

Cluster Resources

Max

Maximums (per DRS cluster): Hosts = 32 VMs (powered on) = 3000 (512 per host)
Maximums (per Resource Pool): Children = 1024 Tree depth = 8
Maximums (other): Hosts per datacenter = 500 RPs per host = 1600 RPs per cluster = 1600
Concurrent vMotions per host = 4 (1GbE), 8 (10GbE)

Terminology: Datacenters - mark organizational & vMotion boundaries. Clusters - gather host resources. Resource Pools - apply policies to clusters. DRS cluster is implicitly a resource pool. Resources include CPU, memory, power, storage & networking. See Storage section for Storage DRS, Profiles & Datastore Clusters. See Networking section for NIOC & Network Resource Pools.

☞ List resource group settings: `esxcfg-resgrp -l`

Resource pools: • *Shares* - low, normal, high & custom • *Reservations* - MHz(CPU)/MB(RAM) • *Limits* - MHz/MB • *Expandable reservation* - yes (can draw from parent's pool) - no (only from own pool) • *Shares* only apply during contention. Shares are relative to siblings (VMs or Res Pool). *Reservations* guarantee a minimum, can be allocated more. Only checked when VM is powered on. *Limits* are an upper bound, never exceeded; manage user expectations but can waste idle resources. *Resource pool admission control* prevents violations when VM is powered on or child pool created. *Fixed reservations* create strict isolation. *Expandable reservations* can borrow resources, do not automatically hunt upwards, but defines if reservations are considered by admission control. More flexible but provides less protection. Child pools actively reserve resources from parent even if VMs are powered off. Hierarchical resource pools require DRS enabled.

DRS: • *Manual* • *Partial* (Initial VM placement) • *Fully Automated* (Initial VM placement & Dynamic balancing). Migration threshold slider sets allowable host load imbalance. *Current Host Load Standard Deviation* - load imbalance (higher number increases priority level). *Current < Target* unless recommendations unapplied. Priority levels 1-5 (1 is highest). *Grafted from* pools created when adding host to DRS cluster & keeping host's resource pool hierarchy. Maintenance mode only clears VMs off host if DRS cluster is fully automated. Disabling DRS deletes resource pools & affinity rules - set DRS to manual to keep settings. DRS can be overcommitted/yellow (host failure) or invalid/red (usually direct client changes).

Affinity Rules: VM-VM keep VMs together/apart. VM-Host keep VMs on/off specific hosts. *Should* rule best effort. *Must* rule mandatory (for licensing). Rule conflicts - older wins, newer rule disabled. Obeying anti-affinity ranks over affinity. Disabled rules ignored.

DPM: uses IPMI, iLO or WOL (in that order). DRS & DPM thresholds are independent. Verify host's *DPM Last Time Exited Standby*. DPM level • *Off* • *Manual* (makes recommendations) • *Automatic*.

Links: DRS Deep Dive <http://www.yellow-bricks.com/drs-deepdive>
EVC (Enhanced vMotion Compatibility) FAQ <http://kb.vmware.com/kb/1005764>
EVC CPU compatibility <http://kb.vmware.com/kb/1003212>