Resource Management in VMware ESX Server 3

Mark Fei

Technical Instructor, VMware

Objectives

- To understand:
  - How resource pools allow you to define resource policies that are enforceable regardless of server heterogeneity or VMotion activity
  - How to use standalone resource pools for single-host resource policy control
  - The purpose and elements of a DRS cluster

Agenda

- How are VMs' CPU and memory resources managed?
- What is a resource pool?
- Managing a pool's resources
- A resource pool example
- Admission control
- DRS benefits and how it works
- DRS settings

VMs' CPU resource settings

- Limit
  - A cap on the consumption of CPU time by this VM, measured in MHz
- Reservation
  - A certain number of CPU cycles reserved for this VM, measured in MHz
  - The VMkernel chooses which CPU(s), and may migrate
- Shares
  - More shares means that this VM will win competitions for CPU time more often
- All the VCPUs in a VM must be simultaneously scheduled

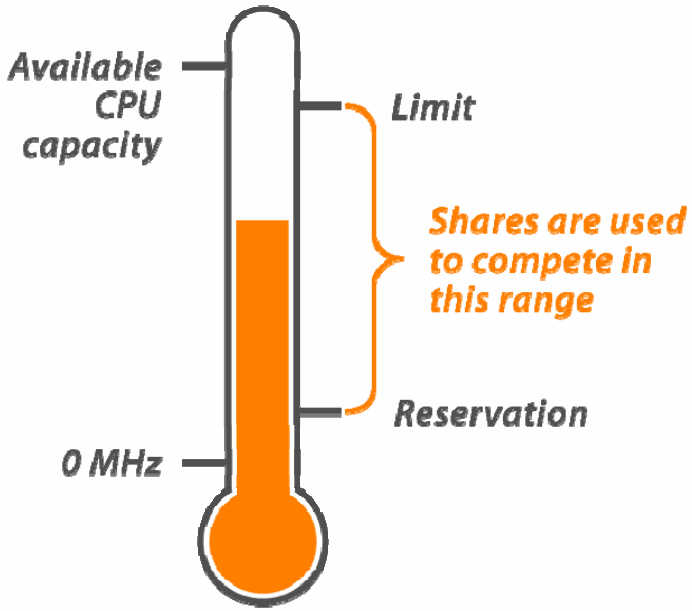

A virtual machine will only start if its reservation can be guaranteed

VMs' memory resource settings

- Limit
  - A cap on the consumption of memory by this VM, measured in MB
- Reservation
  - A certain amount of memory reserved for this VM, measured in MB
- Shares
  - More shares means that this VM will win competitions for memory more often
- VMkernel allocates a per-VM swap file to cover each VM's range between limit and reservation

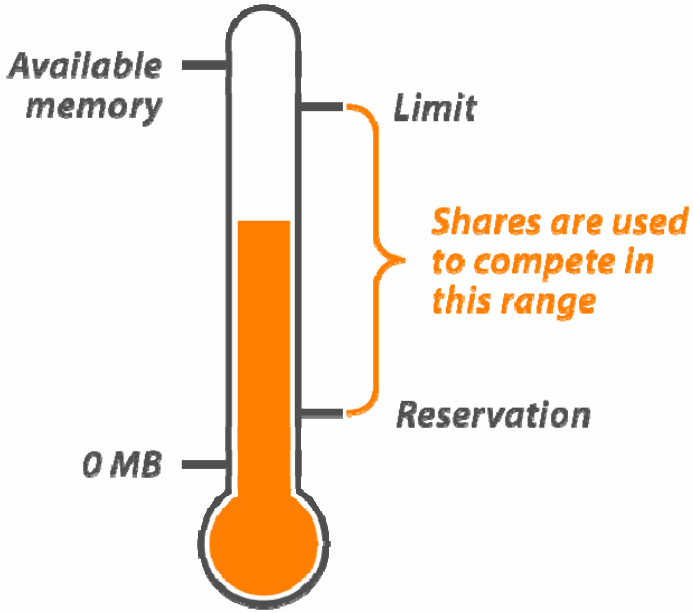

A virtual machine will only start if its reservation can be guaranteed

How VMs compete for resources

- Proportional-share system for *relative* resource management
  - Applied during resource contention
  - Prevents VMs from monopolizing resources
  - Guarantees predictable resource shares

