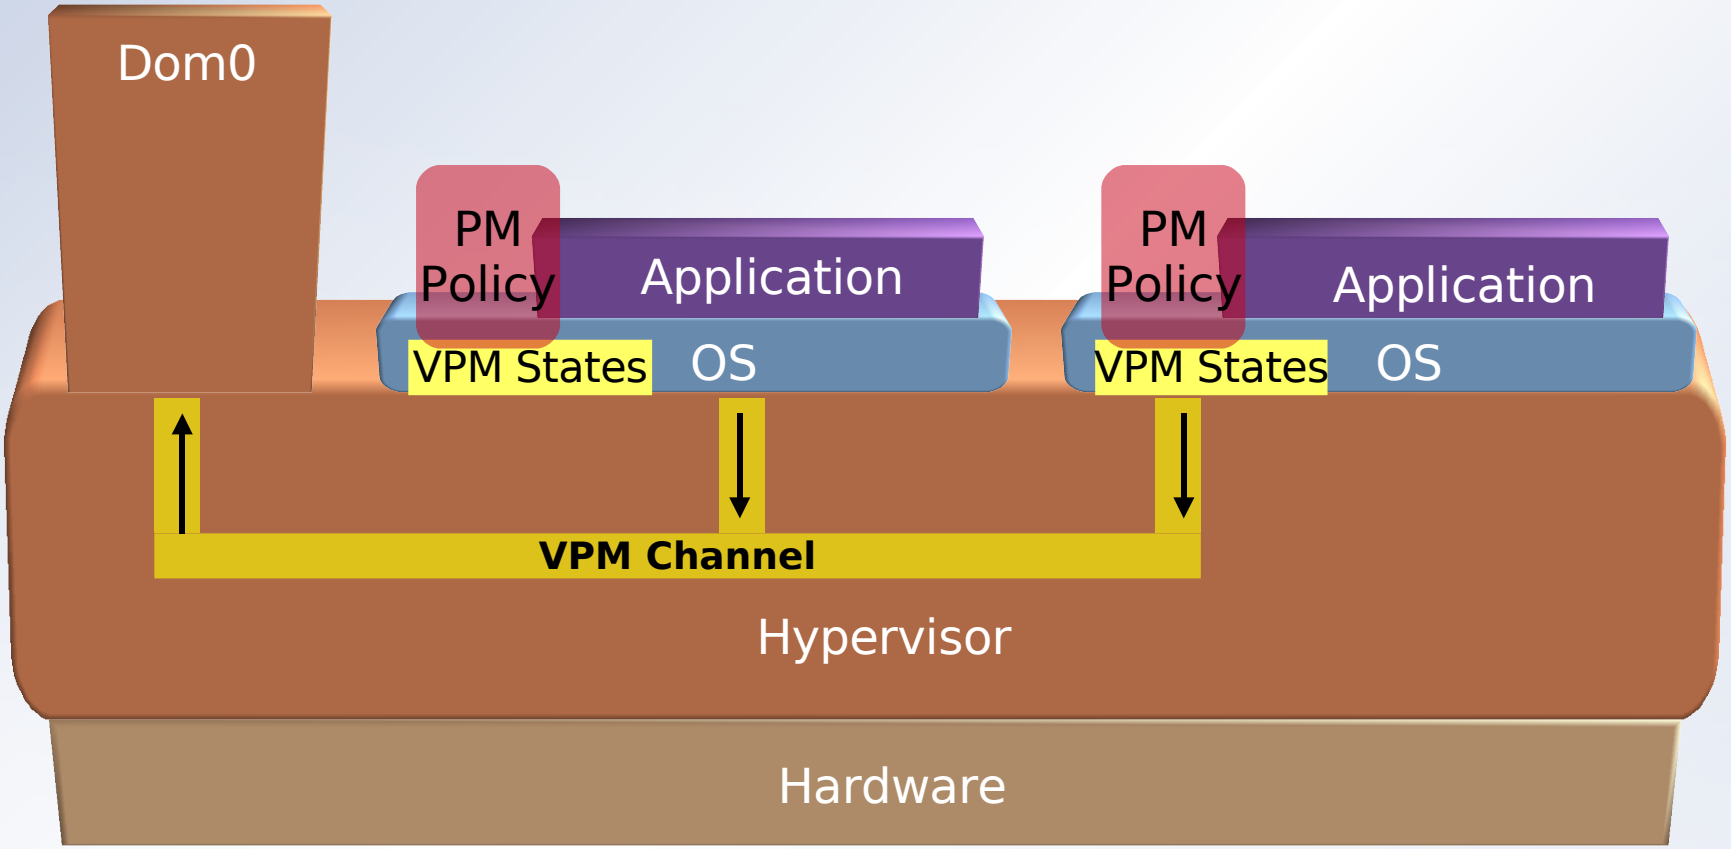

VPM Channel

- Forward VM policy actions to management domain
- Virtualization layer policies manage hardware power states

