INF-VSP1683

VMware vSphere Cluster Resource Pools Best Practices

Frank Denneman, VMware, Inc.

Rawlinson Rivera, VMware, Inc.

Disclaimer

- This session may contain product features that are currently under development.
- This session/overview of the new technology represents no commitment from VMware to deliver these features in any generally available product.
- Features are subject to change, and must not be included in contracts, purchase orders, or sales agreements of any kind.
- Technical feasibility and market demand will affect final delivery.
- Pricing and packaging for any new technologies or features discussed or presented have not been determined.

Talk Outline

- Resource Distribution
- Resource Controls
- Resource Pool Cookbook

Introduction

- **Frank Denneman**
  - Senior Architect at Technical Marketing VMware
  - VCDX #29
  - frankdenneman.com
  - Co-Author of vSphere 4.1, 5.0 and 5.1 clustering deepdive books
- **Rawlinson Rivera**
  - Senior Consultant at VMware Professional Services
  - VCDX #86
  - punchingclouds.com
  - Co-Author of Mastering VMware Virtual Infrastructure 3
  - Author of MCITP Microsoft Exchange 2007 Messaging Design and Deployment Study Guide

DRS Cluster

- **Aggregates ESXi host capacity into one large pool**
- **Functions as the root resource pool**
- **Provide an abstraction layer between resource providers and resource consumers**
  - ESXi Host: Resource provider
  - Virtual machine: Resource consumer
  - Resource pool: Both

Resource Consumers and Providers

- Resource pools: allocation and isolation of groups of VMs
- Consume resources from cluster
- Provide resources to their consumers