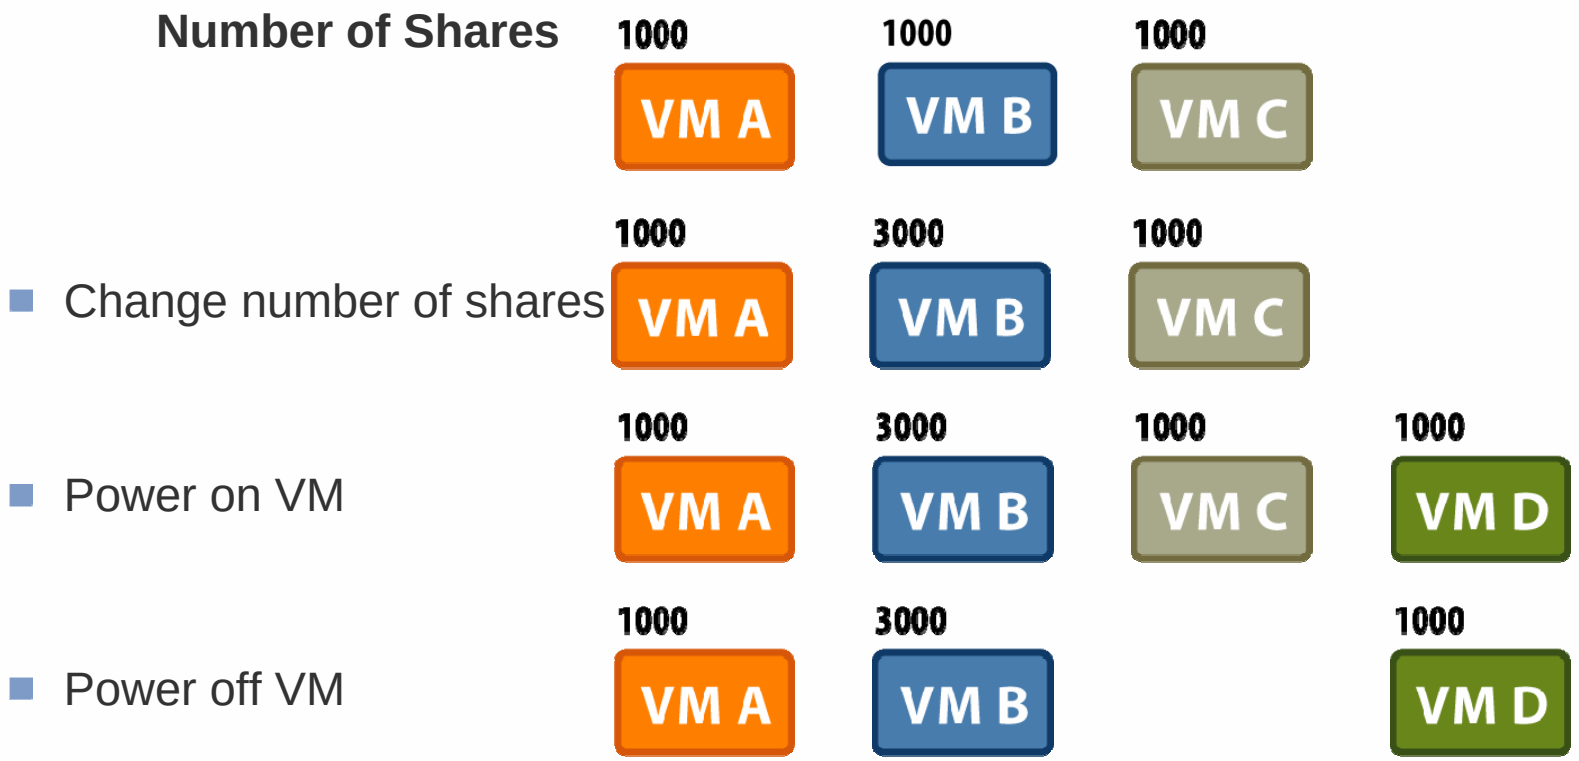

What is a resource pool?

- An object in the VirtualCenter inventory
  - A pool of CPU and memory for VMs
  - Can have associated access control and permissions
- Can be used on a stand-alone host or in a cluster (group of hosts)

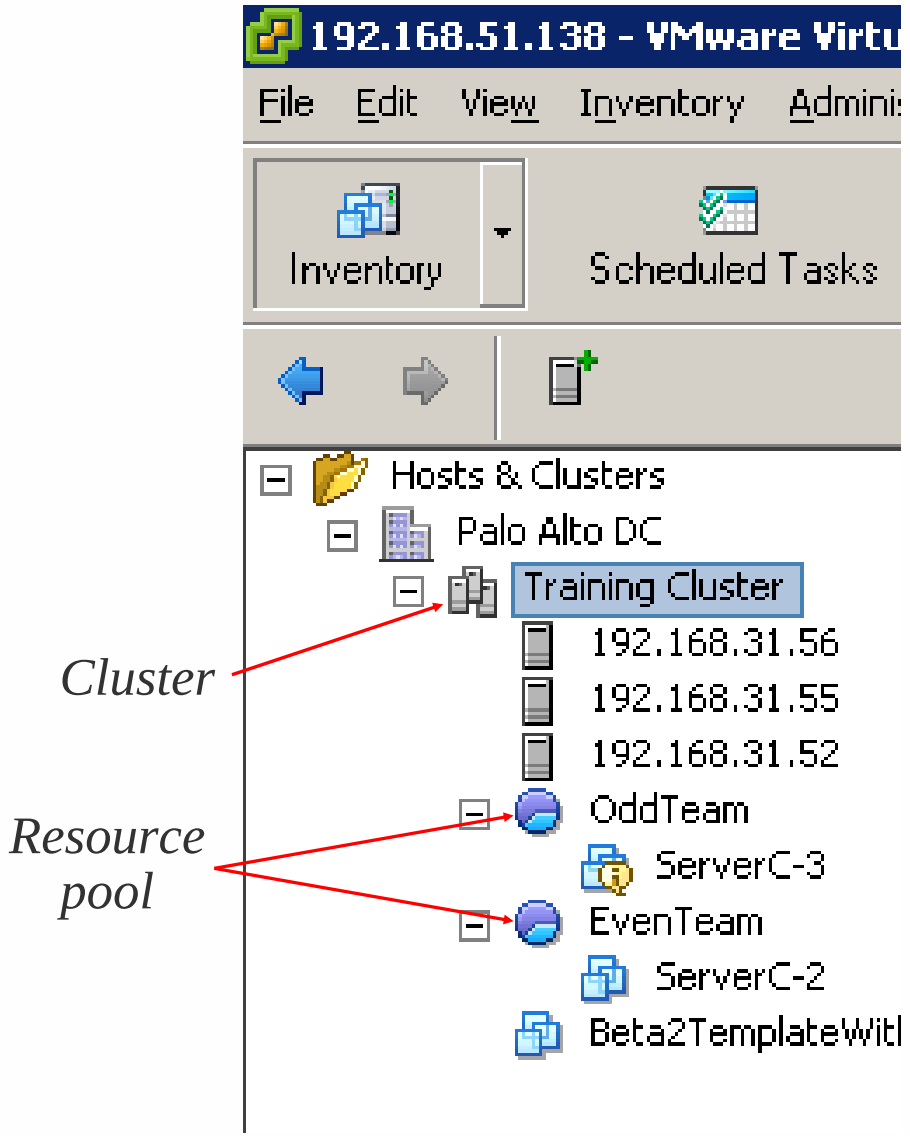

Managing a pool's resources

- Resource pools have the following attributes:
  - Shares
    - Low, Normal, High
  - Reservations, in MHz and MB
  - Limits, in MHz and MB
  - Expandable Reservation?
    - **Yes**: VMs and sub-pools may draw from this pool's parent
    - **No**: VMs and sub-pools may *only* draw from this pool, even if its parent has free resources