Problem: Limited Hardware PM Benefits

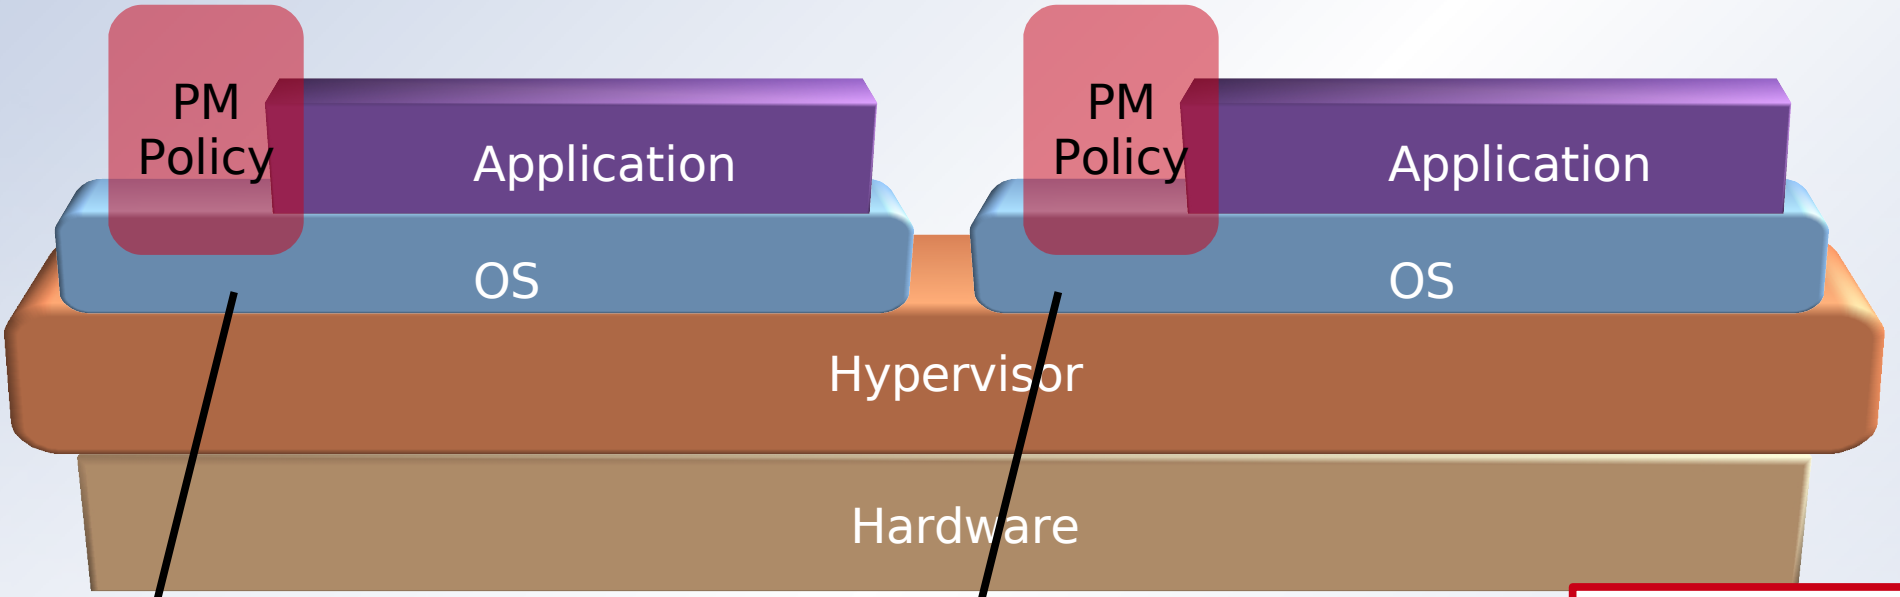

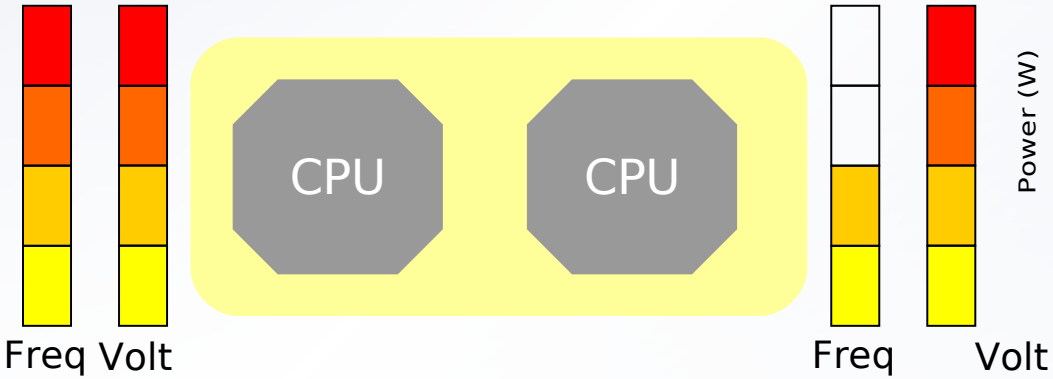

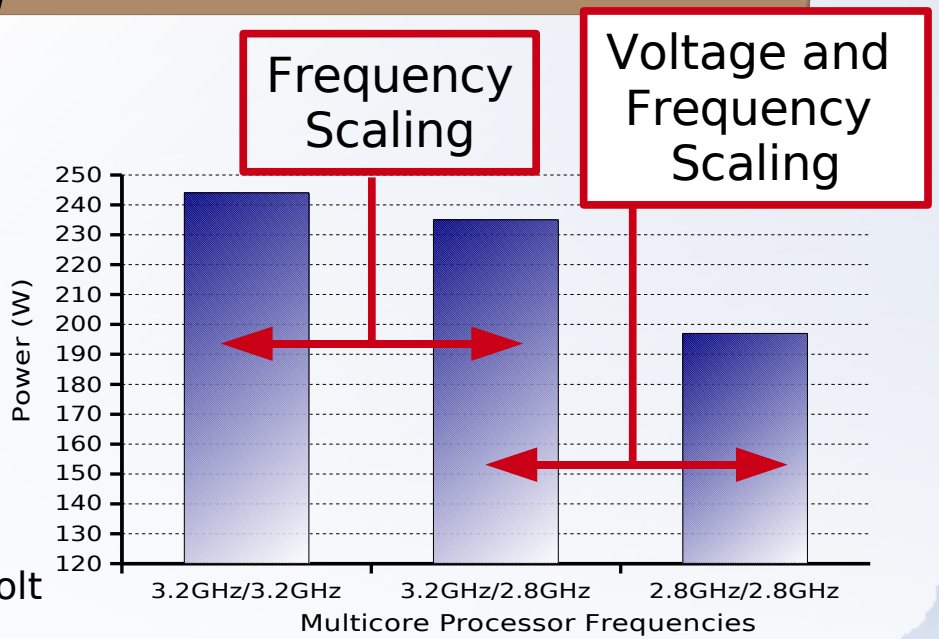

Solution: VPM Mechanisms

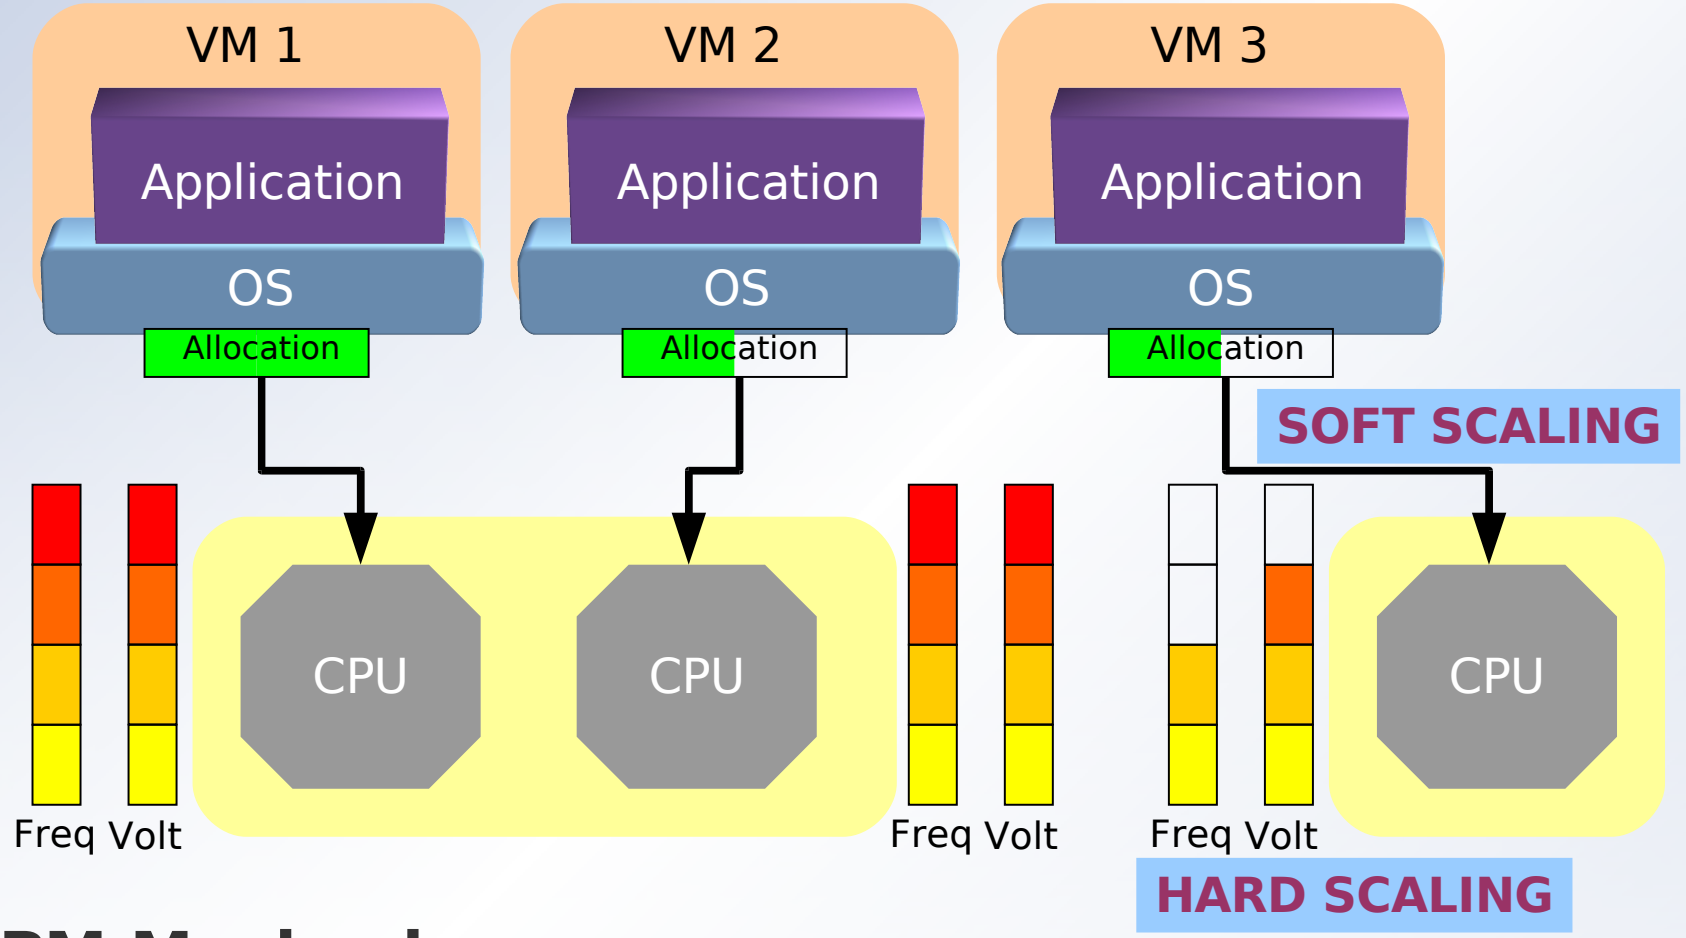

VPM Mechanisms