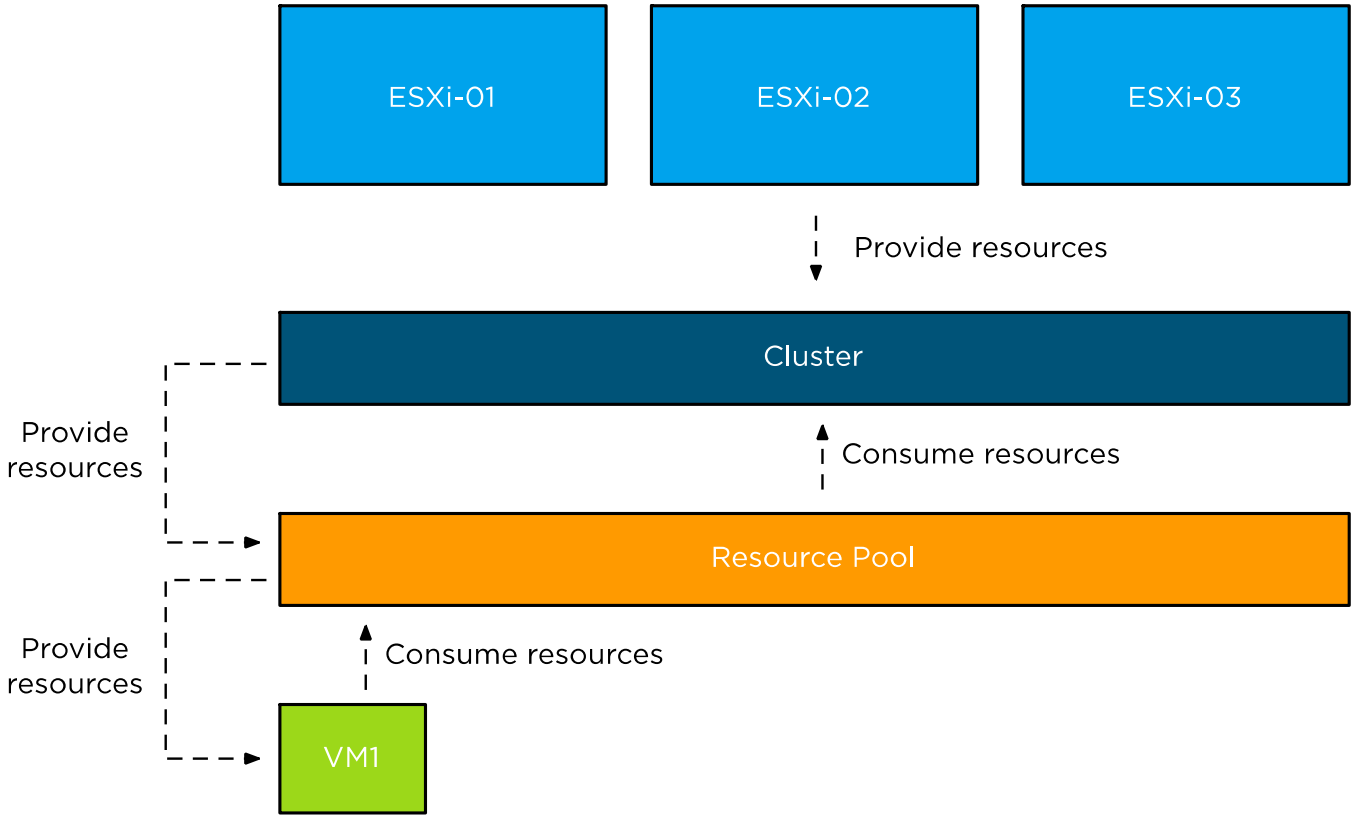

Resource Distribution

- Resource pools are always first stop in resource distribution

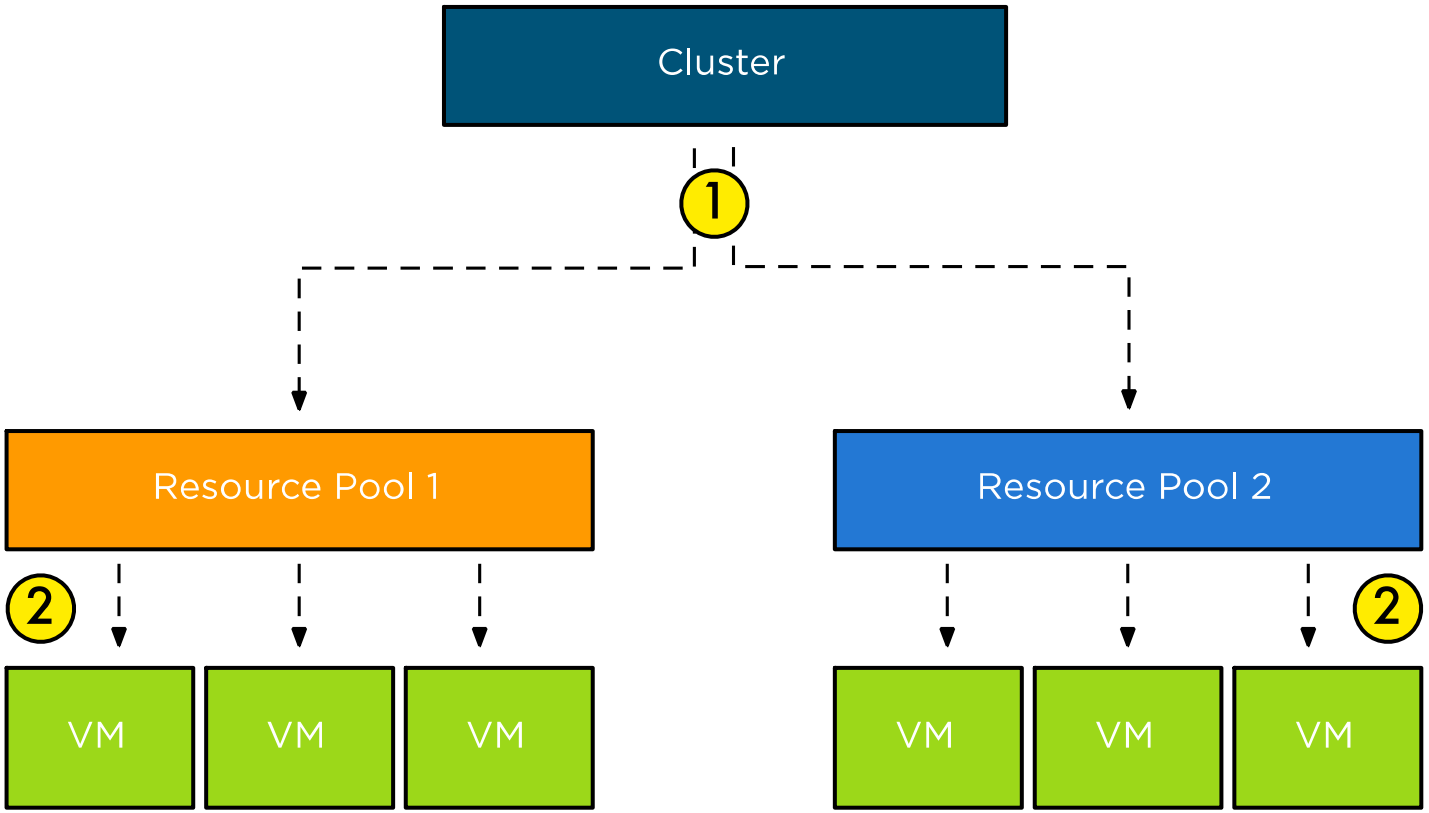

Sibling Rivalry

- **This means that RP1 and RP2 are the first to compete for resources**

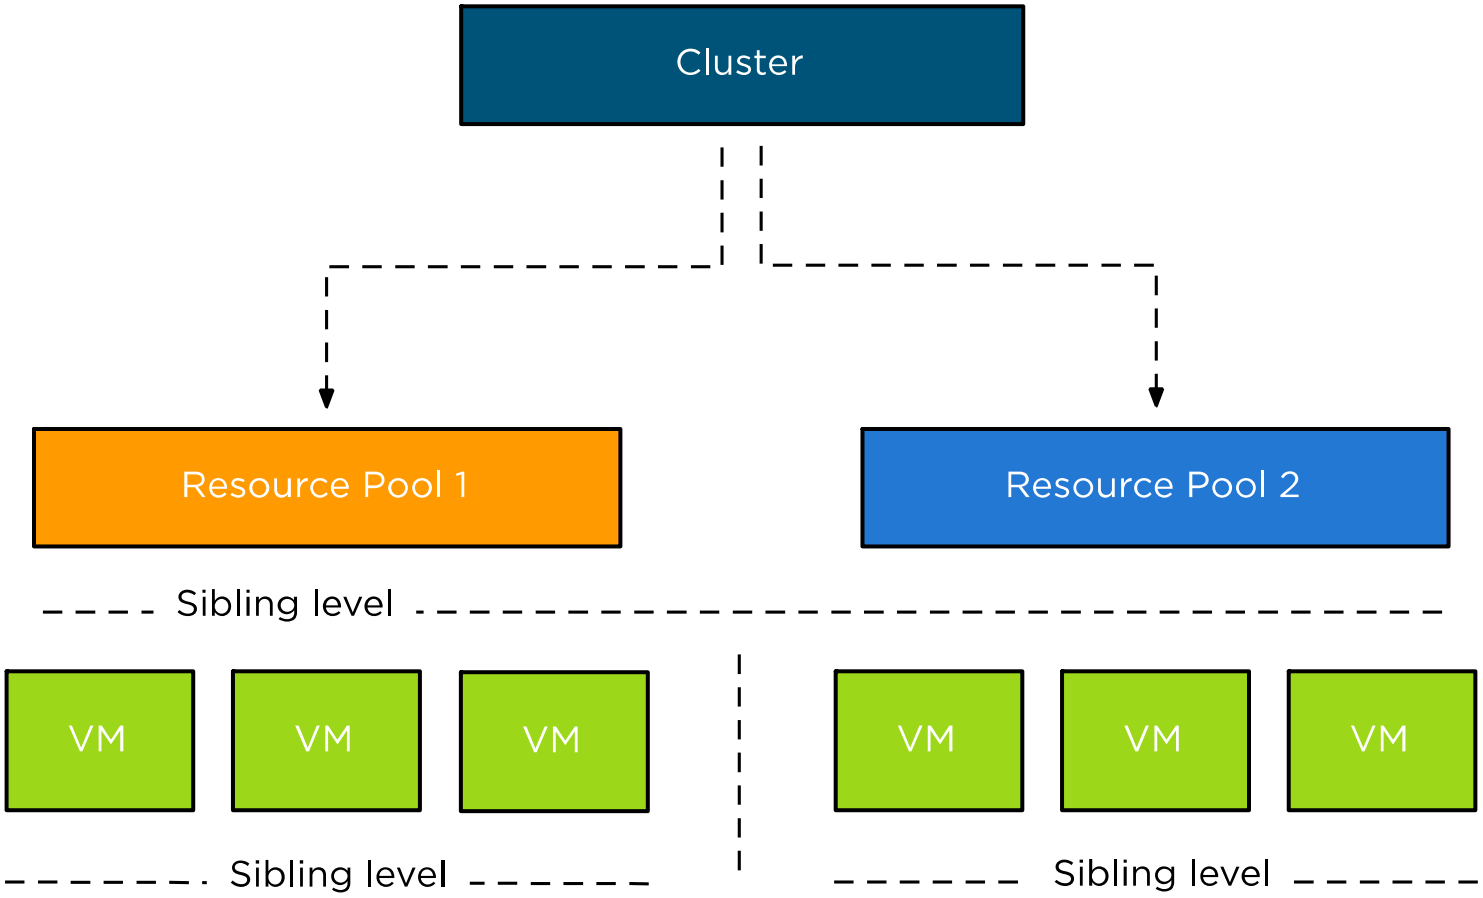

Resource Controls

Why Need Resource Controls?

- What if: Sum of VM Demand > Cluster capacity
- Need a way to decide who gets what.
- This is where resource controls come in
- Goal: Enable high consolidation and over-commitment safely!

Resource Controls

- Reservation: Guaranteed physical allocation (MIN)
- Shares: Allocation in between
- Limit: Guaranteed upper bound (MAX)
- Available on VM level and Resource Pool level
- Expandable reservation