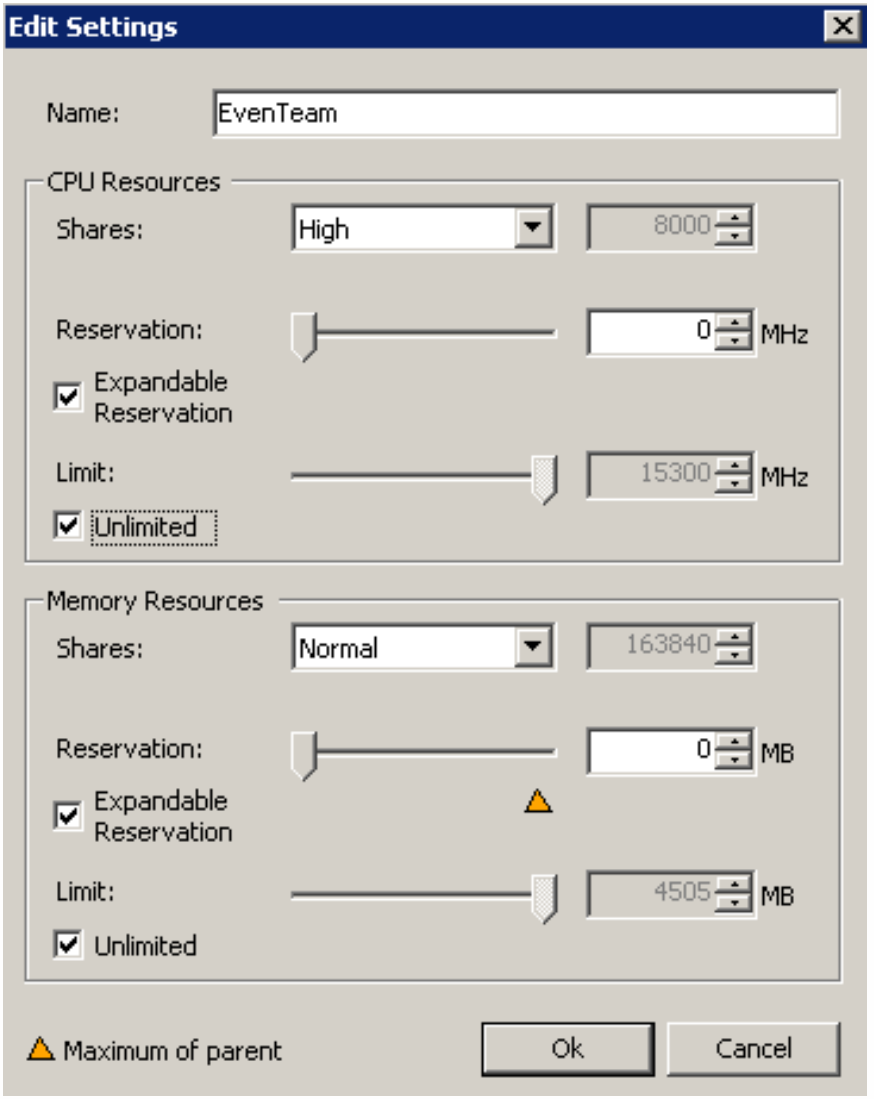

Scenario

- Company X's IT department has two internal customers
  - The finance department supplies 2/3 of the budget
  - The engineering department supplies 1/3 of the budget
- Each internal customer has both production and test/dev virtual machines
- We must cap the test/dev VMs' resource consumption

Resource pool example

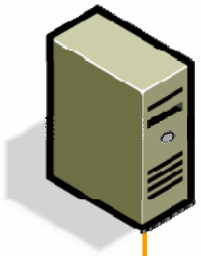

stand-alone host – Svr001
(root resource pool)
CPU: 12000 MHz
Memory: 4 GB

Engineering (Resource Pool)

CPU Shares: 1000
Reservation: 1000 MHz
Limit: 4000 MHz
Expandable Reservation: Yes

Eng-Test (VM)

CPU Shares: 1000
Reservation: 0 MHz
Limit: 4000 MHz

Eng-Prod (VM)

CPU Shares: 2000
Reservation: 250 MHz
Limit: 4000 MHz

Resource pools example: CPU shares

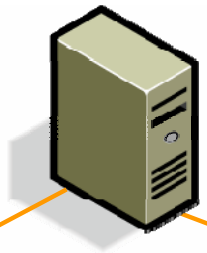

stand-alone host – Svr001
(root resource pool)

Engineering (Resource Pool)
CPU Shares: 1000

Finance (Resource Pool)
CPU Shares: 2000

Eng-Test (VM)
CPU Shares: 1000

Eng-Prod (VM)
CPU Shares: 2000

Fin-Test (VM)
CPU Shares: 1000

Fin-Prod (VM)
CPU Shares: 2000

Resource pools example: CPU contention

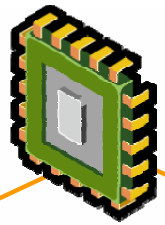

Svr001
All VMs below are running on same physical CPU (PCPU)

