



HA and DRS in VMware Cloud on AWS

Table of contents

HA and DRS in VMware Cloud on AWS	3
Introduction	3
vSphere DRS	4
.....	4
Elastic DRS	4
vSphere HA	5
VM and Application Monitoring	5
Monitoring	7
.....	8
To Conclude	9

HA and DRS in VMware Cloud on AWS

Introduction

VMware Cloud on AWS uses, and turbocharges, vSphere Distributed Resource Scheduler (DRS) and vSphere High Availability (HA) for unmatched availability and performance enhancements. Because VMware Cloud on AWS deploys SDDCs on top of global AWS infrastructure, customers benefit from the enhancements developed using HA and DRS capabilities.

vSphere Distributed Resource Scheduler (DRS) and vSphere High Availability (HA) are key features for vSphere environments, including VMware Cloud on AWS. DRS works to ensure workloads are getting the resources that they are entitled to by placement or balancing. vSphere HA ensures workload availability when a host in a cluster fails, is isolated, or when a VM heartbeat is interrupted.

As VMware Cloud on AWS is a managed service, both DRS and HA are enabled by default on all the clusters in all SDDCs. Customers cannot change these settings in VMware Cloud on AWS. This blog post I wrote to shed a bit more light on how these features are configured and cover some questions around virtual machine (VM) and application monitoring settings in VMware Cloud on AWS.

vSphere DRS

vSphere DRS is enabled and set to fully automated. The migration threshold is set for priority 1, 2, and 3 recommendations. The migration threshold specifies how aggressively DRS recommends balancing VMs using vMotion. The general recommendations are generated automatically based on resources demanded by the virtual machines, resource allocation settings (reservations, limits, and shares) the resources provided by each host, and the cost of migrating VMs.

vSphere DRS is Turned ON

vSphere DRS is Turned ON	
SCHEDULE DRS... RESTORE RESOURCE POOL TREE... EDIT...	
DRS Automation	
Automation Level	Fully Automated DRS automatically places virtual machines onto hosts at VM power-on, and virtual machines are automatically migrated from one host to another to optimize resource utilization.
Migration Threshold	Apply priority 1, priority 2, and priority 3 recommendations. DRS provides recommendations when workloads are moderately imbalanced. This threshold is suggested for environments with stable workloads. (Default)
Predictive DRS	Predictive DRS is disabled. vCenter will not respond to forecasted metrics from VMware Aria Operations.
Virtual Machine Automation	Individual virtual machine automation levels enabled.
> Additional Options	Expand for policies
> Power Management	Off
> Advanced Options	Expand for advanced options

Elastic DRS

Elastic DRS (EDRS) is unique to VMware Cloud on AWS. EDRS brings a policy-based approach for true cloud elasticity, for scaling out and scaling in scenarios. It is built on top of our developed Autoscaler logic. Autoscaler is the backend capability that is used to automatically add, and remove, hosts from VMware Cloud on AWS SDDCs. It is driven resource consumption and triggered by either pre-defined policies or custom policies configured by customers.



On vSAN enabled clusters, the storage baseline policy always makes sure to add host capacity if vSAN storage capacity is > 80%, another mechanism to ensure application availability for customers. More information about custom EDRS policies is [found here](#).

vSphere HA

The same Autoscaler logic is put to good use for vSphere HA in VMware Cloud on AWS. Typically, customers need to immediately remedy a failing host in a cluster by fixing or replacing it. Or customers incorporate spare hosts, which is not cost-effective. vSphere HA takes care of re-registering and powering-on VMs on surviving hosts in the cluster. But, depending on HA settings, a cluster is at risk of resource constraints or VM downtime if another host fails in the same timeframe.

VMware Cloud on AWS solves this by automatically adding a host if a host failure is detected. The goal is to have the number of hosts available that customers configured for a cluster. In the event of a host failure and vSphere HA triggers, a new host is immediately added to the cluster.

With vSAN enabled instance types, the data is resynced and vSphere HA powers on the VMs on the new host. The faulty host is removed. No customer interaction required, this is done automatically and maximizes VM availability.