Shares

- **Specify the relative importance of object on sibling level**

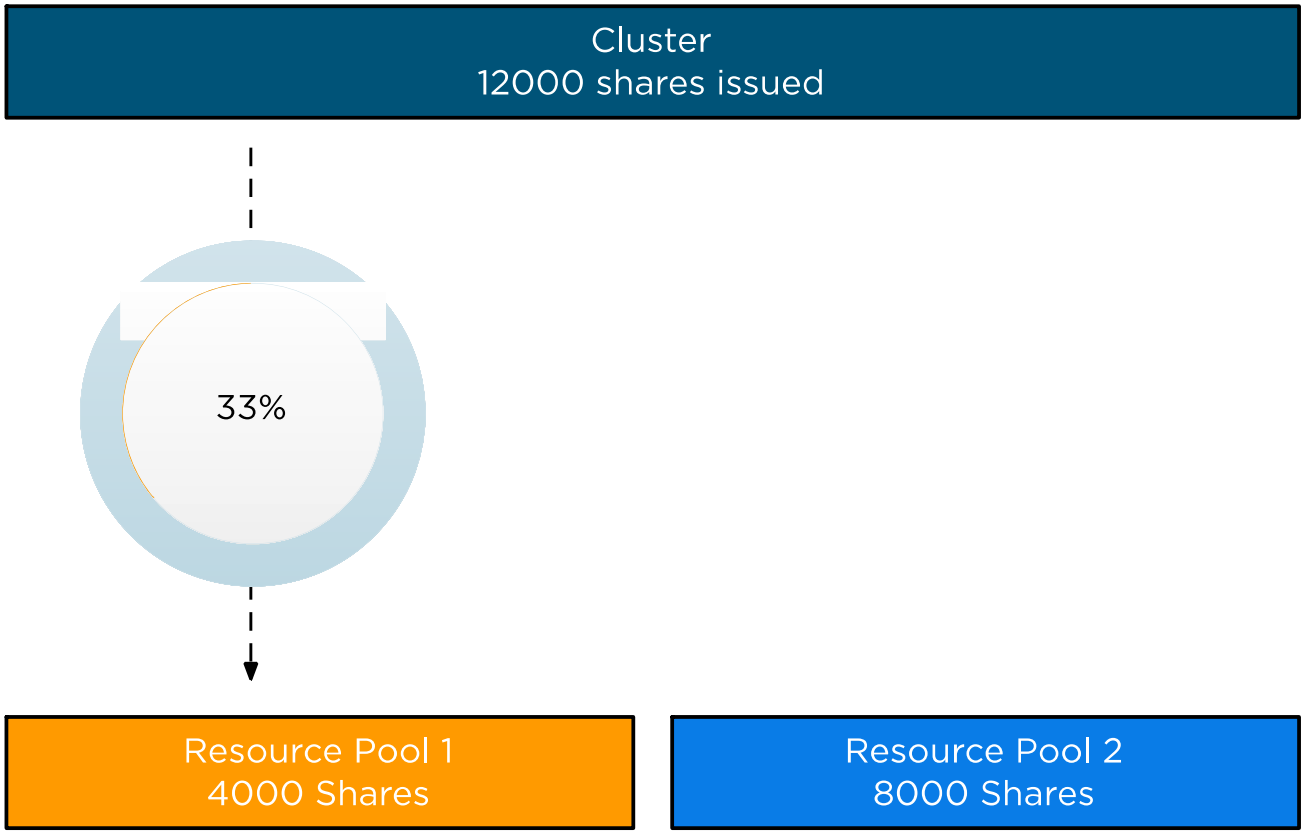

Shares

- **Four share levels**
  - Low (1)
  - Normal (2)
  - High (4)
  - Custom
- **Normal share level on Resource Pool equals:**
  - 4000 shares of CPU (1000 shares per vCPU)
  - 163,840 shares or memory (10 shares per MB)
- **Monster VM (64 vCPU, 1 TB memory High share level)**
  - 128,000 shares of CPU
  - 2,048,000,000 shares