Engineering
CPU Shares: 1000
~33% of PCPU

Finance
CPU Shares: 2000
~66% of PCPU

Eng-Test
CPU Shares: 1000

Eng-Prod
CPU Shares: 2000

Fin-Test
CPU Shares: 1000

Fin-Prod
CPU Shares: 2000

Eng-Test gets ~33% of Engineering's CPU allocation = About 10% of the PCPU

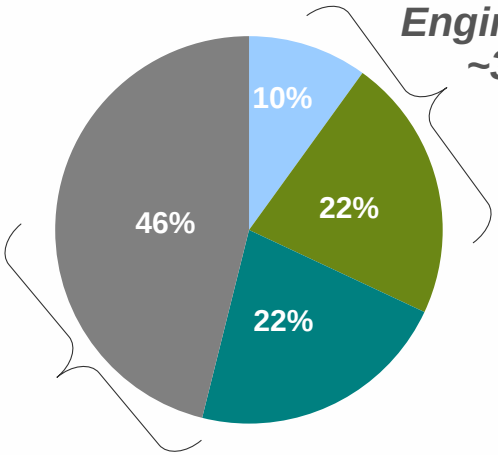

%age of PCPU allocation

Admission control for CPU and memory reservations

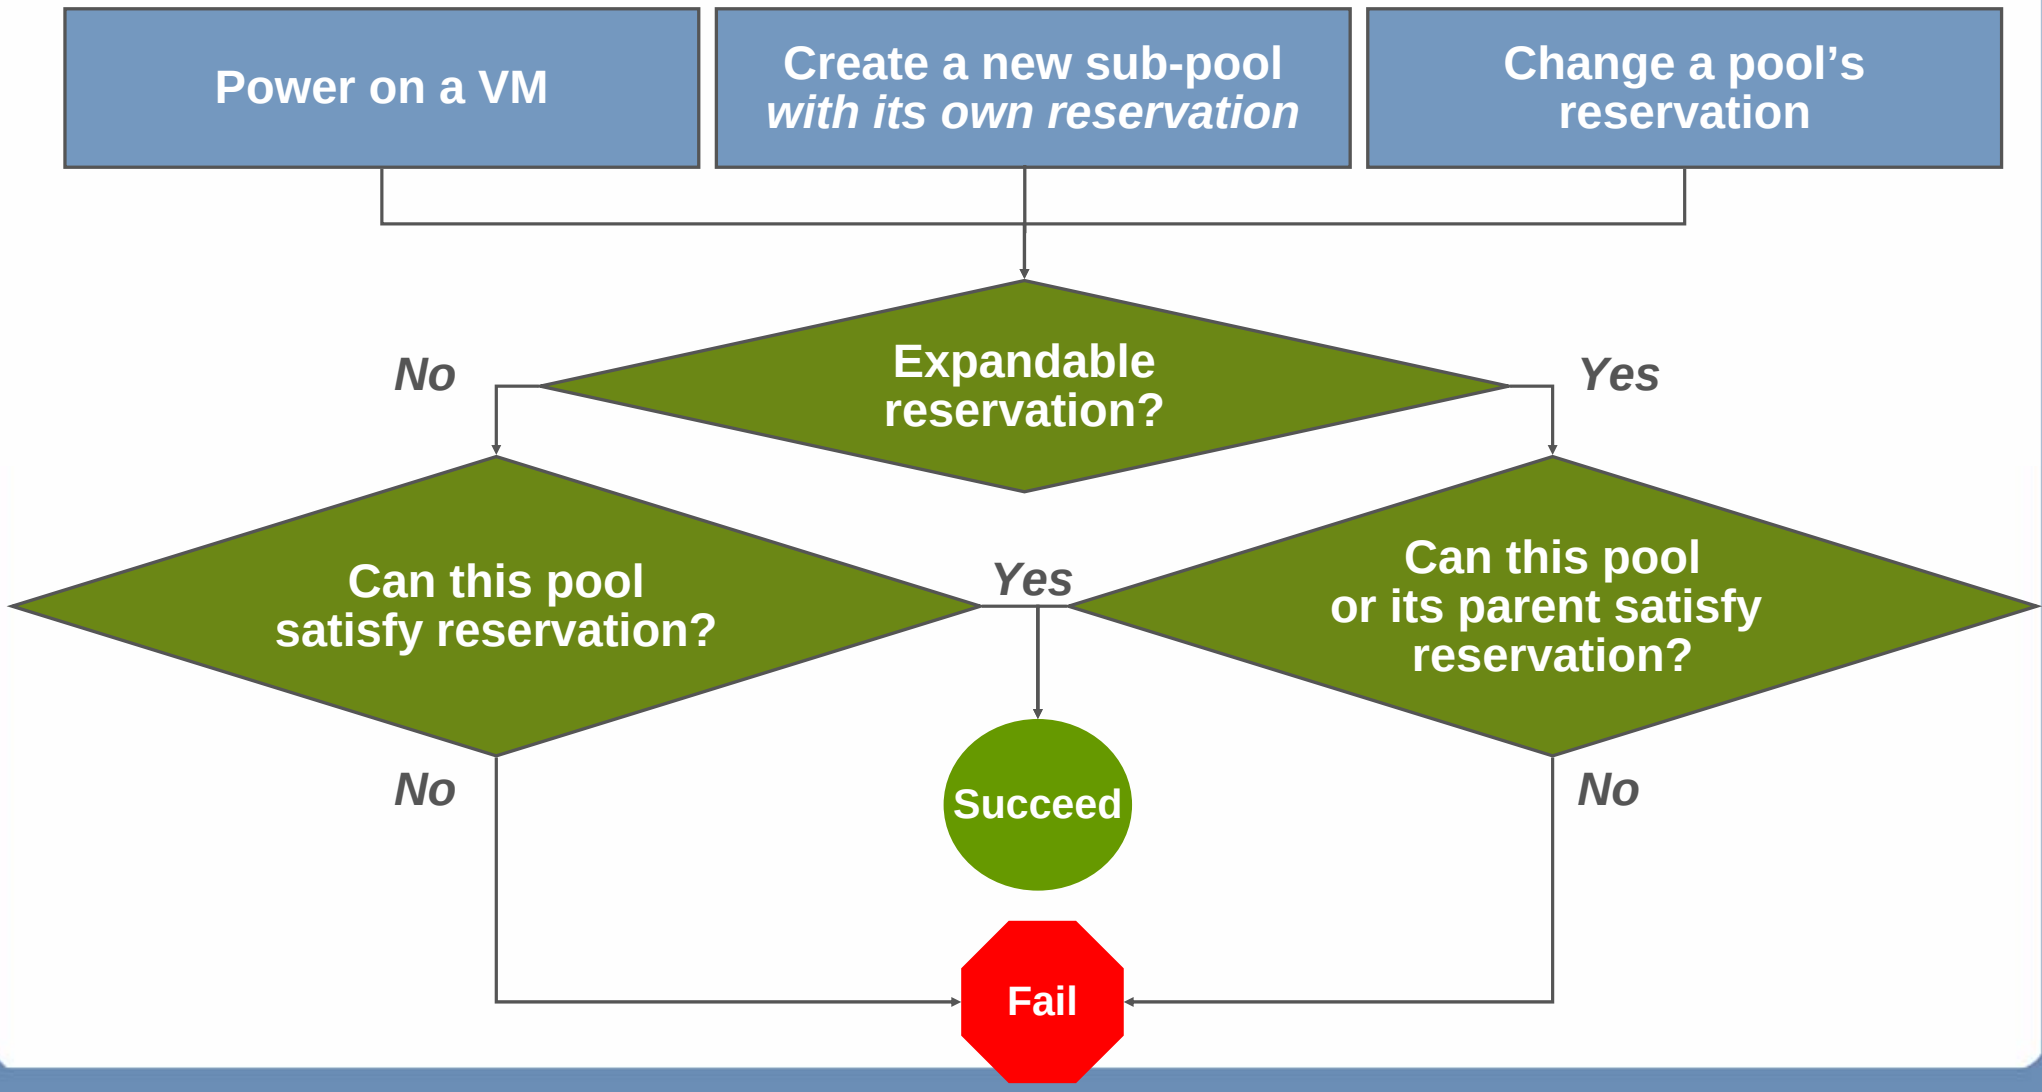

DRS cluster in the VirtualCenter inventory

- DRS allows you to aggregate several hosts' resources into one resource pool
- Create a cluster, enable DRS, add hosts
  - A DRS cluster is implicitly a resource pool
- You may divide each resource pool into sub-pools
  - And grant other administrators the privilege to make VMs and/or subpools there