- Soft scaling restricts resource allocations

Solution: VPM Mechanisms

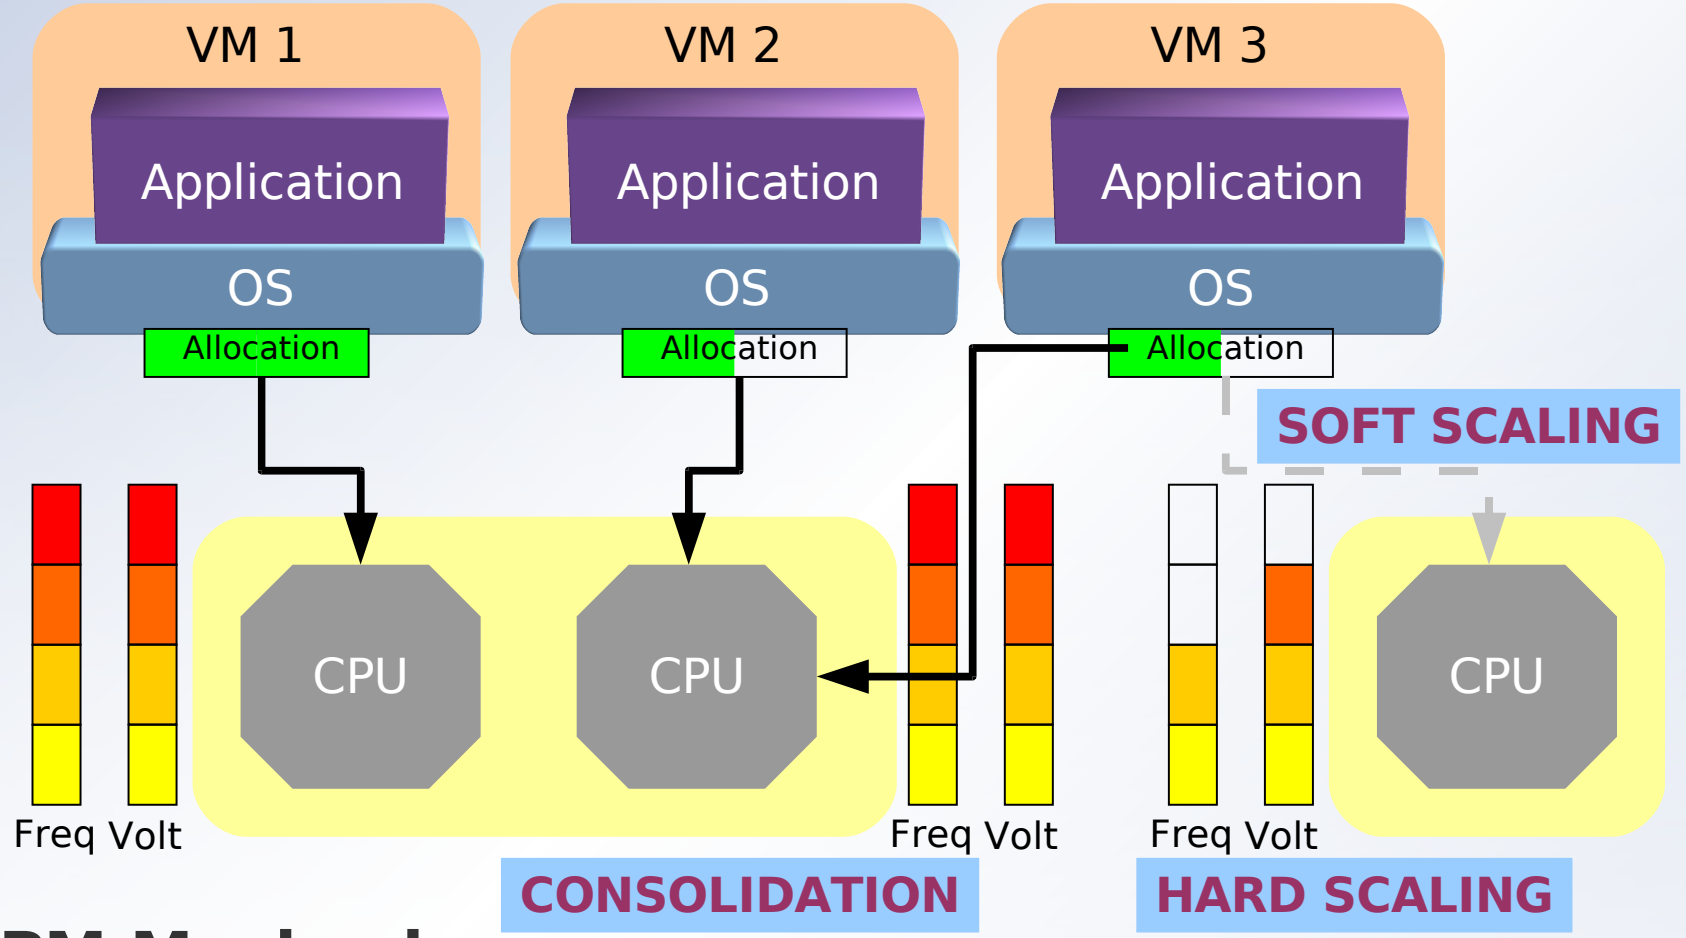

VPM Mechanisms

- Soft scaling restricts resource allocations
- Multiple soft scaled virtual resources can be consolidated