Common Cluster Tree

- **Common cluster-tree equals:**

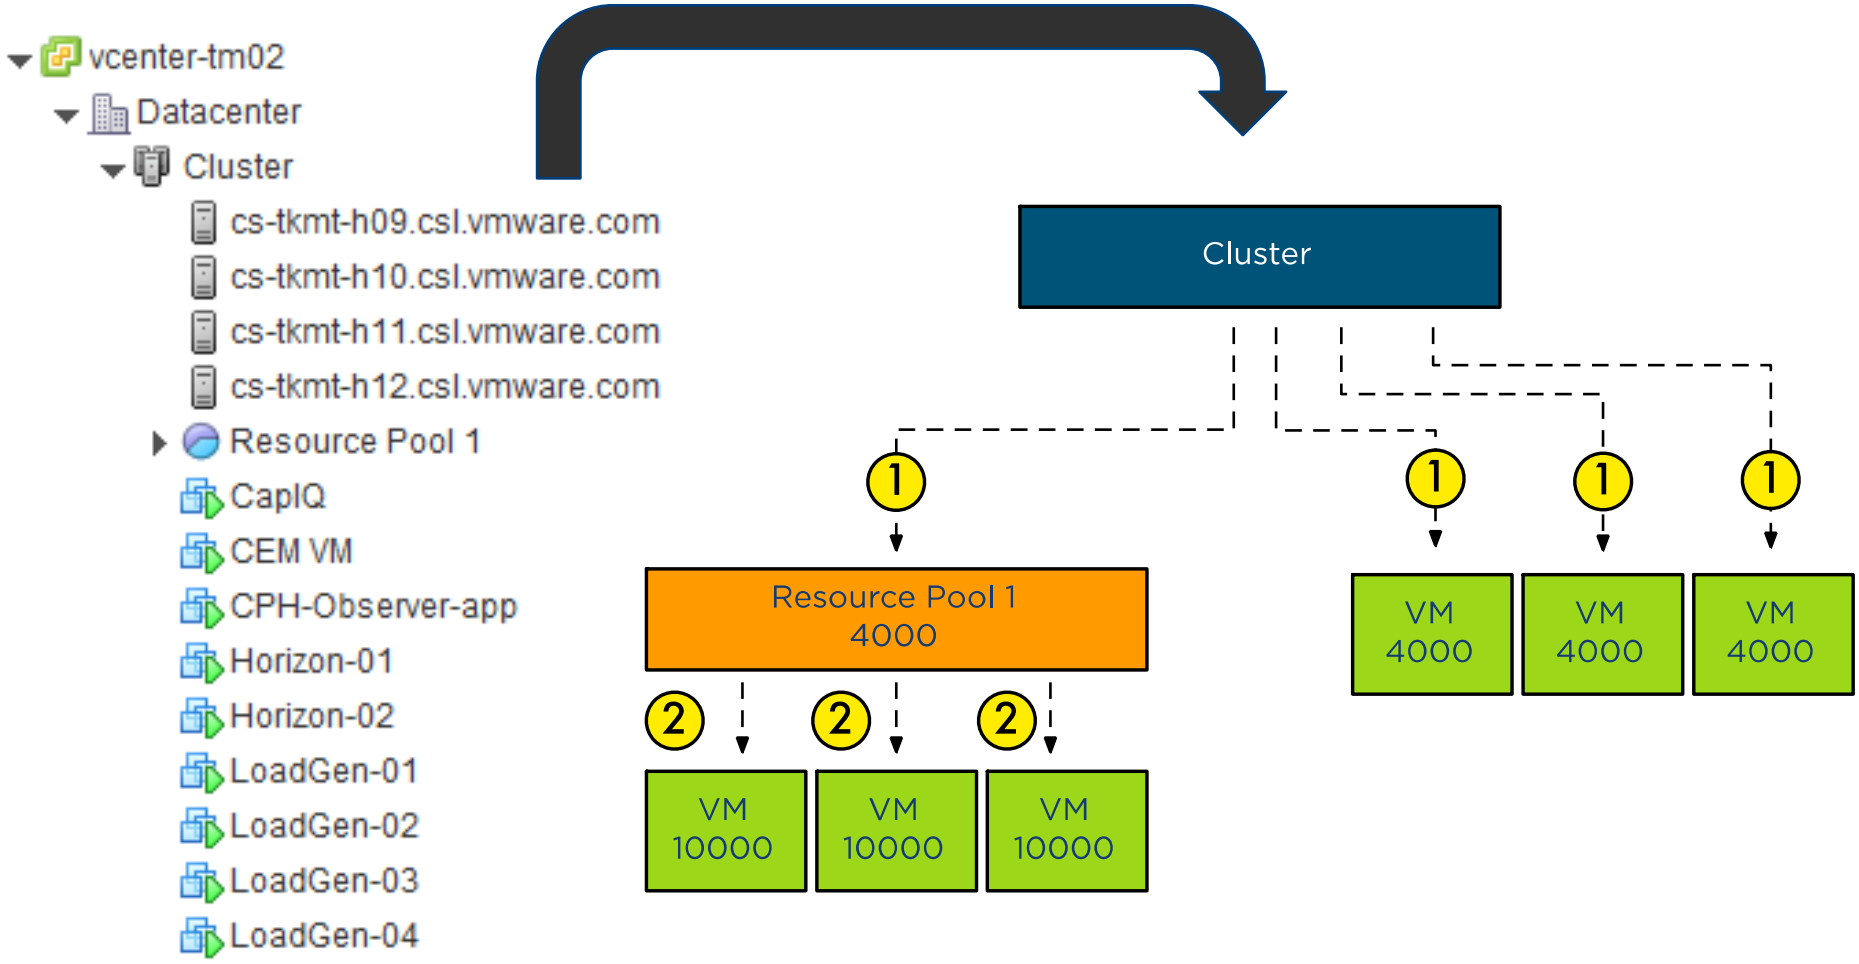

vmware

Resource Distribution

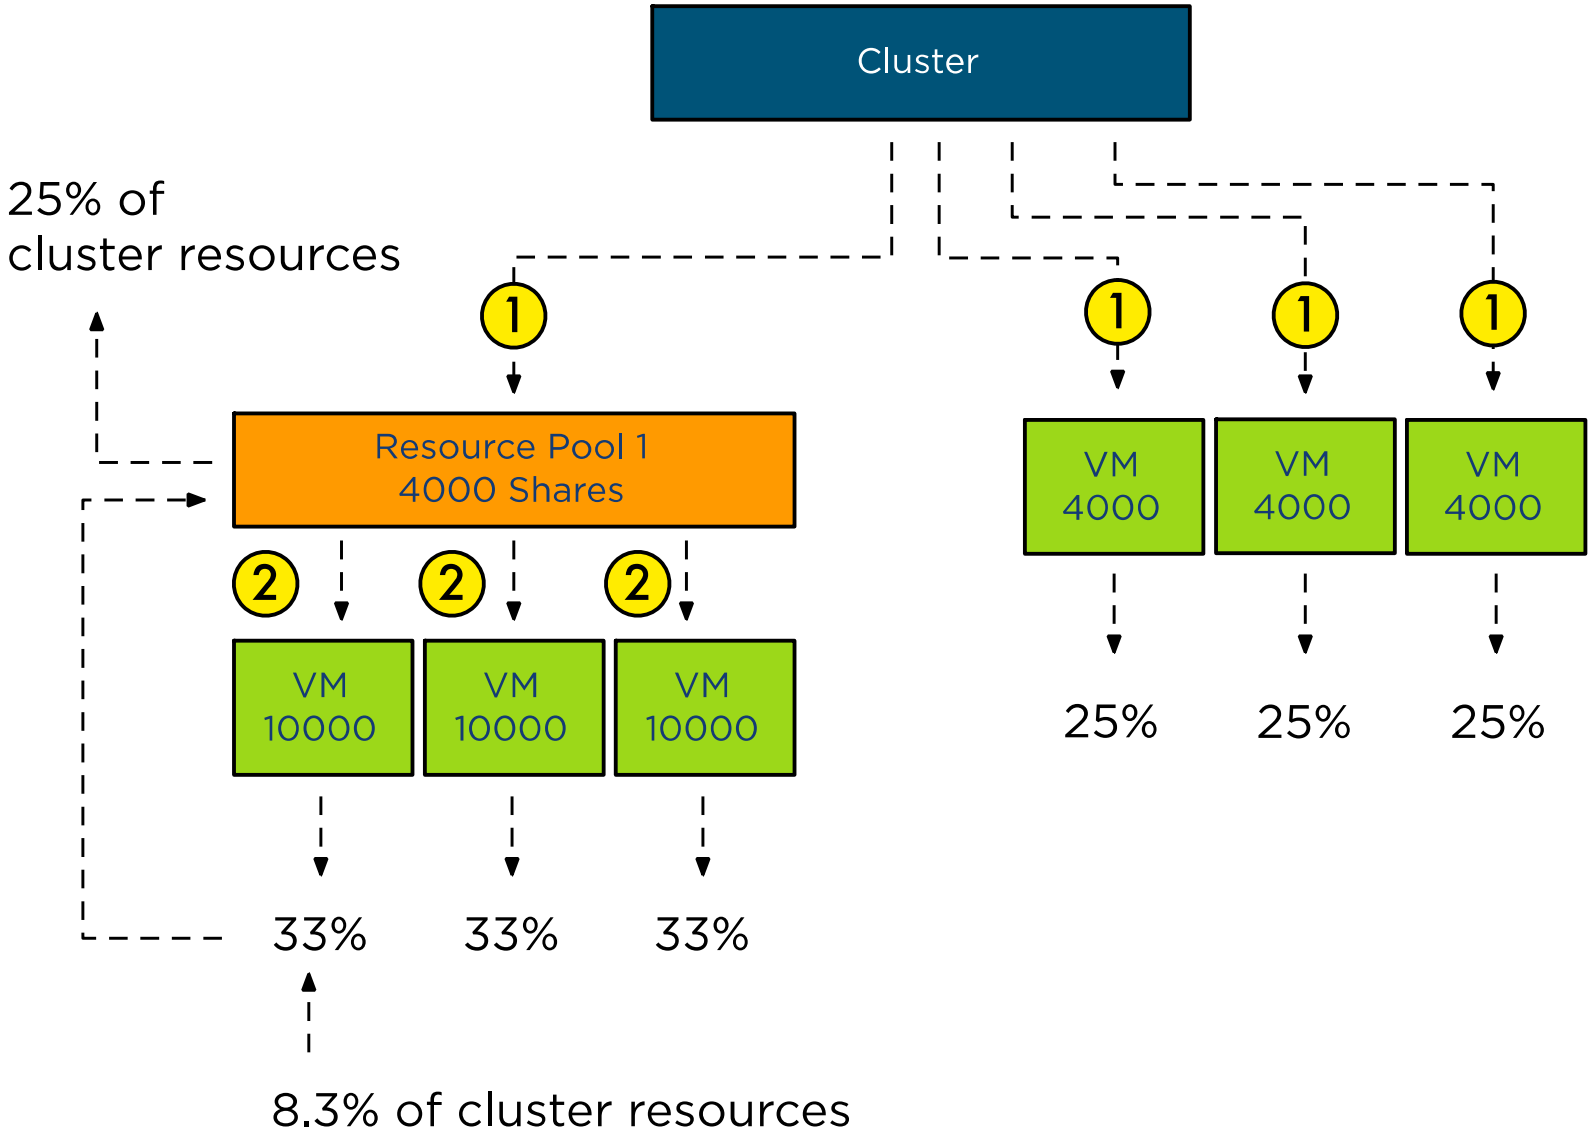

Test-DEV Resource Pool Design

- Typical Test/DEV and production RP design
- 1 RP for Production (High)
- 1 RP for Staging, test and Development (Low)