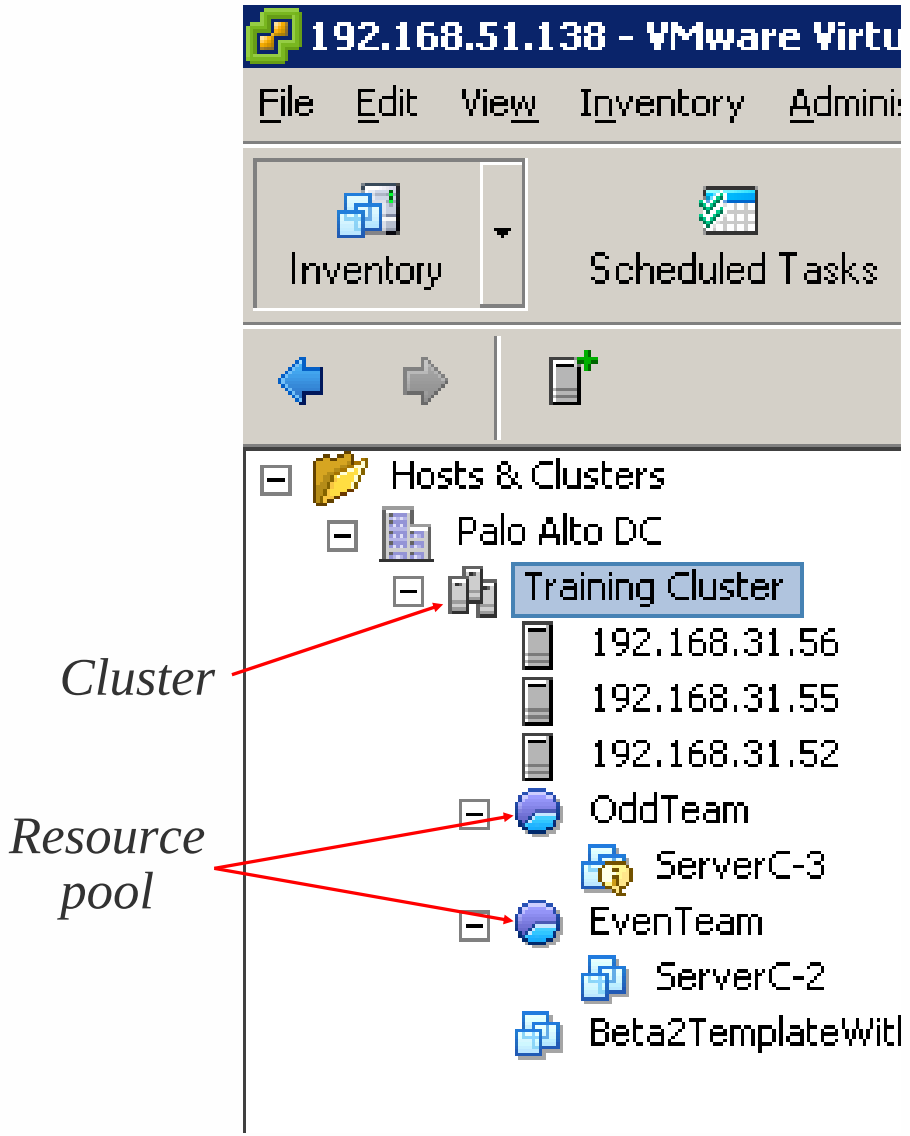

DRS: purpose and features

- Goals of DRS
  - Balance virtual machine load across hosts in cluster
  - Enforce resource policies accurately (reservations, limits, shares)
  - Respect placement constraints
    - Affinity and anti-affinity rules
    - VMotion compatibility (CPU type, SAN and LAN connectivity)
- Initial placement
  - Power on virtual machine in resource pool
  - Recommend host (prioritized list)
- Dynamic balancing
  - Monitor key virtual machine, pool, and host metrics
  - Deliver entitled resources to pools and VMs
  - Recommend migrations (prioritized list)

Move VM between ESX servers: VMotion migration

- A VMotion migration moves a VM that is powered on
- Why migrate using VMotion?
  - Improve overall hardware utilization
  - Allow continued VM operation while accommodating scheduled hardware downtime

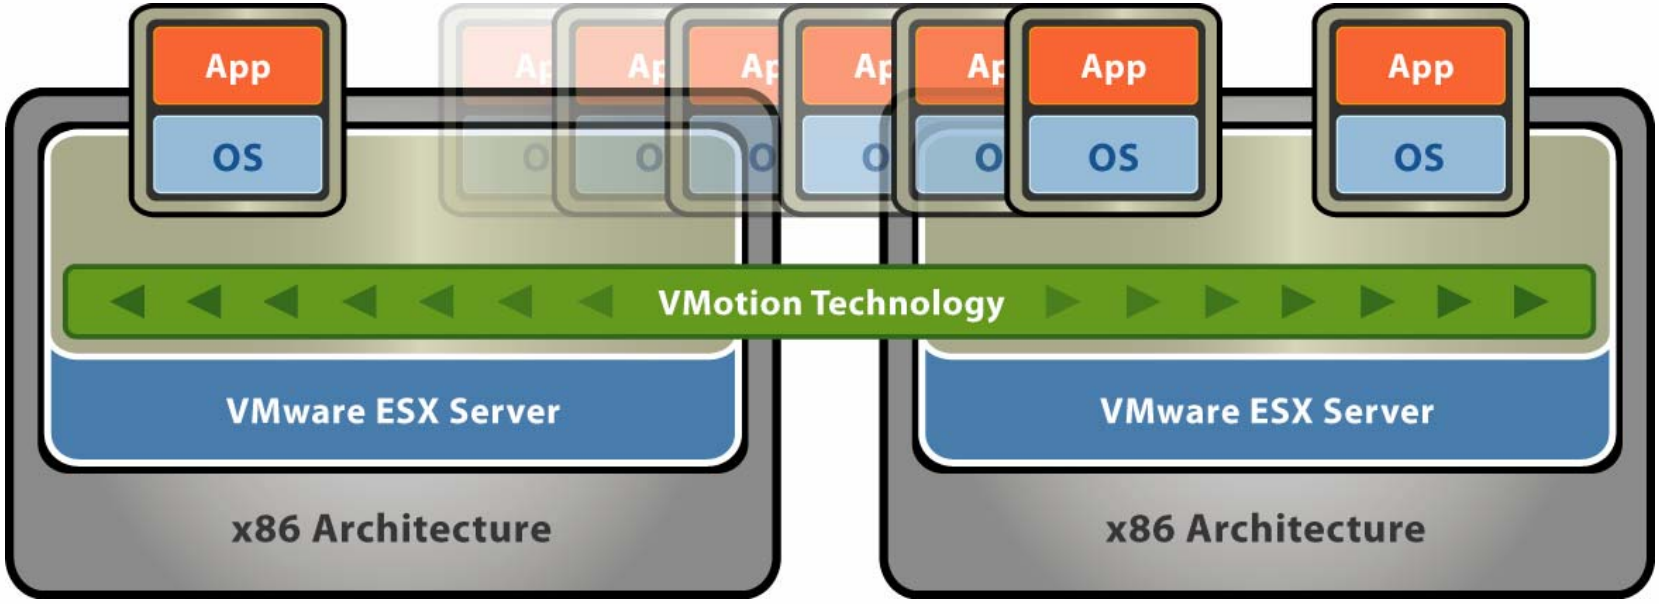

DRS cluster settings—automation level

Configure the automation level for initial placement of VMs and dynamic balancing while VMs are running

Automation level	Initial VM placement	Dynamic balancing
Manual	Manual	Manual
Partially-automated	Automatic	Manual
Fully-automated	Automatic	Automatic