VirtualPower Architecture

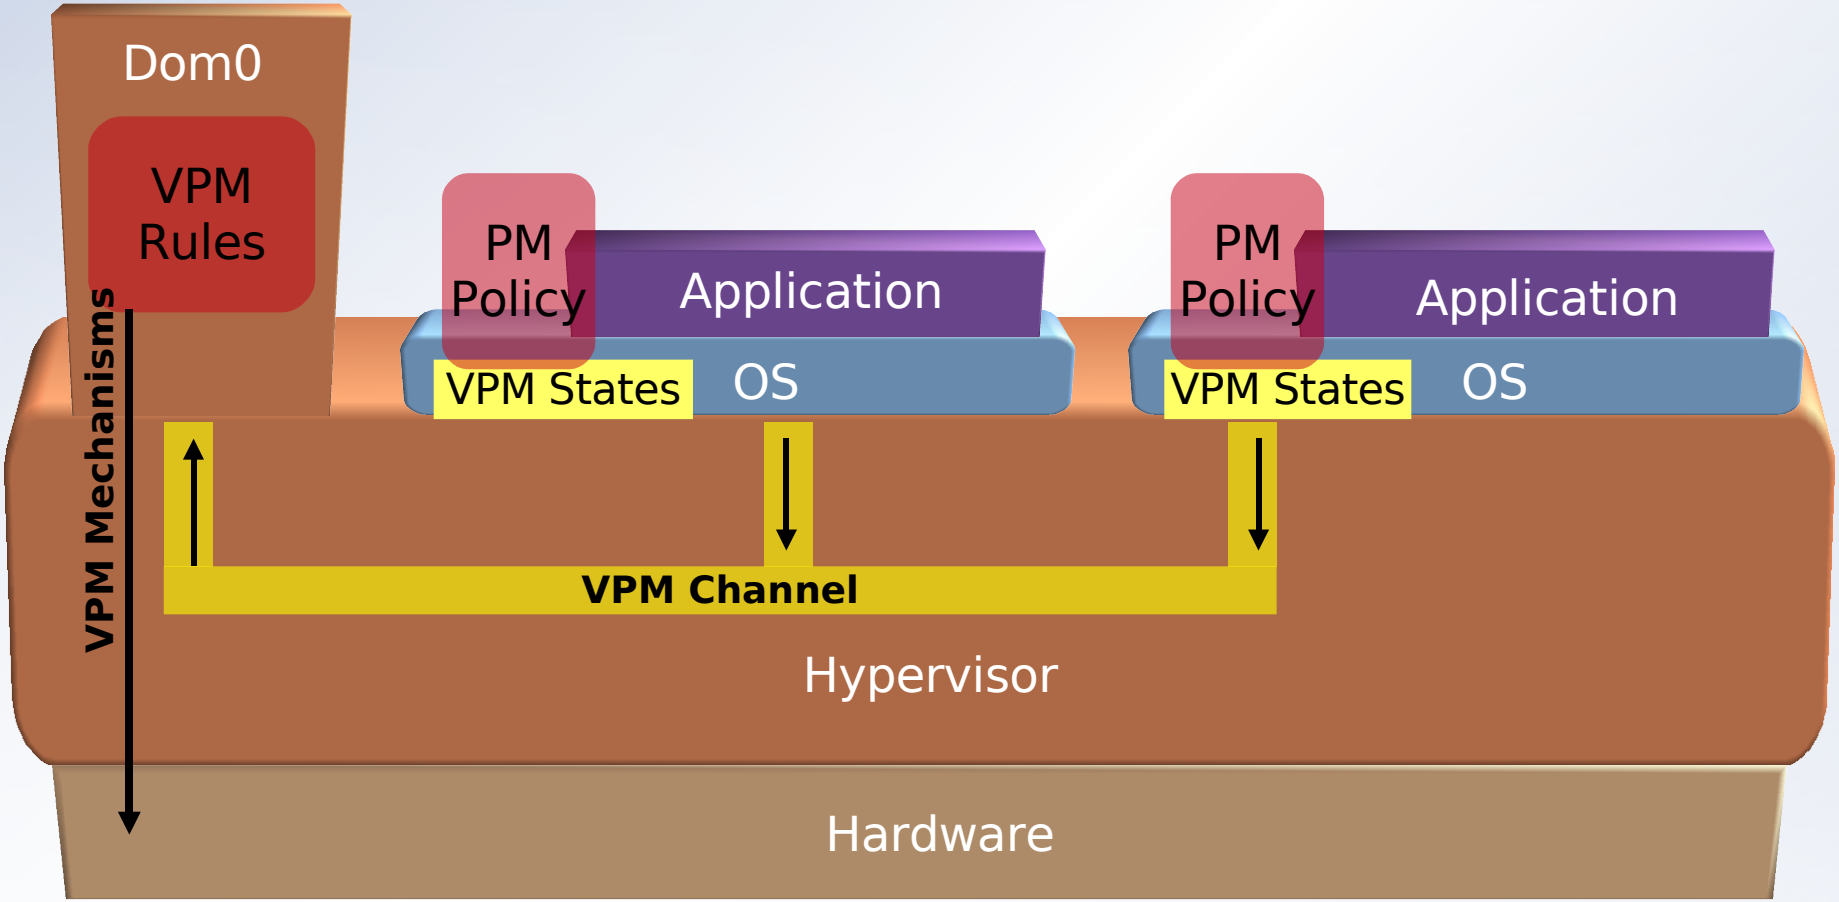

Key Idea: State Based Guidance for VPM Rules

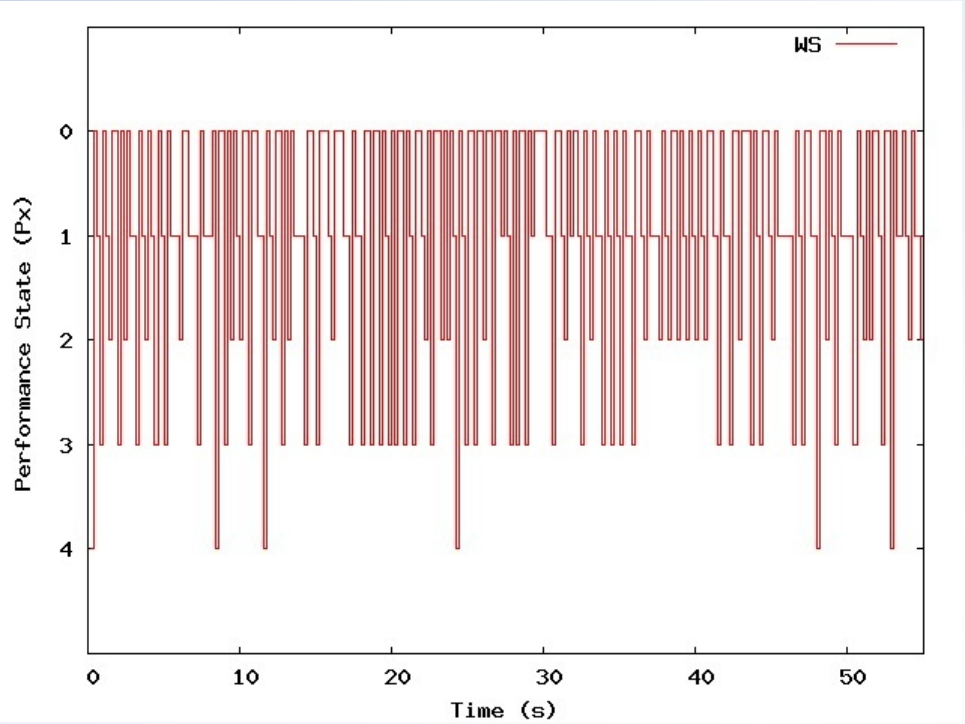

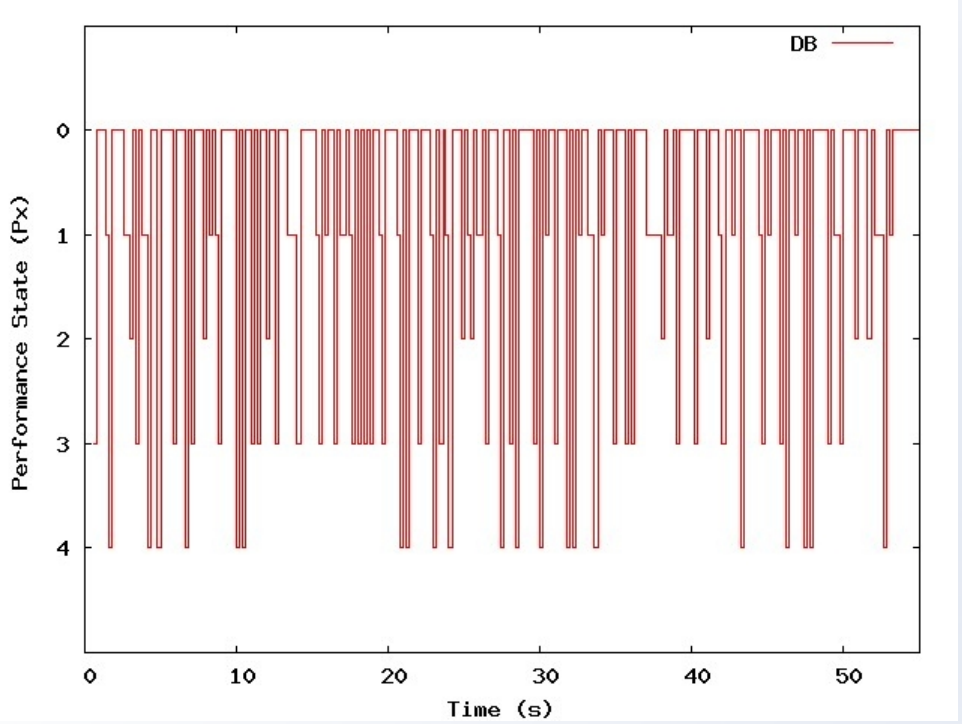

Transparently leverage application specific policies

- VPM state requests from VMs drive virtualization layer policies: Implicit feedback loop
- Requests based upon application specific policies: Feedback allows for SLA compliance under PM

Example: PM-L Rule with State Based Guidance

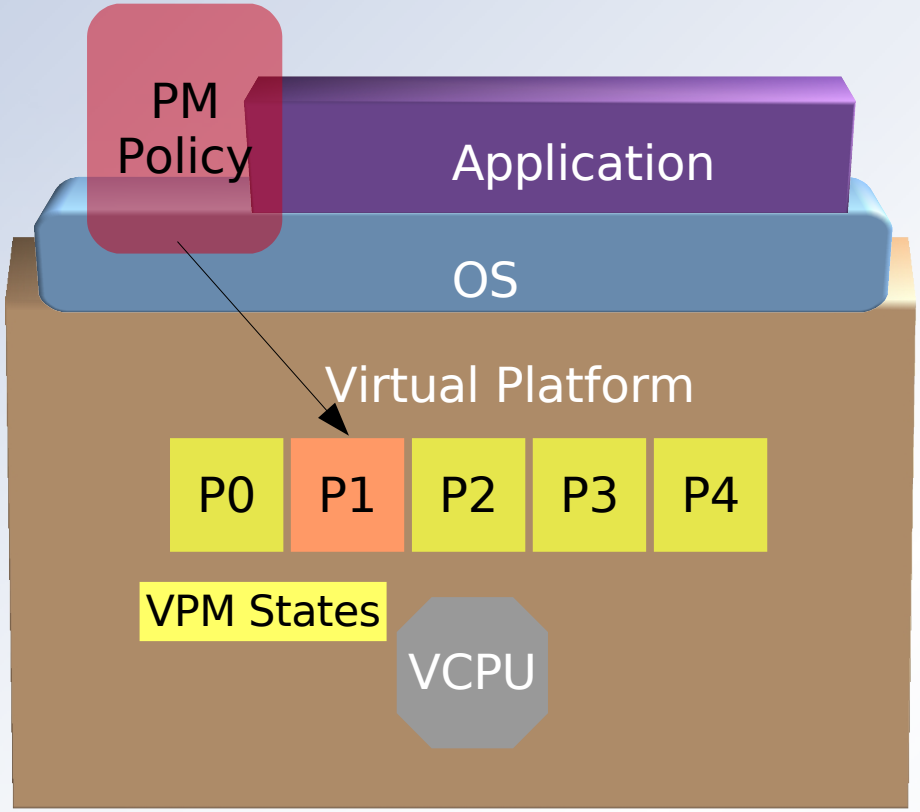

VPM State	CPU Freq	Soft Scaling
3.2GHz	3.2GHz	100%
2.8GHz	2.8GHz	100%
2.0GHz	2.8GHz	71%
1.6GHz	2.8GHz	57%
800MHz	2.8GHz	29%

PM-L Rule

VPM Rules

- Allows for flexibility in datacenter management: different rules for different types of VMs
- Can be simple (e.g. simple translation), or rely upon more complicated analysis for state based guidance