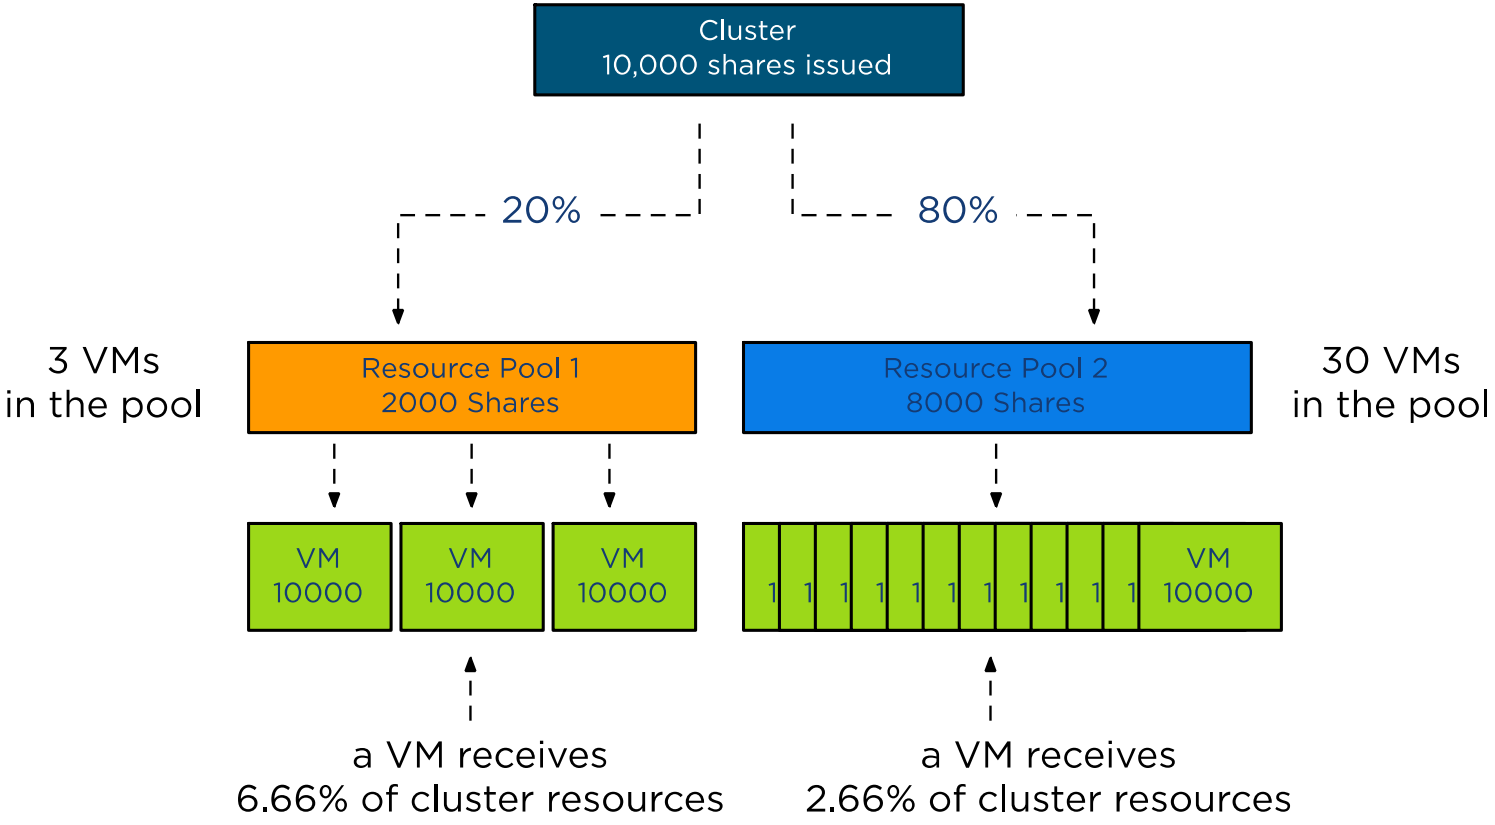

Solution

- Use custom shares
- Use script to define custom shares based on a pre-defined weight and the amount of VMs / vCPUs in the resource pool
- Use Resource Pool reservations
- Use Resource Pool limits

Reservations and Limits

- Reservations: Minimum physical resource allocation for the pool
- Limit: Maximum physical resource allocation for the pool
- The pool is limited or guaranteed an X amount of physical resources

Resource Pool Reservation

- **Reservations are not limits**
- **Shares are used for access non-protected resources**

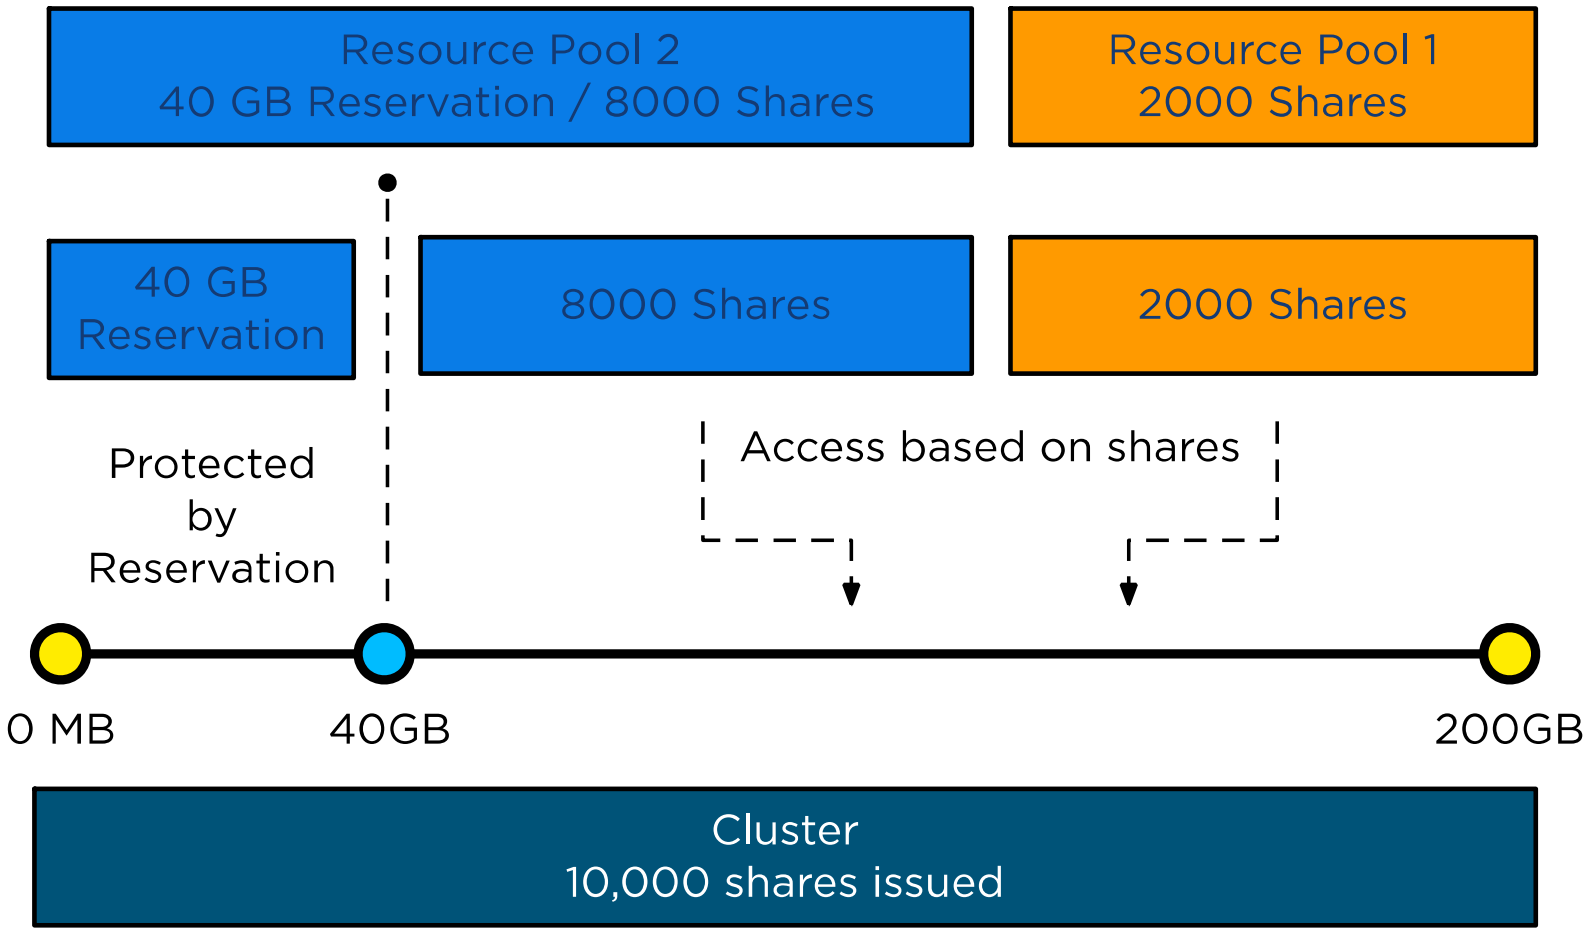

Resource Pool Reservation

- **Resource Pool “lends out” portion of physical resource to VM**
- **Dynamic entitlement of VM determines resource allocation within the pool**
  - Activity
  - Resource allocation settings defined on VM-level

VM Reservation in a Resource Pool

- **Extracted from RP-level reservation**
- **Static solution**
- **VM protected memory unavailable for other VMs**
- **Impact on HA Admission Control**
  - VM-Level Reservation impacts slot sizes
  - RP-Level Reservation is ignored by HA