The default vSphere HA settings in VMware Cloud on AWS are accounting for host failure, host isolation, datastore protection, and VM/Application monitoring. Admission control, depicting the failover capacity in a cluster, is set to percentage-based, with the actual percentage value depending on the number of hosts in a cluster. This example screenshot is created on 3-host cluster.

vSphere HA is Turned ON

EDIT...

Runtime information for vSphere HA is reported under [vSphere HA Monitoring](#)

Proactive HA is Turned OFF

EDIT...

Failure conditions and responses

Failure	Response	Details
Host failure	✓ Restart VMs	Restart VMs using VM restart priority ordering.
Proactive HA	⊘ Disabled	Proactive HA is not enabled.
Host Isolation	✓ Power off and restart VMs	VMs on isolated hosts will be powered off and restarted on available hosts.
Datastore with Permanent Device Loss	✓ Power off and restart VMs	Datastore protection enabled. Always attempt to restart VMs.
Datastore with All Paths Down	✓ Power off and restart VMs	Datastore protection enabled. Ensure resources are available before restarting VMs.
Guest not heartbeating	✓ Reset VMs	VM and application monitoring enabled. VMs will be reset.

6 items

Admission Control

Failover capacity is defined by reserving a percentage of the cluster resources.
 Reserved failover CPU capacity: 33%
 Reserved failover Memory capacity: 33%

Datastore for Heartbeating [Expand for details](#)

VM and Application Monitoring

The vSphere HA failure scenario of 'Guest not heartbeating' is also configured. VM and application monitoring is enabled. It helps

to monitor the guest OS and applications running inside a VM, leveraging VMtools. If heartbeats from VMtools are interrupted, or no IO is generated by the VM, it is likely the Guest OS has crashed. A vSphere HA event can be triggered, resetting the VM, depending on settings. The default setting in VMware Cloud on AWS are not visible for customers. But are configured as the following screenshot;

Edit Cluster Settings | Cluster-1



VM Monitoring Only

Turns on VMware tools heartbeats. When heartbeats are not received within a set time, the VM is reset.

VM and Application Monitoring

Turns on application heartbeats. When heartbeats are not received within a set time, the VM is reset.

VM monitoring sensitivity

Preset

Low ● High

Custom

Failure interval seconds

Minimum uptime seconds

Maximum per-VM resets

Maximum resets time window








No window

Within hrs

Monitoring

While customers have no access to change settings (see the 'edit' button greyed out), they can monitor the behavior. The vSphere Client provides insights into both DRS and HA specifics like DRS history and overall HA information.

DRS History

Time	DRS Actions
03/01/2024, 1:02:45 PM	 Place db101 on host 10.21.3
03/01/2024, 1:02:43 PM	 Place db100 on host 10.21.3
03/01/2024, 1:02:40 PM	 Place app102 on host 10.21
03/01/2024, 1:02:37 PM	 Place app101 on host 10.21.
03/01/2024, 11:38:58 AM	 Place nh-jumphost on host
03/01/2024, 10:57:52 AM	 Place nh-jumphost on host
03/01/2024, 10:44:11 AM	 Place nh-jumphost on host

Virtual Machines ⋮

Protected	57
Unprotected	0

Advanced Runtime Info ⋮

Cluster total memory	3 TB
Failover capacity (Memory)	33% (1,013.6 GB)
Cluster total CPU	556.8 GHz
Failover capacity (CPU)	33% (183.74 GHz)
Failover resource auto-computation	On

Hosts ⋮

Primary	10.21.32.9
Hosts connected to primary	2
Hosts not connected to primary	0
vSphere HA agent not reachable	0
vSphere HA agent configuration error	0
Hosts failed	0
Network isolated	0
Network partitioned	0
vSphere HA agent initializing	0
Disconnected from vCenter	0
Standby mode	0
Maintenance mode	0
vSphere HA agent unconfiguration failures	0
vSphere HA agent scheduled for retry	0

To Conclude

VMware Cloud on AWS ensures workload performance and availability so our customers don't have to worry about that.