Example: Reacting to VM Policy Actions

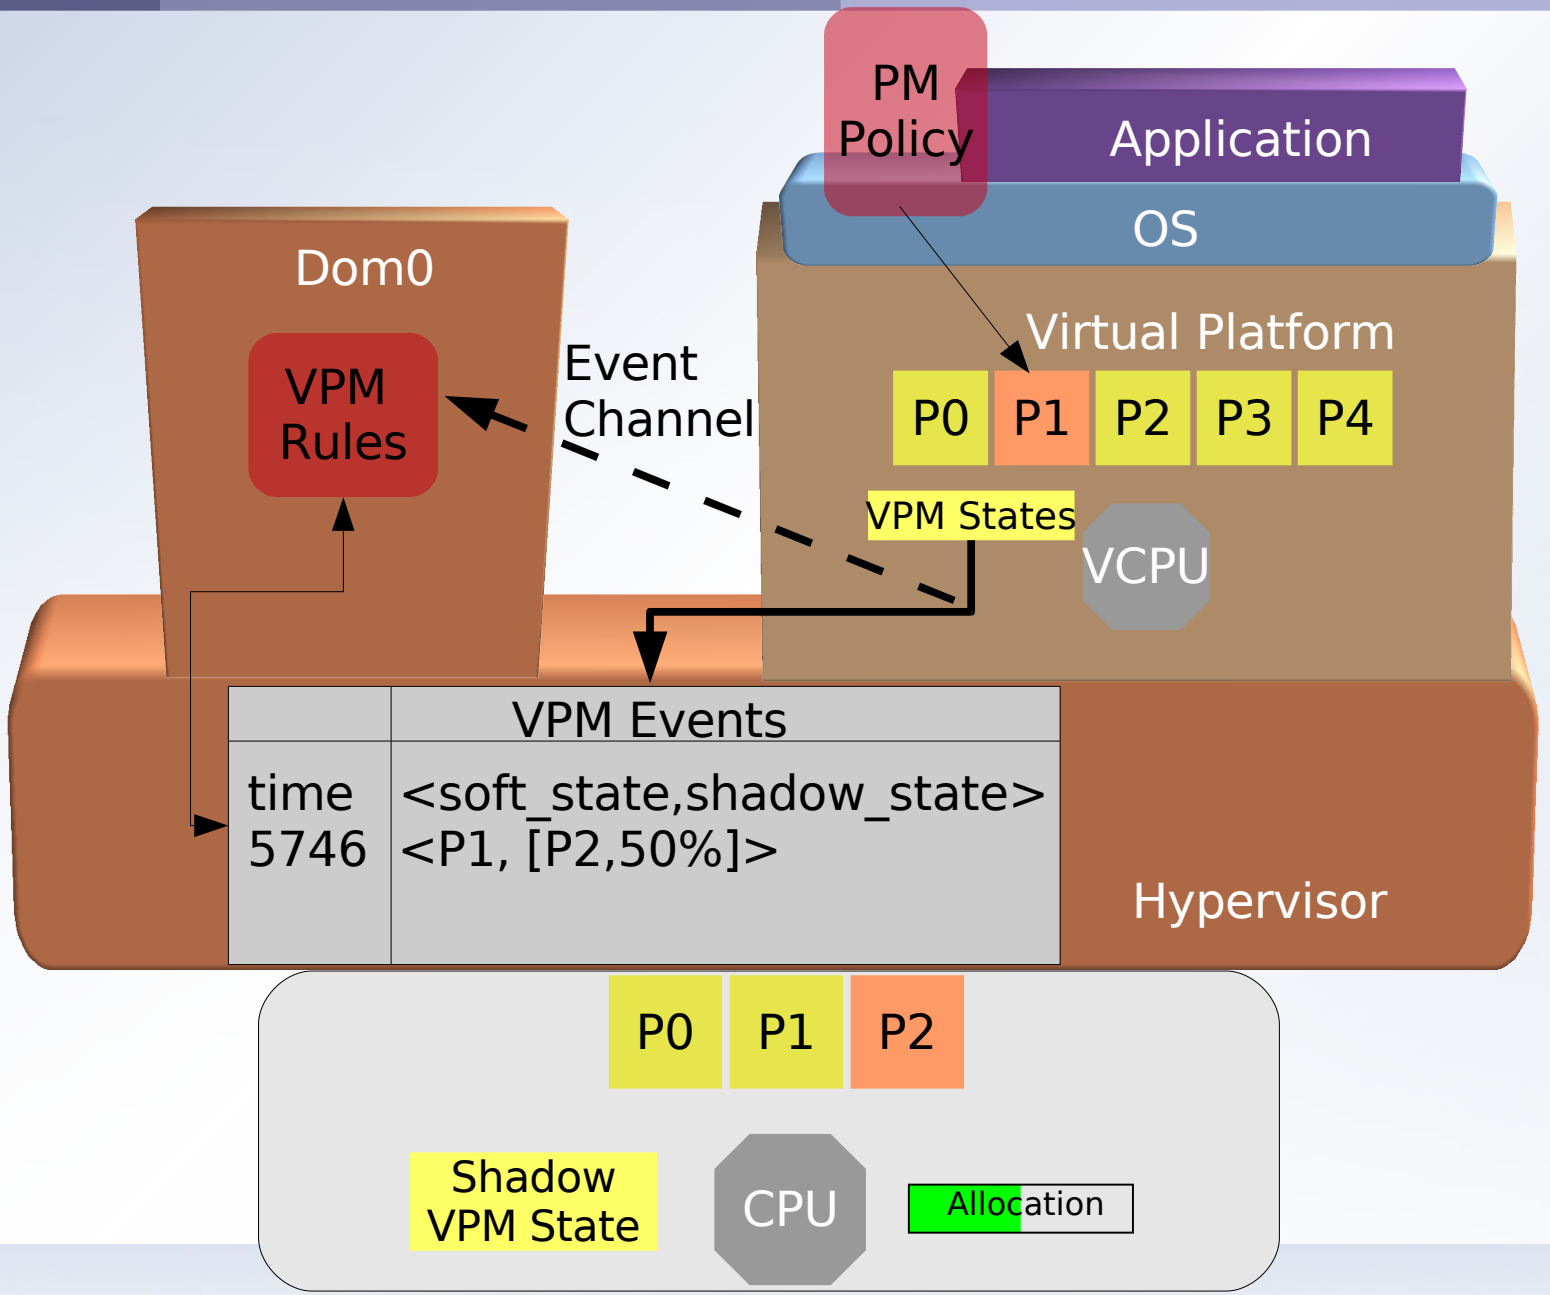

Example: Reacting to VM Policy Actions

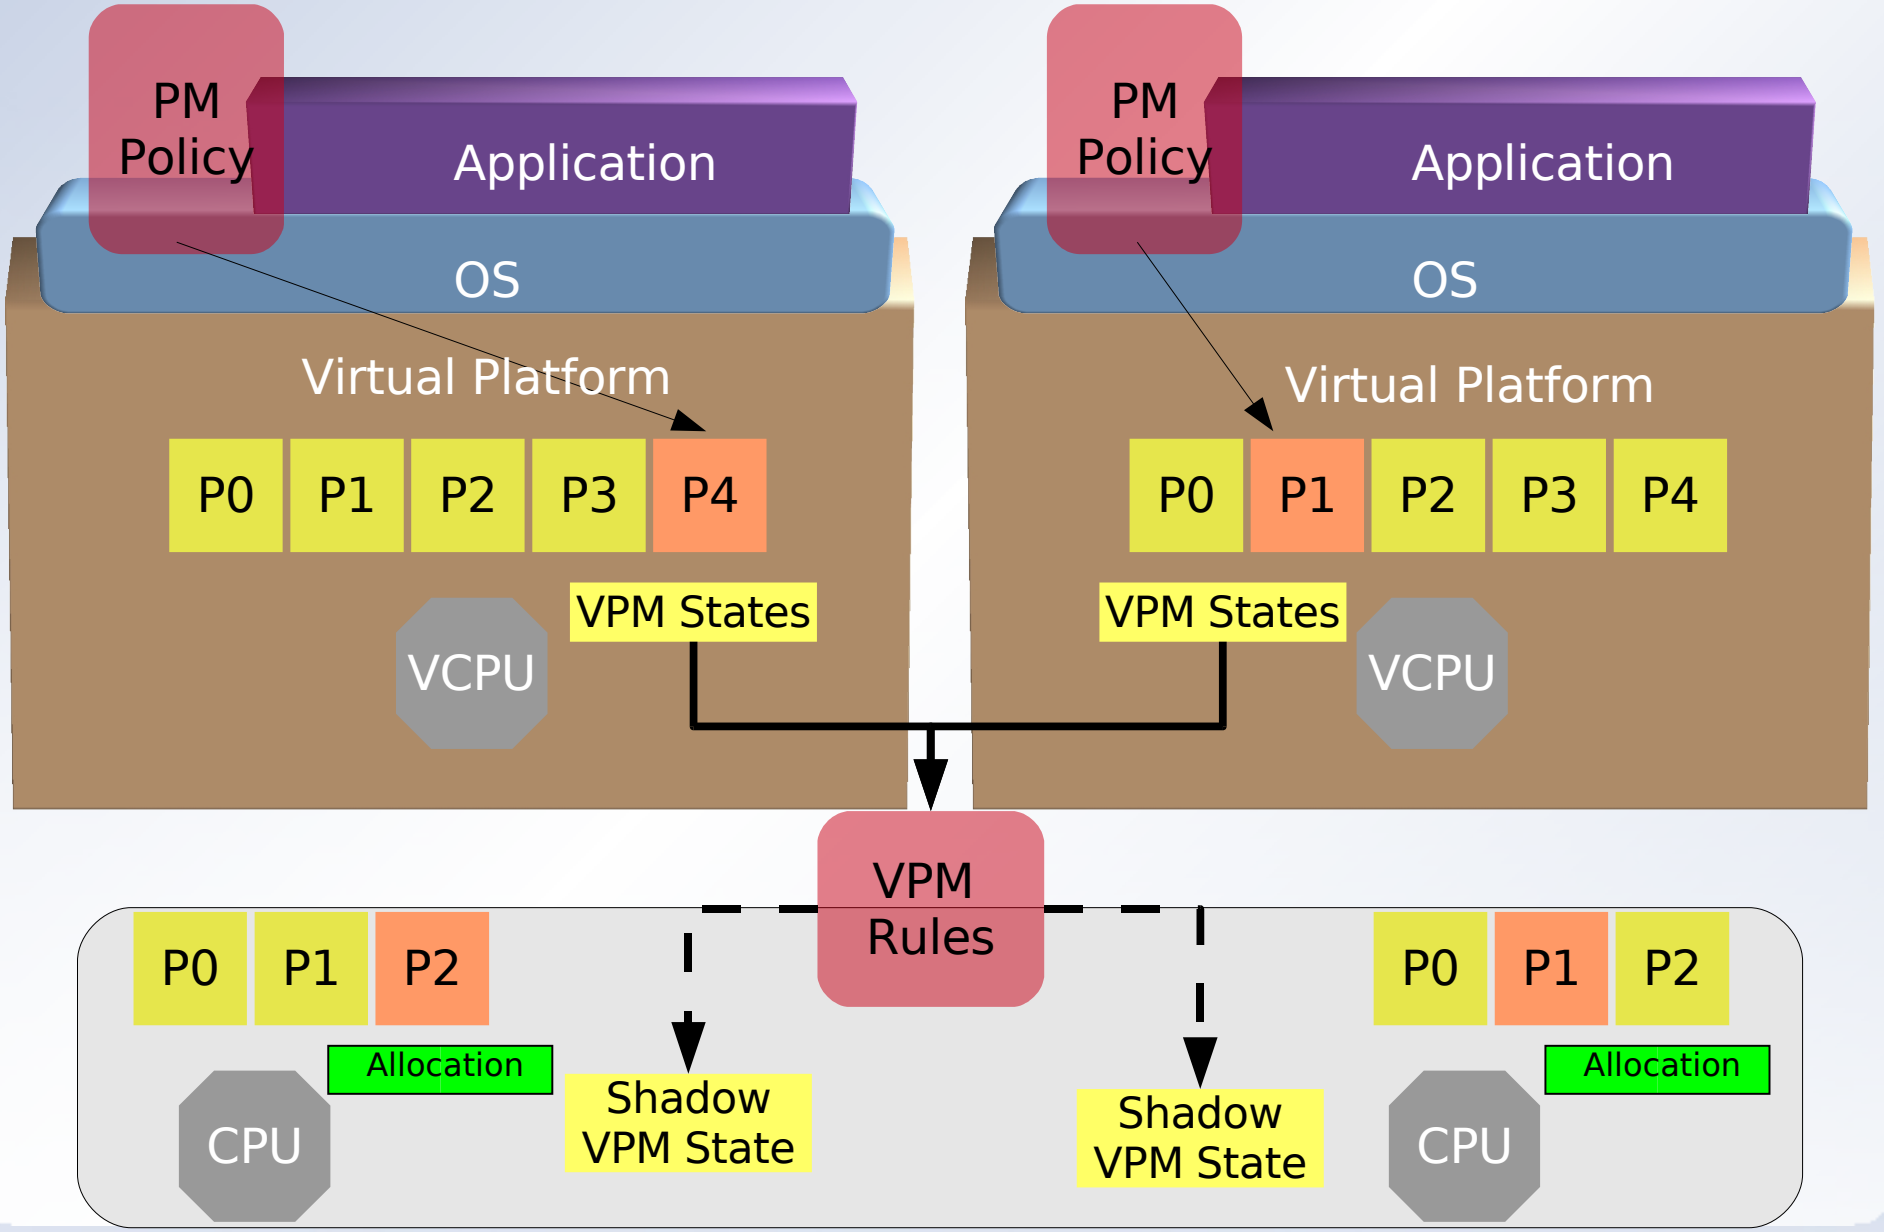

Meeting SLA Constraints with State Based Guidance

Workloads

- Tiered web service (RUBiS)
  - VM policy: Linux ondemand governor
- Transactional workloads
  - VM policy: monitors transaction processing rate and selects state based upon “slack”
- Web service (Nutch) with Quality of Information metric (based upon actual application --Travelport)
  - VM policy: monitors “slack” in QoI and processing time of requests across different client classes