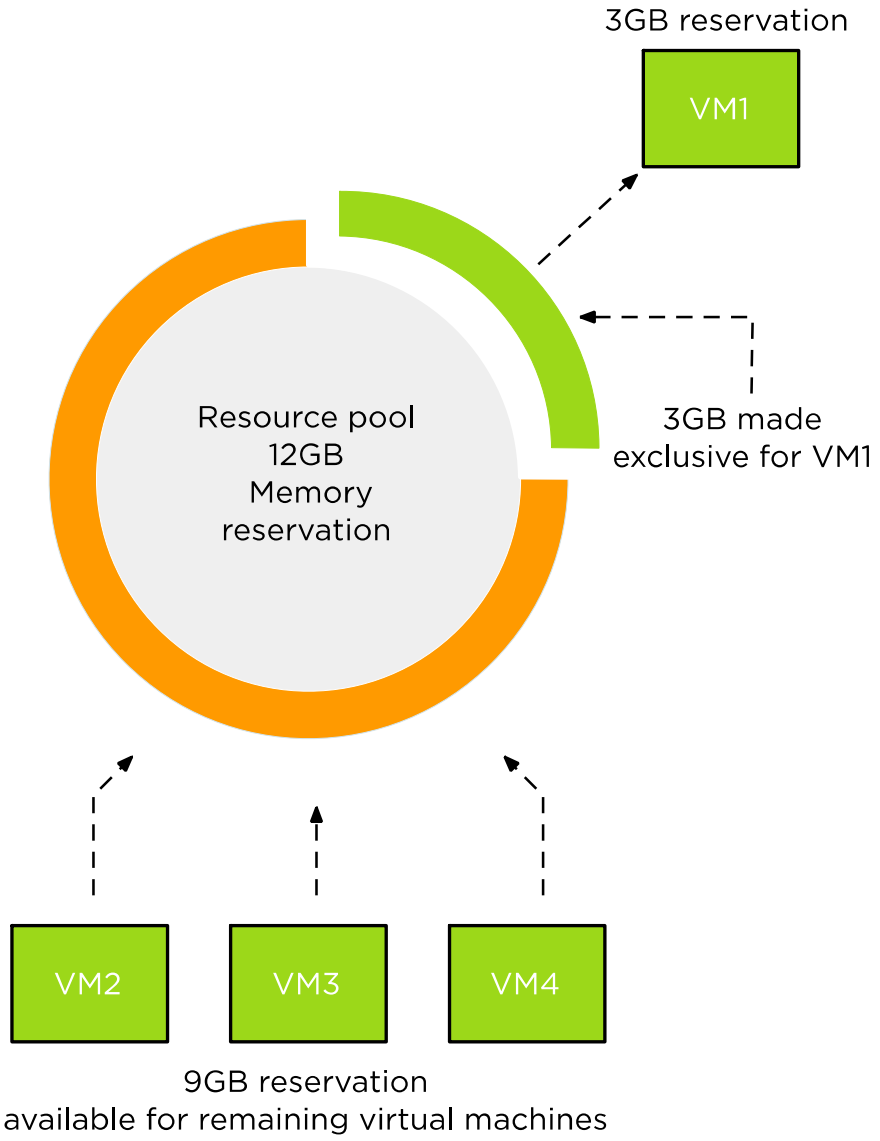

VM Level Reservation in a Resource Pool—Admission Control

- **Reduces Consolidation Ratio**
- **Resource Pool Reservation satisfies VM reservation**