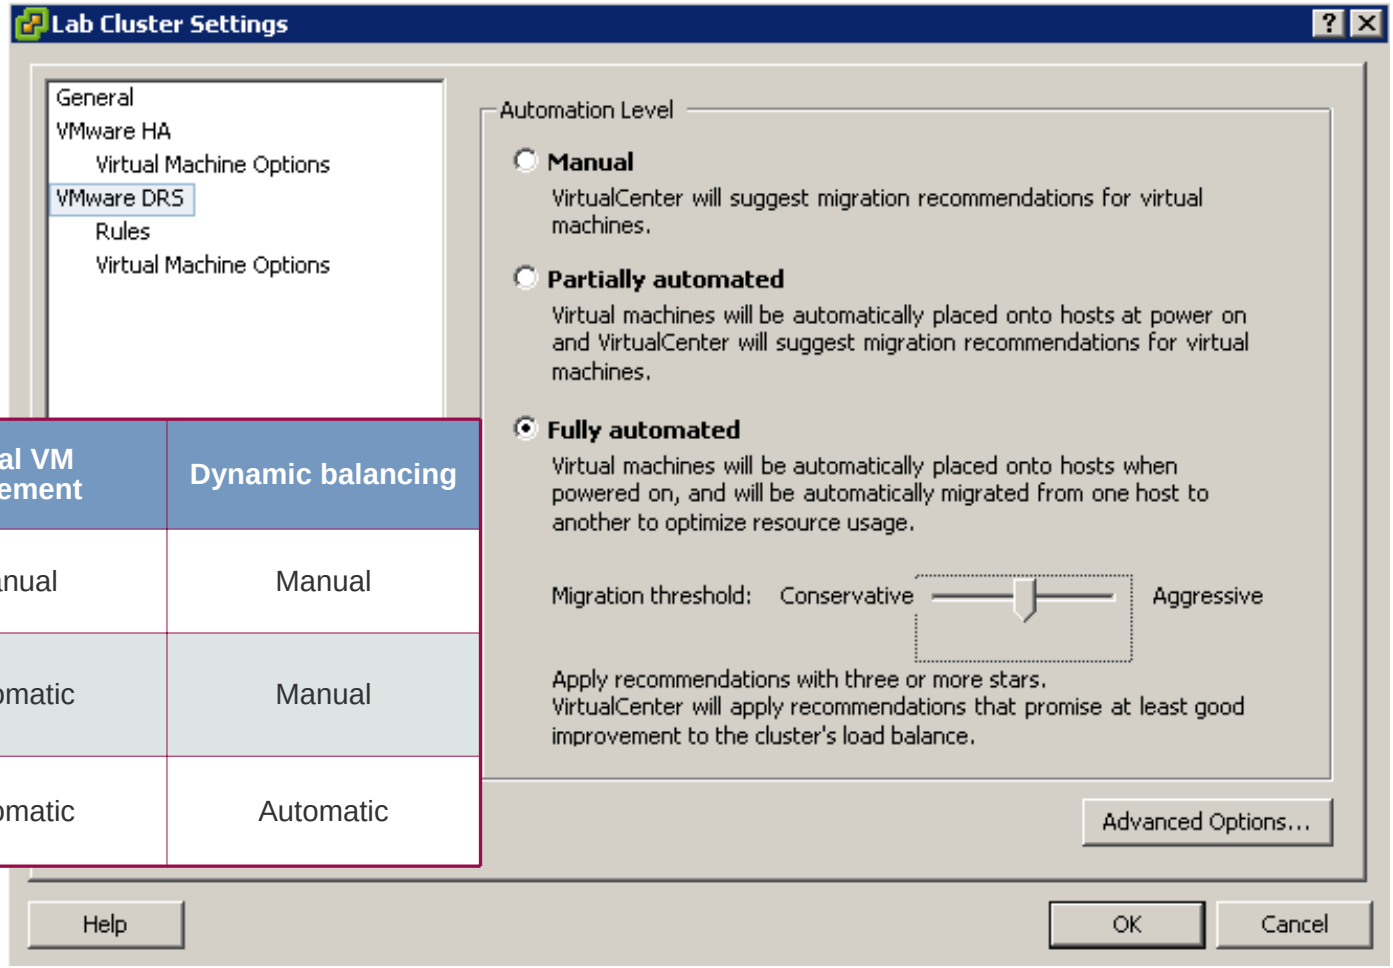

DRS cluster settings – placement constraints

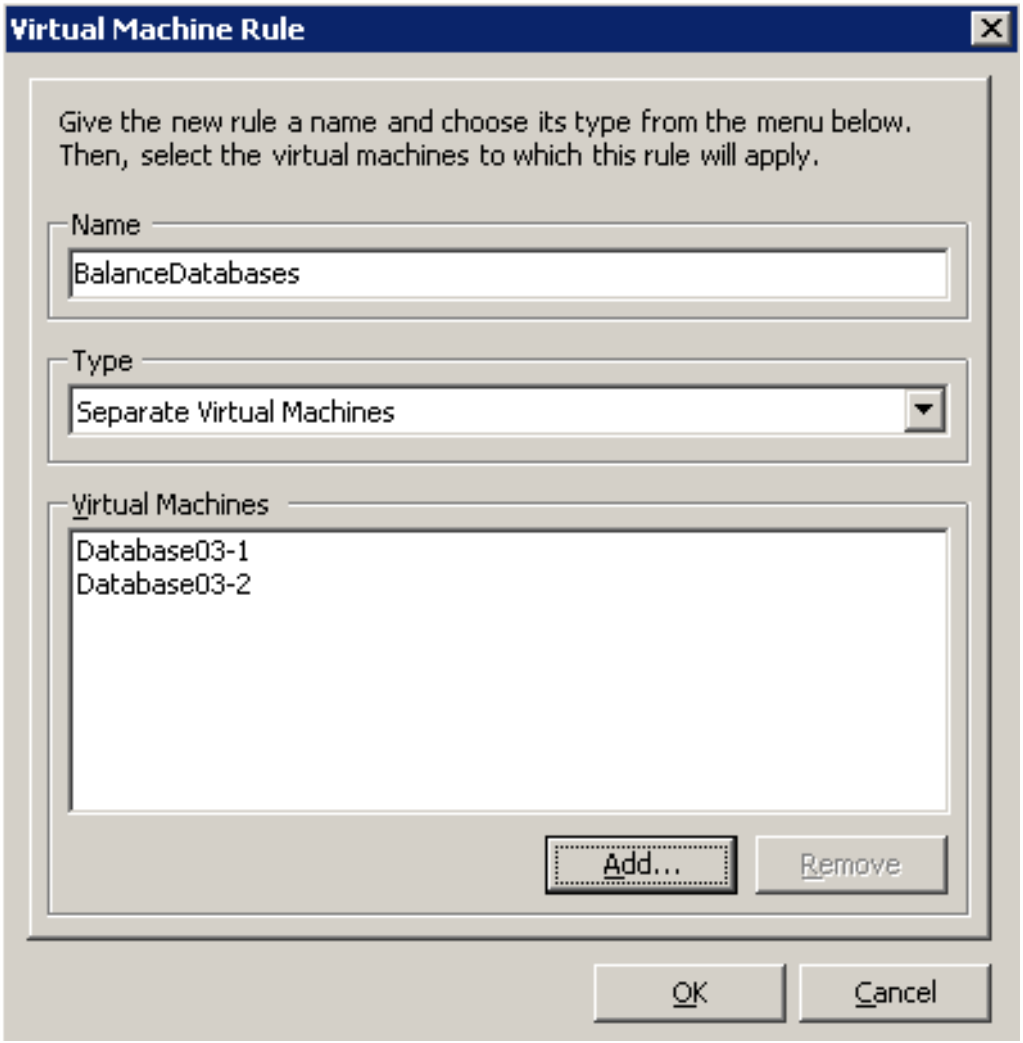

- Affinity rules
  - Run virtual machines on same host
  - Use for multi-VM systems where performance benefits from keeping network traffic internal
- Anti-affinity rules
  - Run virtual machines on different hosts
  - Use for multi-VM systems that load balance