RUBiS: Utilizing Different VPM Rules

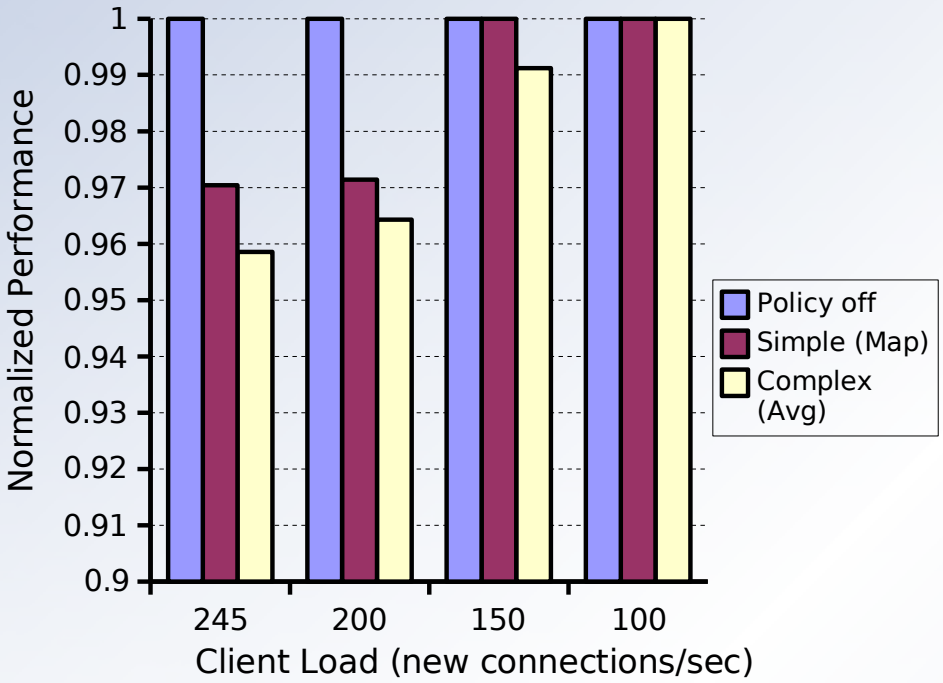

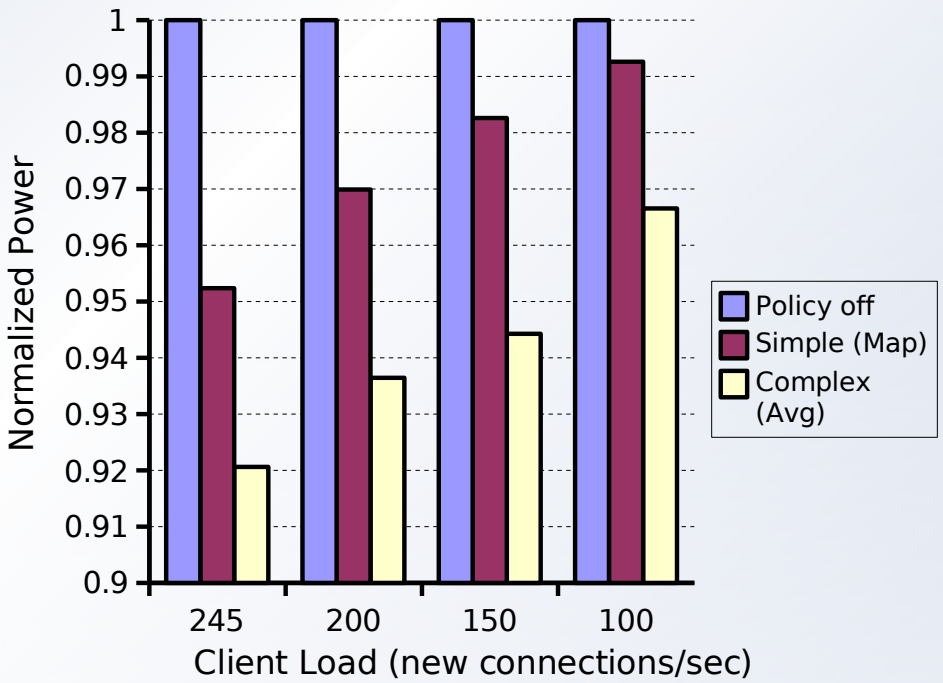

- Necessary to use different VPM rules for different applications
- VPM rules can be sophisticated
  - Adaptive
  - Complex analysis
  - Learning methods

Transactional Workloads: Meeting Varying Demands

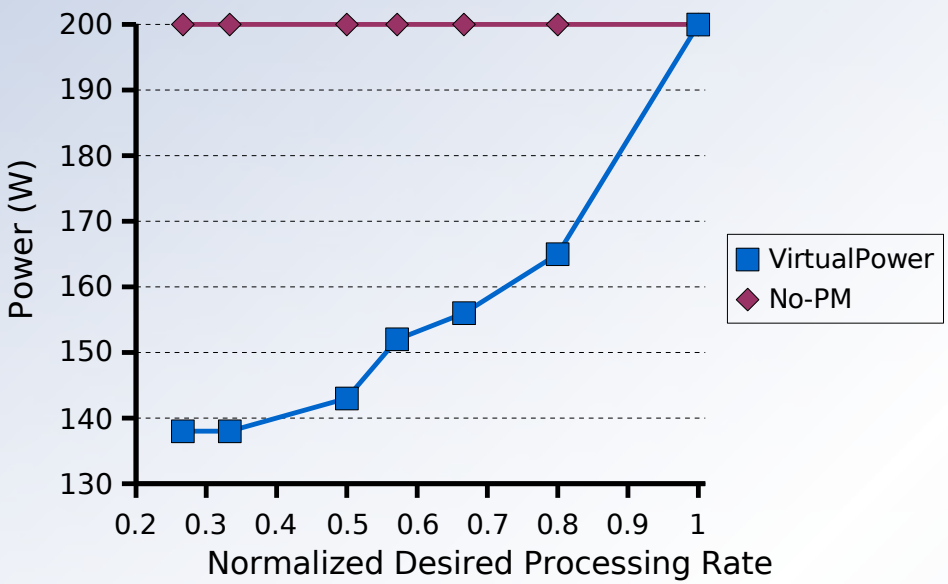

Single VM

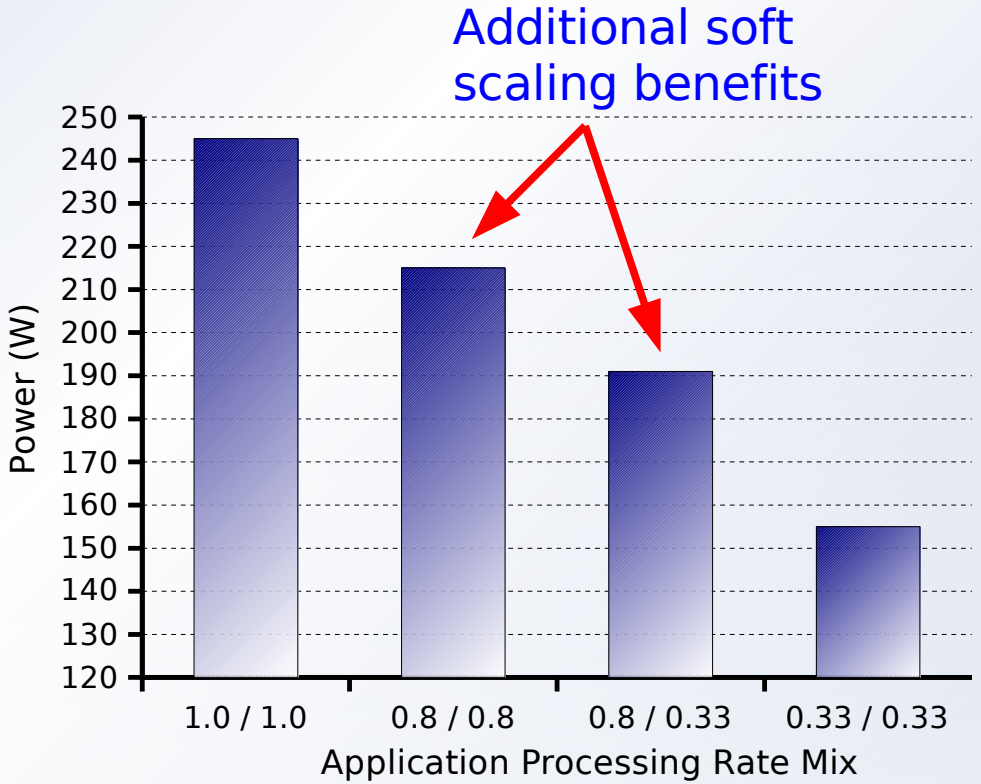

Two VMs

- Single VM: Obvious power benefits for reduced rates
- Multi-VM: VPM rules can obtain substantial savings across VMs with identical or different demands