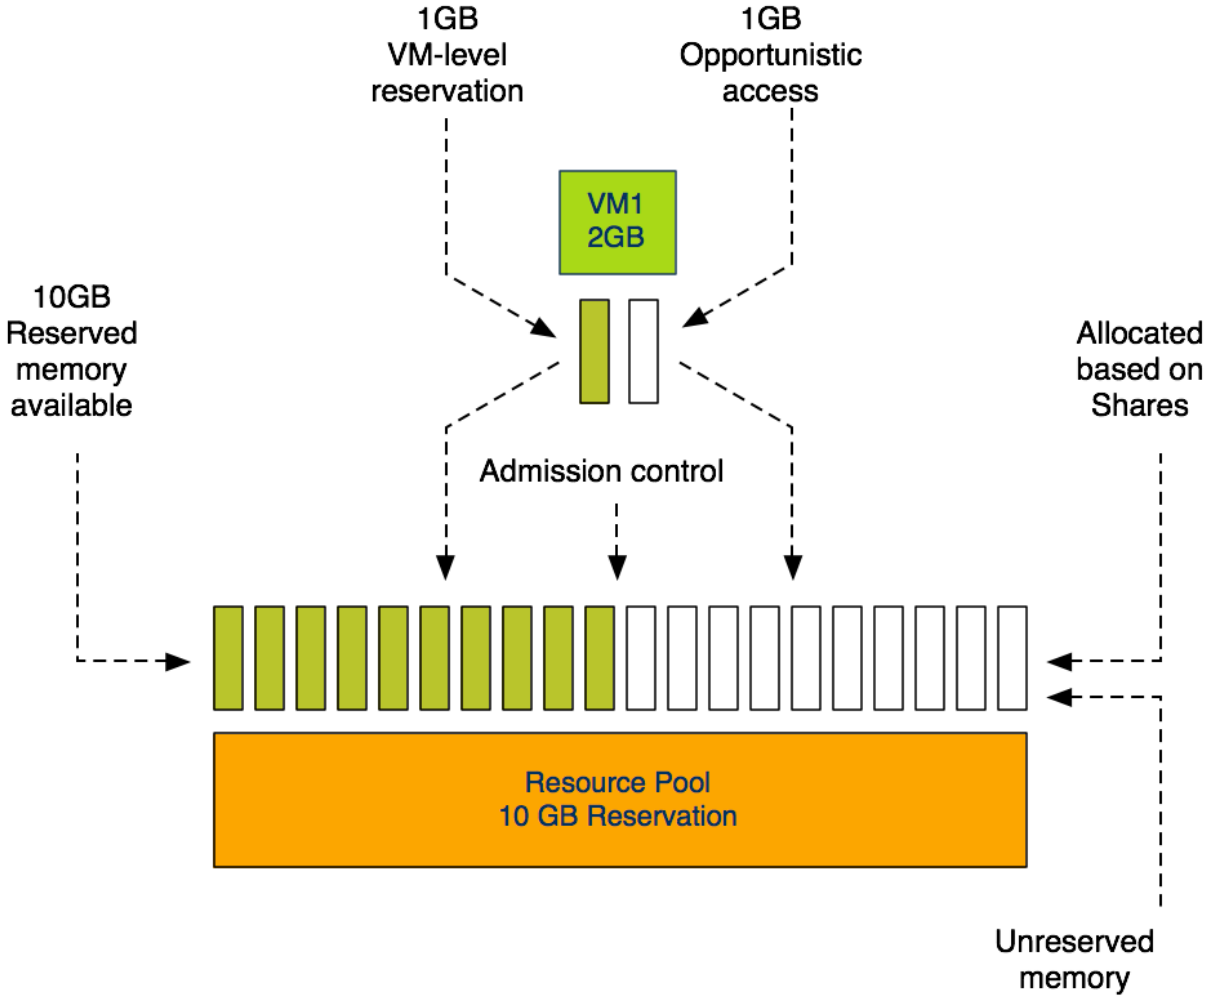

VM Level Reservation in a Resource Pool—Admission Control

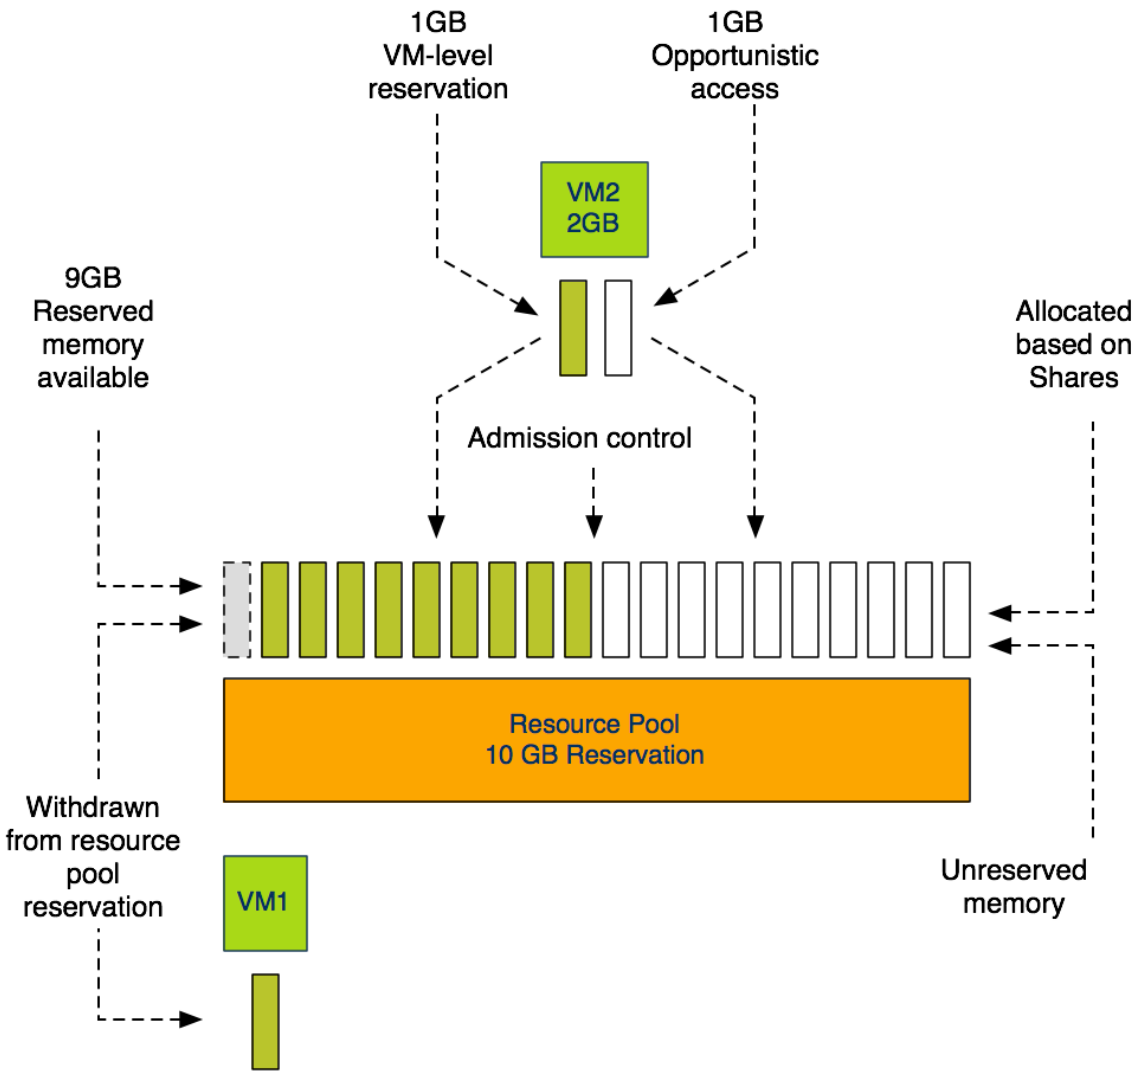

VM Level Reservation in a Resource Pool—Admission Control

- **No more Resource Pool reserved resource to satisfy VM reservation**
- **Result: Power-on operation failure**
- **Possible Solution:**
  - Expandable Reservation

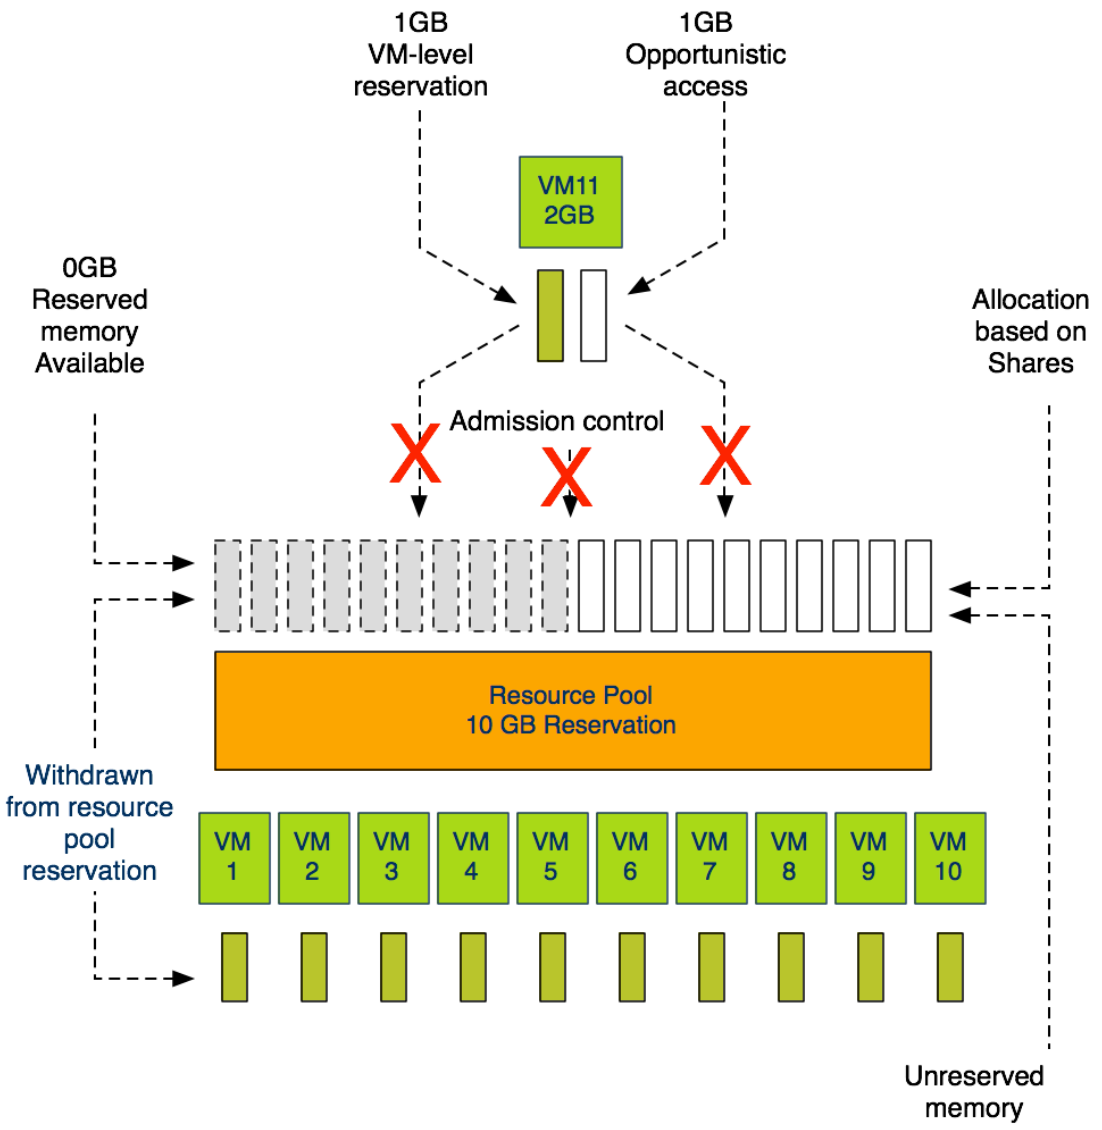

Expandable Reservation

- Allows resource pool to “confiscate” reserved resources from parent
- Does not consider reserved resources from siblings
- Search halts at resource pool with limit or disabled Expandable Reservation
- Is dynamic, once demand drops reserved resources are returned to parent

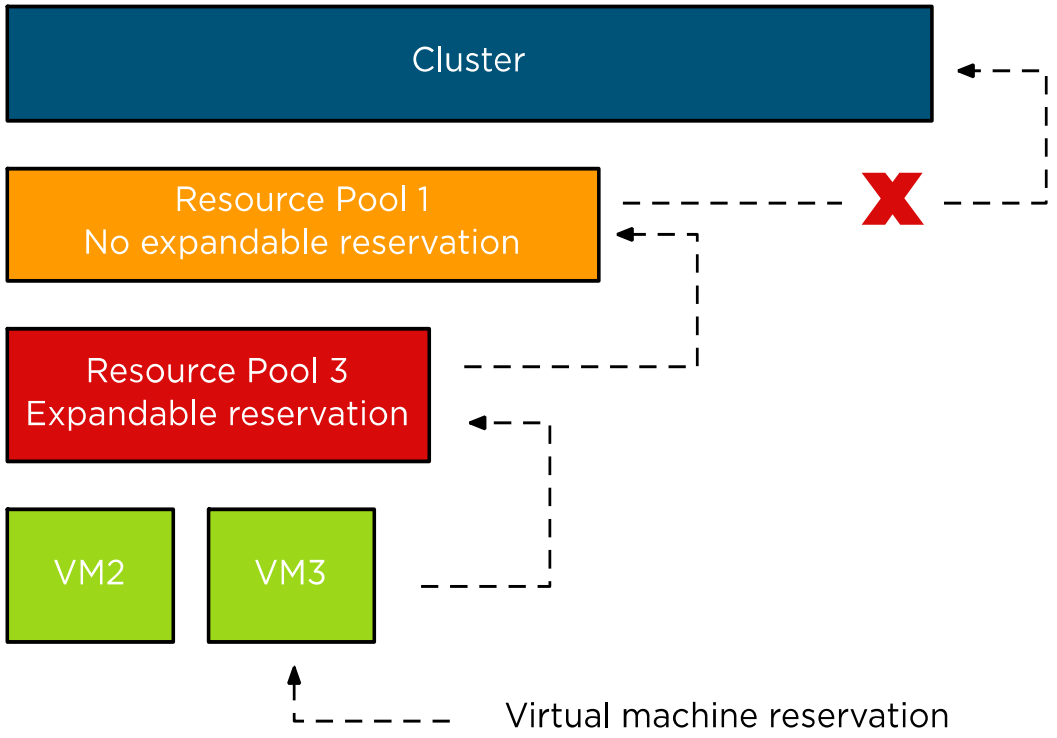

Limits

- Apply to child objects within the resource pool collectively
- Adjust the maximum amount of resource available per child object
- Depends on dynamic entitlement of child object
- When calculating limit on a resource pool take reservation and memory overhead reservation into account!