DRS cluster settings – automation level per VM

- Optionally set automation level per VM

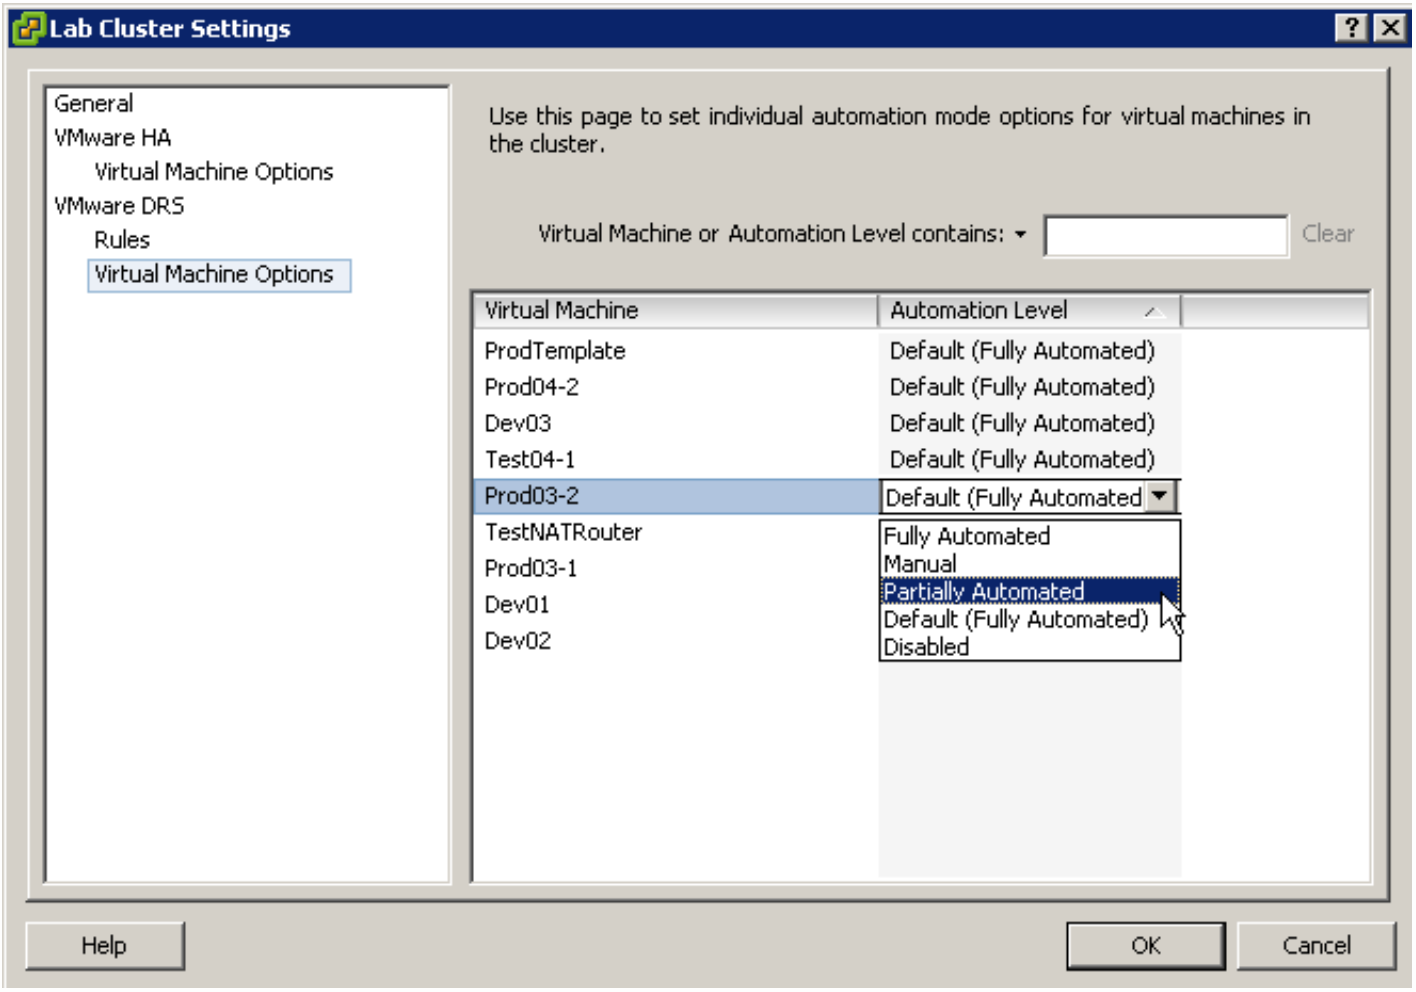

Best practices for DRS

- Because adding a host to a DRS cluster requires maintenance mode, plan to use VMotion to evacuate the host
- When DRS makes strong recommendations, follow them
  - Otherwise, balance and fairness may deteriorate
  - Some VMotion is necessary
- Enable automation
  - Choose default based on environment, comfort level
  - Let DRS autonomously manage most VMs
  - Use per-VM automation level overrides to accommodate sensitive VMs

Resource pools in a DRS cluster

Resource pools are used to subdivide the computing resources in a cluster

Cluster

Root Resource Pool
5 x (4.8 GHz, 4 GB)
(CPU = 24000 GHz, Memory = 20 GB)

Resource Pool 1
(CPU = 16 GHz, Memory = 12 GB)

CPU Shares: 2000
Reservation: 4 GHz
Limit: 16 GHz

Resource Pool 2
(CPU = 3 GHz, Memory = 3 GB)

CPU Shares: 4000
Reservation: 0
Limit: 3 GHz

VM

CPU Shares: 1000
Reservation: 0
Limit: 16 GHz

VM

CPU Shares: 3000
Reservation: 1 GHz
Limit: 2 GHz

VM **VM** **VM**

Delegated administration

- Joe administers cluster
  - Carves up cluster resources into pools, provides bulk allocations to pool admins
  - Has “Datacenter Administrator” VC role
- Jane administers Resource Pool 1
  - Carves up pool resources into smaller pools for users
  - Has “Resource Pool Administrator” VC role
- Ted administers VMs in Resource Pool 3
  - Allocates resources to VMs
  - Has “Virtual Machine Power User role”