Nutch: Flexibility in Application Performance Metrics

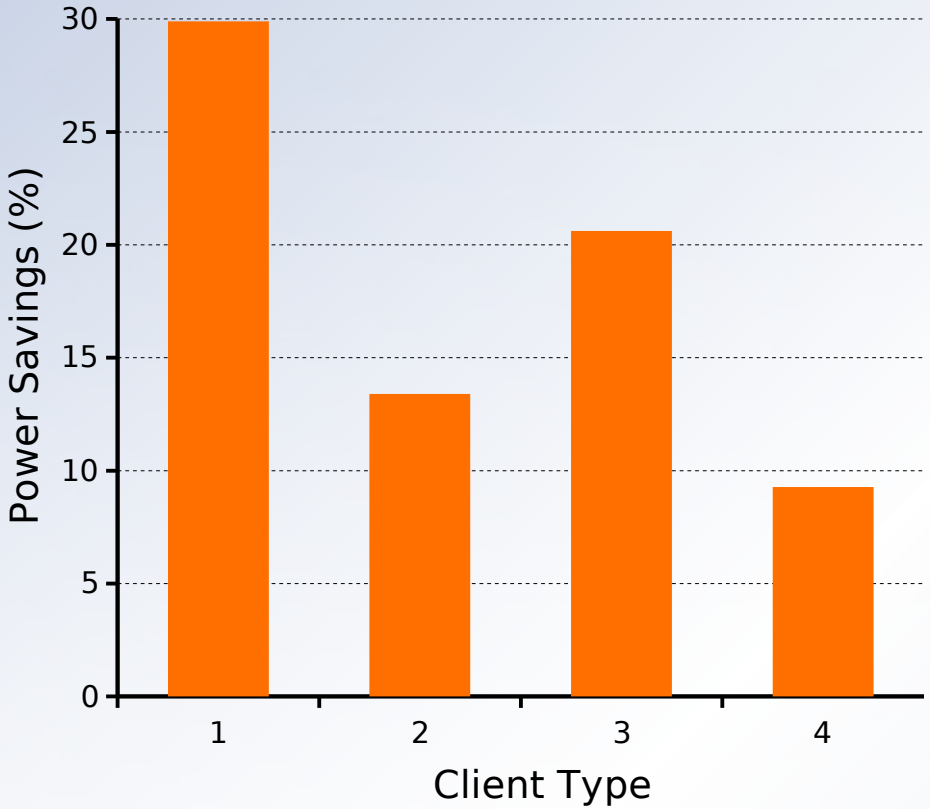

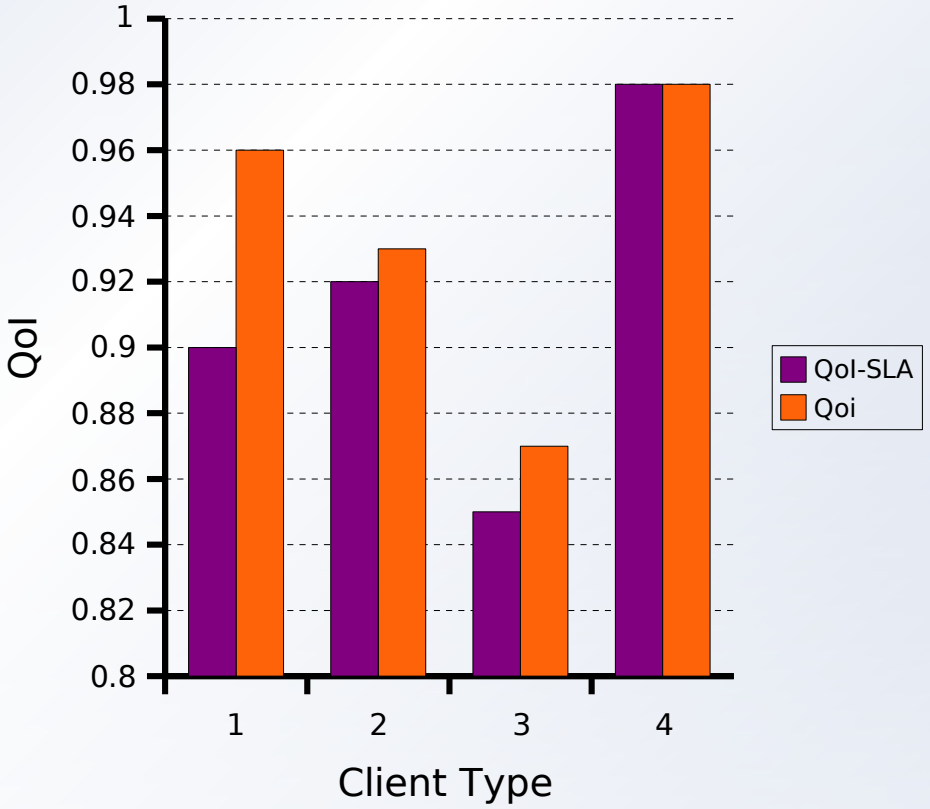

- PM driven by Quality of Information (QoI) metric
  - QoI based on Travelport application
- Use of VM policies for state based guidance
  - SLA compliance across variety of metrics

Consolidation with Heterogeneous Systems (1)

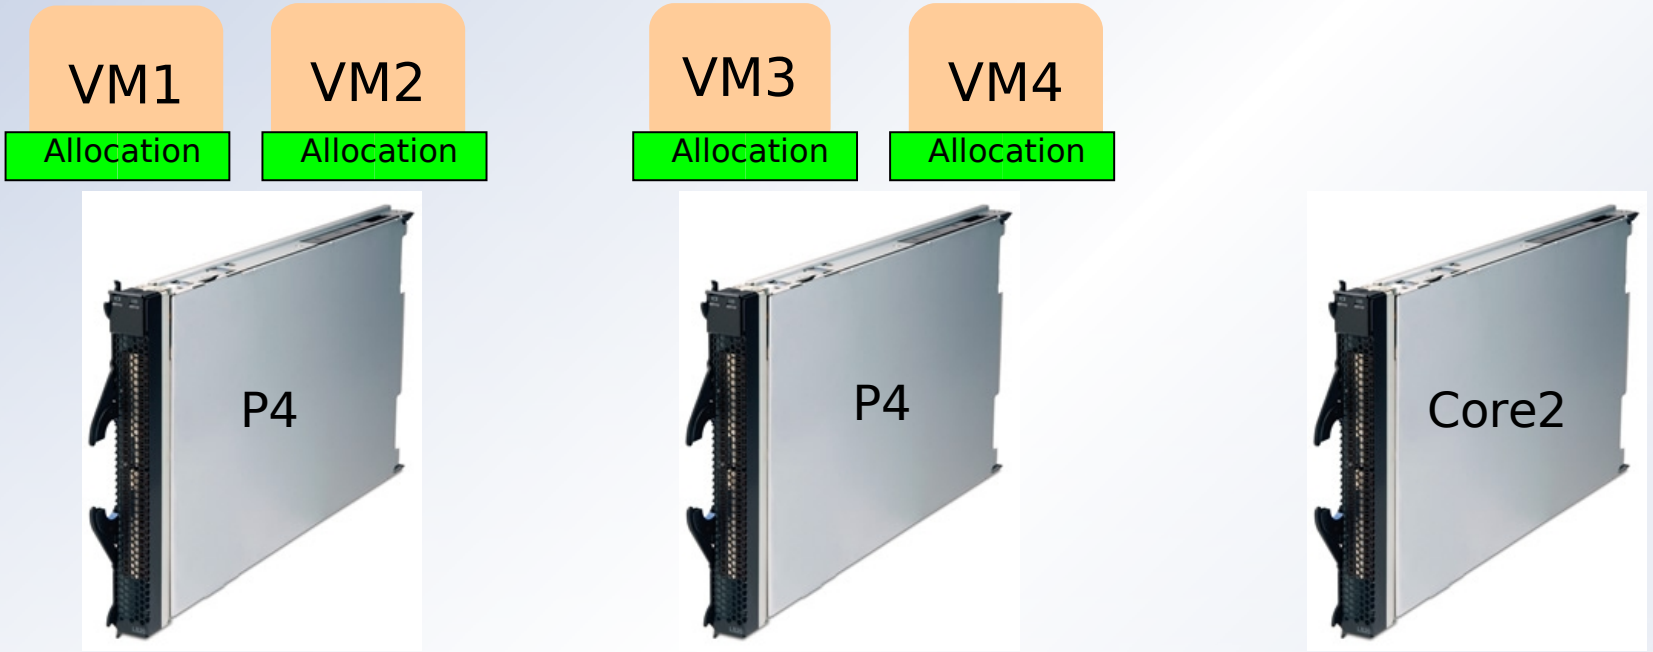

- Three dual core platforms, four deployed VMs
  - Heterogeneous systems
  - Workloads require full performance of P4 core
- PM-G policy heuristic: utilize more power efficient hardware (Core2)

Consolidation with Heterogeneous Systems (2)

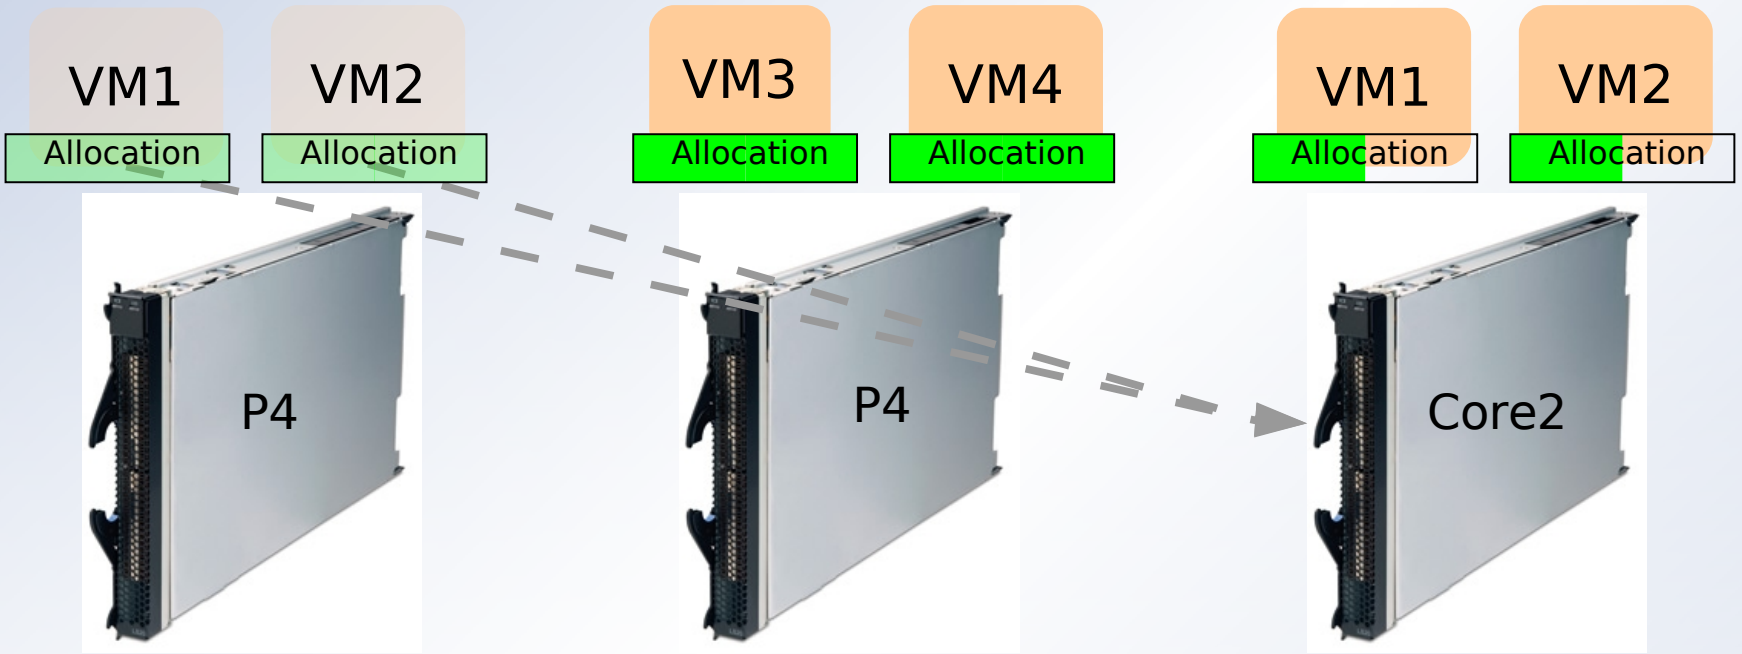

- Migrate two VMs to Core2 system
- Local PM-L policy on Core2 performs soft scaling based upon observed requests
- Soft scaling provides room for further consolidation