Tier 1 Application—Contention Mitigation Strategy

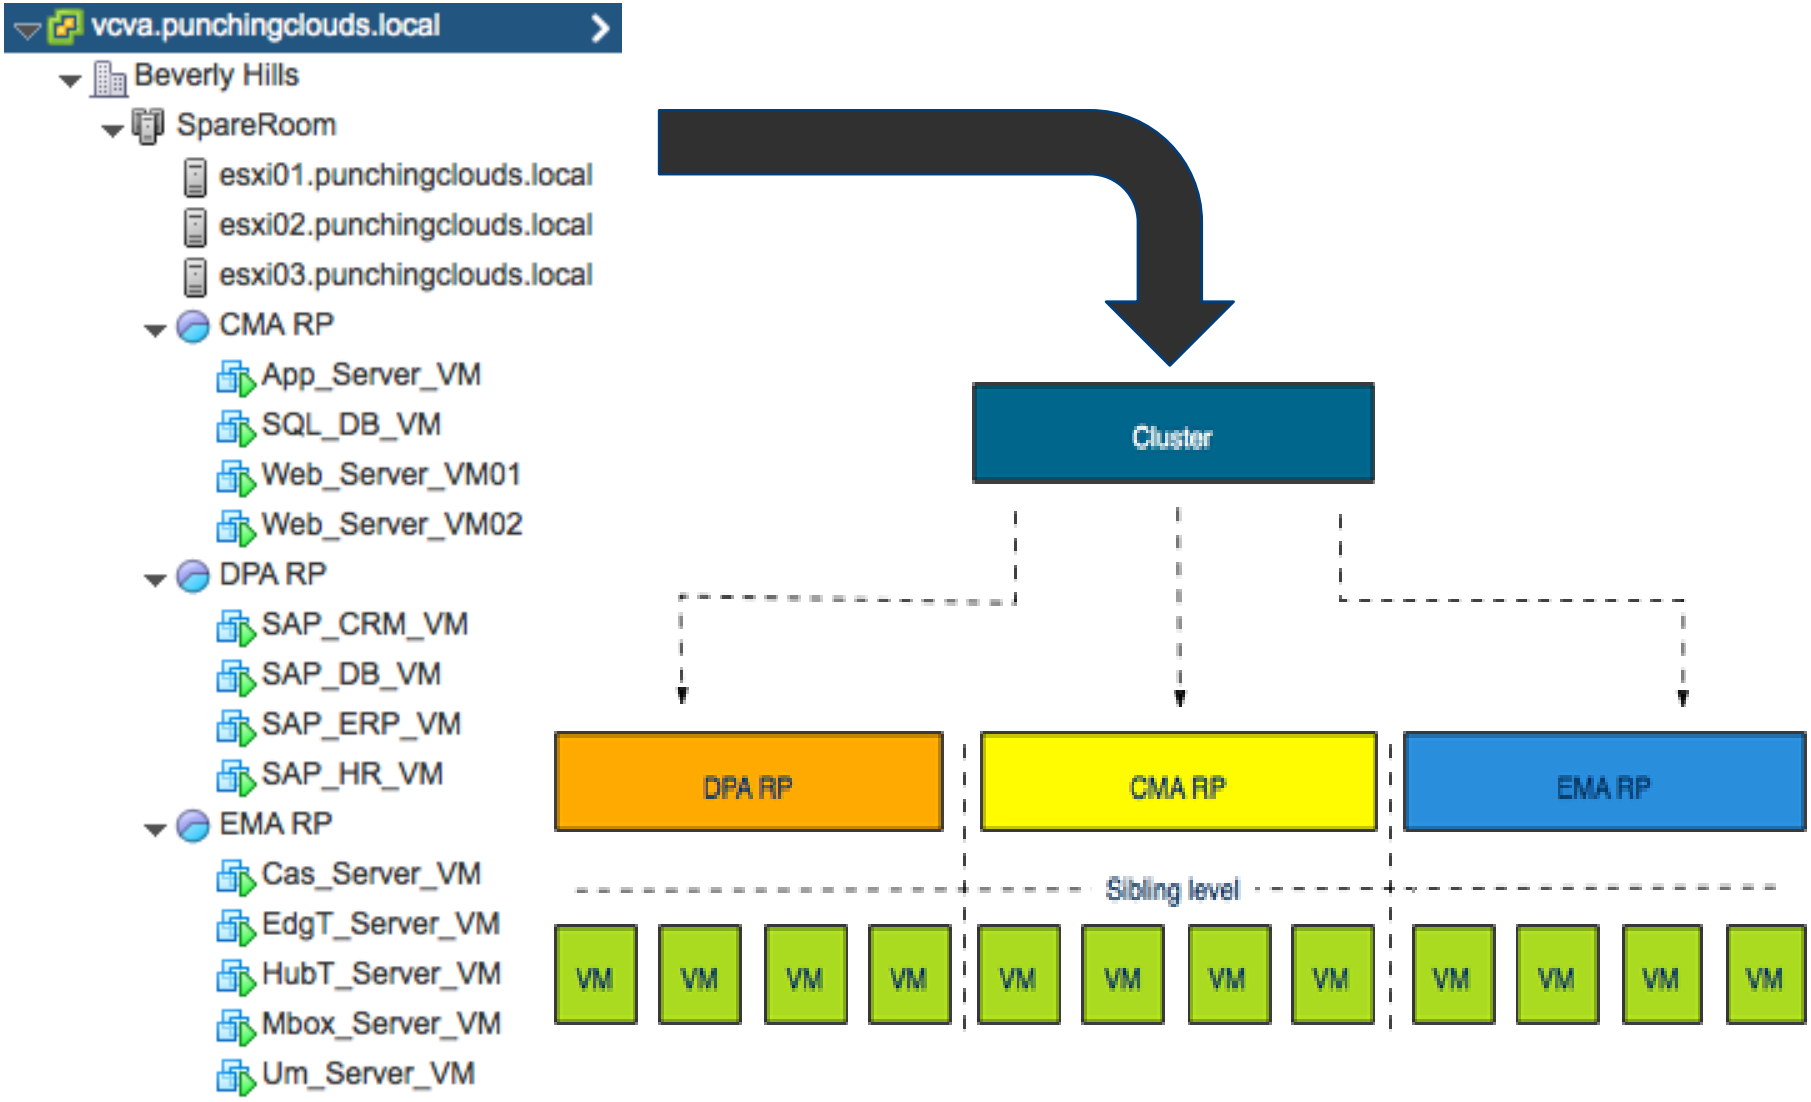

Contention Mitigation Strategy

- Apply Contention Mitigation Strategies at parent pool levels only
- Configure adequate resource and share value to suit application requirements and number of VMs
- Individual per VM shares and resource settings are not maintain in BCDR scenarios by SRM

Recommended Resource Pools Practices

- Consider the effect of over populating resource pools
- Consider hardware resource limitations
- Maintain a reservation of 10 to 20 resources available for host when allocation resources to resource pools

FILL OUT A SURVEY

EVERY COMPLETE SURVEY IS ENTERED INTO DRAWING FOR A $25 VMWARE COMPANY STORE GIFT CERTIFICATE

INF-VSP1683

VMware vSphere Cluster Resource Pools Best Practices

Frank Denneman, VMware, Inc.

Rawlinson Rivera, VMware, Inc.