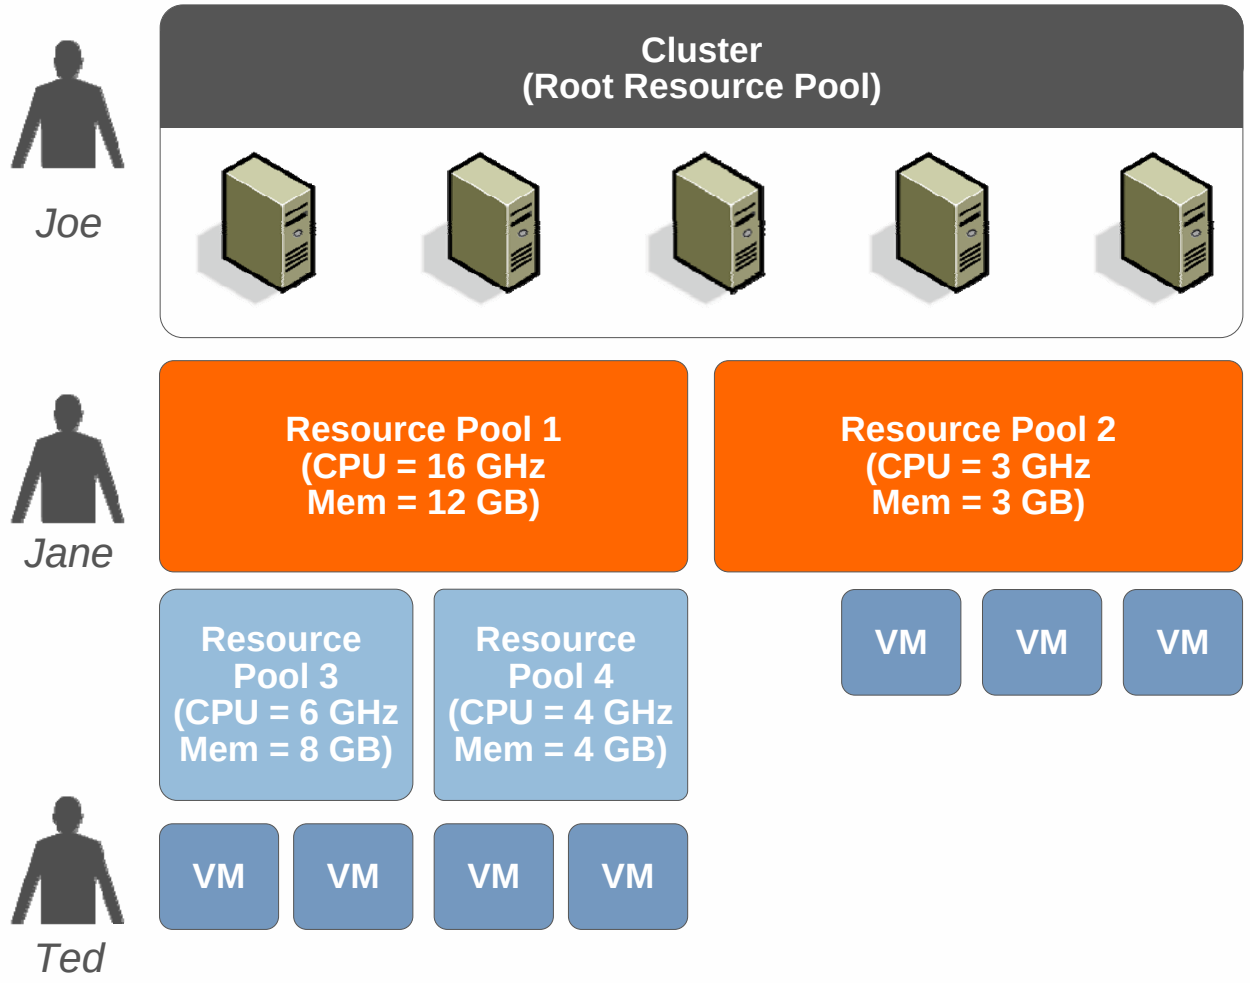

When to use expandable reservations

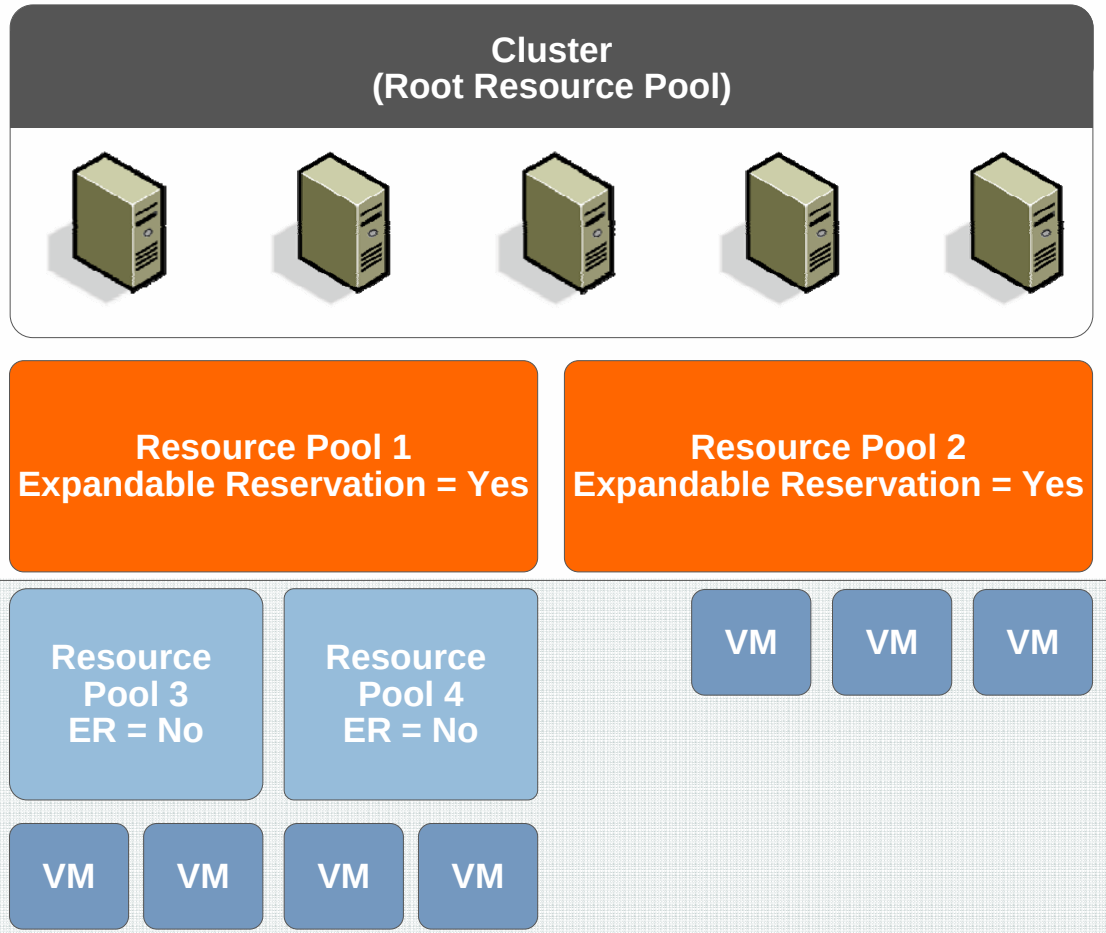

Pools for use within our team

Pools created by us
for use by our customers

Summary

- A resource pool has three attributes – reservation, limit and shares
- Resource pools can be created on standalone hosts or in DRS clusters
- VMotion is the underlying technology of VMware DRS
- A DRS cluster provides initial placement of VMs at power on and dynamic load balancing of running VMs

Questions?