Consolidation with Heterogeneous Systems (3)

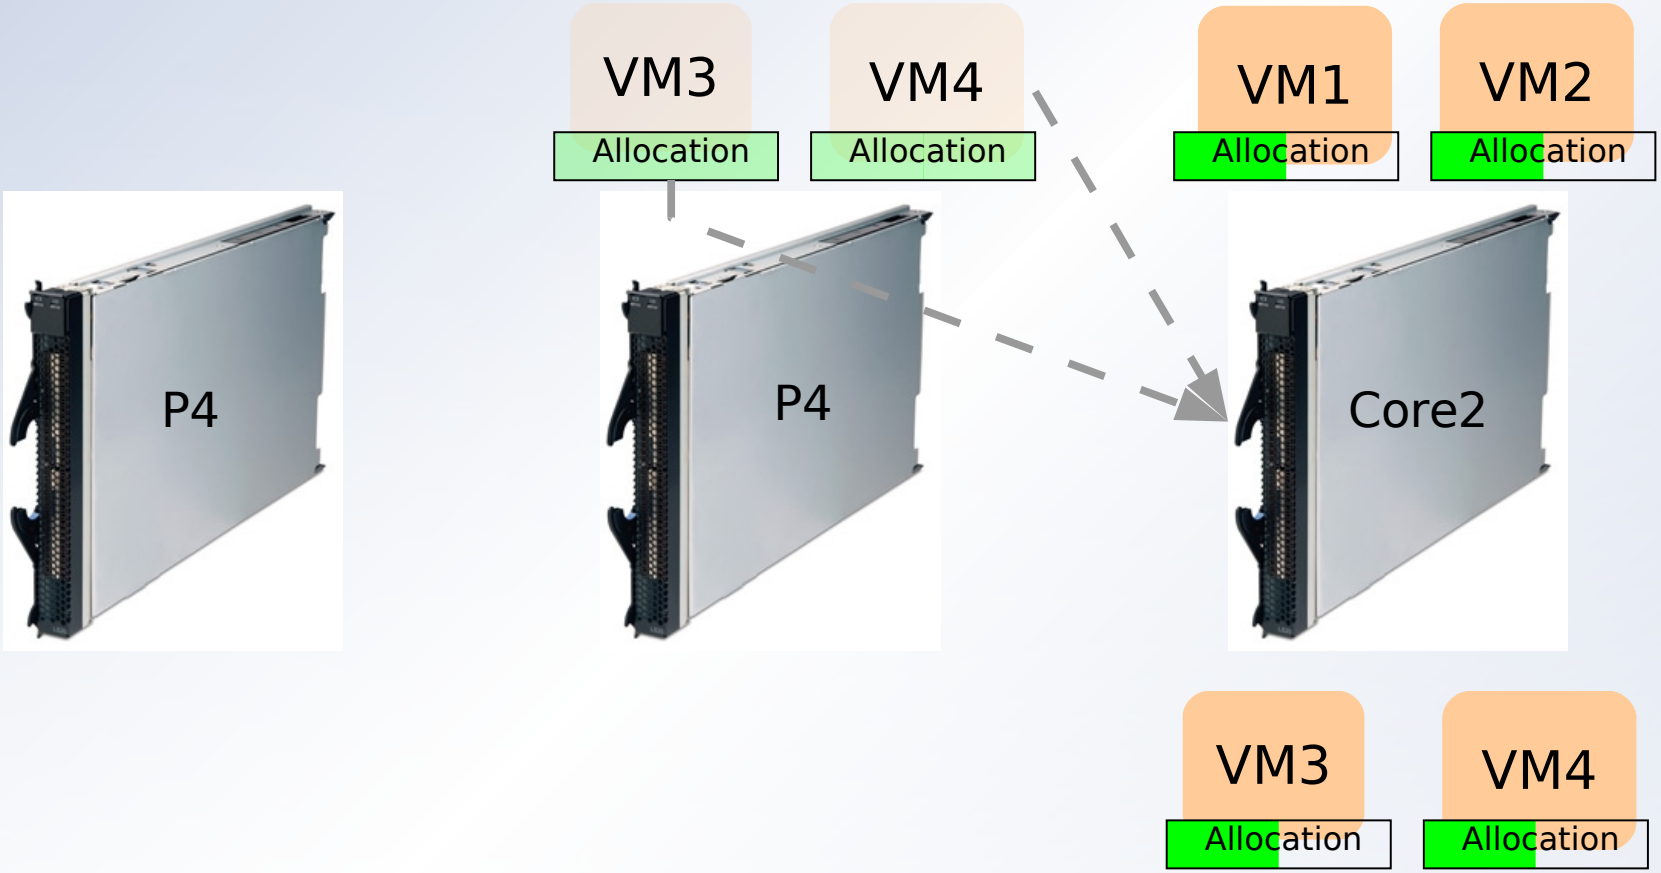

Power Results with Heterogeneous Consolidation

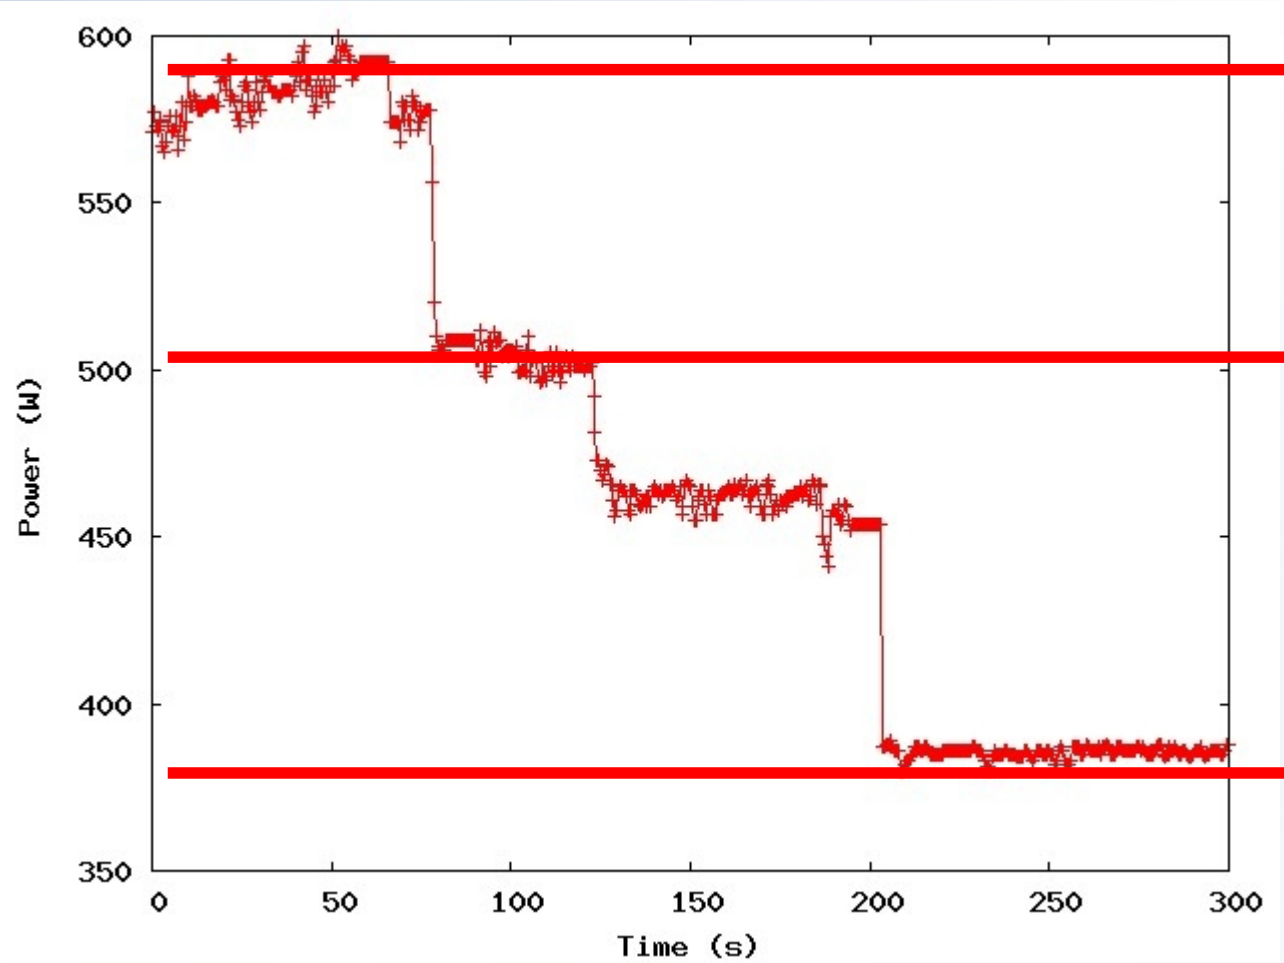

Heterogeneity awareness

Soft scale enabled consolidation

Concluding Remarks / Future Work

Power management in virtualized systems

- Transparently leverage existing application policies
- Deal with heterogeneity in hardware/manageability
- Maintain isolation and independence
- Obtain power savings with VM resource sharing

Solutions/contributions

- Virtualized “soft” PM states
- VPM channels and mechanisms

Future Work

- Distributed power throttling: VPM tokens
- Idle power management: Additional VPM C-states
- Efficient soft-scale consolidations: Hardware extensions

Acknowledgments / Questions

Eugene Gorbatov
Rob Knauerhase

Dilma Da Silva
Freeman Rawson

Partha Ranganathan
Vanish Talwar

Sandip Agarwala
Sanjay Kumar
Himanshu Raj

? ? ? ?