

16 | VM TO HOSTS AFFINITY RULE

JUL stored in: [VMware](#) and tagged:

VMware vSphere 4.1 introduces a new affinity rule, called "Virtual Machines to Hosts" (VM-Host). This new rule is available in vSphere 4.1 DRS clusters in addition to the existing (anti) affinity rule, which is now called VM-VM affinity rule. The new VM-Host affinity rule provides the ability of placing a group of virtual machines on a subset of hosts inside the cluster. The new rule can be very useful in blade system environments and for honoring ISV license requirements. Rules can be created to ensure that virtual machines run on ESX hosts in different blade chassis for availability reasons, or the complete opposite and limit the virtual machines to ESX hosts inside a blade chassis to optimize network speeds by keeping network traffic inside the blade chassis. VM-Host are also very useful to fulfill the requirements of special ISV license models as well, for example restricting Oracle database virtual machines to run only on ESX hosts which are licensed by Oracle.

Difference between VM-Host affinity rules and VM-VM rules

The VM-Host affinity rule differs from the VM-VM rule. A VM-Host (anti) affinity rule specifies the (anti) affinity between a group of virtual machines and a group of ESX hosts inside the cluster, whereas a VM-VM (anti) affinity rule only specifies the (anti) affinity between individual virtual machines.

Components.

A virtual machine to host affinity rule exists out of three components:

- Virtual machine DRS group
- ESX host DRS group
- Designation – "Must" affinity/anti-affinity or "Should" affinity/anti-affinity

Virtual machine DRS groups and ESX host DRS Group are quite self-explanatory so let's dive into the designations component straight away.

Designations

Two different types of VM-Host rules are available, a VM-Host affinity rule can either be a "must" rule or a "should" rule. The must-rule is a mandatory rule for HA, DRS and DPM; it confines or prevents the virtual machines to run on the ESX hosts specified in the ESX host DRS Group.

The "should" rule is a preferential rule for DRS and DPM and expresses a preference. DRS and DPM use their best effort to try to confine or prevent the virtual machines from running on the ESX host they are affinity to, but DRS and DPM can violate "should" rules if it compromises certain key operations. HA is not aware of preferential rules because DRS will not communicate these rules to HA. HA, DRS and DPM must take the mandatory rules into account when generating or executing operations. HA, DRS and DPM will never take any action that results in the violation of mandatory affinity rules. Because of this, mandatory rules place more constraints on VM mobility, making it more difficult for DRS to balance load and enforce resource allocation policies, HA and DPM operations are constrained as well, for example, mandatory rules will:

- Limit DRS in selecting hosts to load-balance the cluster
- Limit HA in selecting hosts to power up the virtual machines
- Limit DPM in selecting hosts to power down

Due to its limiting behavior, it is recommended to use mandatory rules sparingly and only for specific cases, such as licensing requirements. Preferential rules can be used to meet availability requirements such as separating virtual machines between blade centers.

DRS and mandatory rules

DRS takes mandatory rules into account when generating load-balance recommendations. If a rule is created and the current virtual machine placement is in violation with the rule, DRS will create a priority one recommendation (five stars) and executes the recommendation if DRS is set to fully automatic. DRS will not generate a recommendation that will violate the rule, it will not migrate virtual machines to or from an ESX server, even if it places the source ESX host into maintenance mode. VMotion will reject the operation as well if it detects that the operation is in violation of the mandatory rule.

If a reservation is set on the virtual machine, DRS takes both reservation and mandatory affinity rule into account. Both requirements must be satisfied during placement or power on. If DRS is unable to honor either one of the requirements the virtual machine is not powered on or migrated to the proposed destination host. For example, if a new rule is created and the current virtual machine placement is in violation of the rule, it can only migrate to a new host if the virtual machine memory reservation can be satisfied on the new host; if this is not possible, DRS will not generate the



Search

Popular articles

- [Resource pool memory reservations](#)
- [Sizing VMs and NUMA nodes](#)
- [Impact of Memory reservations](#)
- [NFS and IP-Hash load-balancing](#)
- [Left hand SAN-Lessons learned](#)
- [Windows 2008 disk alignment](#)
- [Increasing the Queue Depth](#)

About this blog

Welcome to Frankdenneman.nl, a blog about virtualization with a strong focus on storage technology and resource management.

Disclaimer

The views expressed anywhere on this site are strictly mine and not the opinions and views of VMware.

recommendation.

If a rule is created that conflict with another active, the older rule overrules the newer rule and DRS will disable the new rule.

As you can imagine that mandatory affinity rules can complicate troubleshooting in certain scenarios for example, why a virtual machine is not migrated from a highly utilized host to an alternative lightly utilized host in the cluster.

DPM

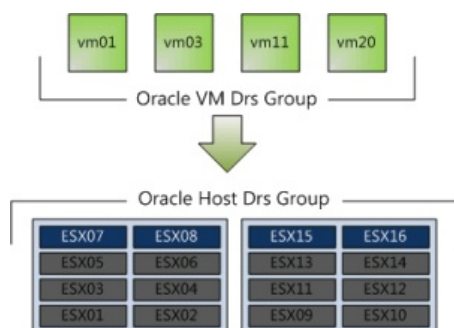
DPM does not place an ESX host into standby mode if it will violate the mandatory rule and will power-on ESX hosts if these are needed to meet the requirements of the mandatory rules.

High Availability

Due to the DRS-HA integration in vSphere 4.1, HA respects mandatory (must) rules. During an ESX host failure event, HA asks DRS to supply the list of hosts and places the virtual machines only on the compatible host, i.e. the host that are allowed by the mandatory rules. HA is unaware of the preferential (should) rules, so HA might unknowingly violate the rule during placement of virtual machines after an ESX failure, but the violation will be corrected by the next DRS invocation.

Let's take a look at a configuration which I think is going to be widely implemented soon, the Oracle Must affinity rule.

1. Place all Oracle virtual machines in a Cluster VM DRS group. (vm01, vm03, vm11, vm20)
2. Place all Oracle licensed ESX host in a Cluster Host DRS Group (ESX07, ESX08, ESX15, ESX16)
3. Select "Must run on Host in Group"



In this scenario, DRS never places, migrates, or recommend placement of a host-affined virtual machine on a host to which is not listed in the Cluster Host DRS Group (ESX01 – ESX06 & ESX09-ESX14). This means that DRS will never ever place the virtual machine on an unlicensed host, not for maintenance mode, not for DPM power saving and not after an ESX host failure event.

This virtual machine to host affinity rule make it possible to run oracle inside big clusters without having to license all the ESX host. I have been involved in a few projects where Oracle license was a constraint. Normally separate smaller clusters were deployed for Oracle database virtual machines, increasing both OPEX and CAPEX of the environment. These rules allows the Oracle virtual machines to run inside the cluster with other virtual machines without having to license all the ESX host inside the cluster. Hereby making the lives easier of both the architect and the administrator. vSphere 4.1, you gotta love it!



12 Responses to "VM to Hosts affinity rule"



Tomi Hakala says:
July 16th, 2010 at 10:39

Have you checked this with Oracle that they actually DO approve vSphere 4.1 host affinity rule for not licensing whole cluster? My experience with Oracle is that they always find a way to charge most from you, and until they publicly acknowledge that vSphere 4.1 host affinity is what they call "hard partitioning" I cannot rely on any other information 😞



Frank Denneman says:
July 16th, 2010 at 10:49

No, I haven't checked, but in theory you will need to license every ESX host in your vCenter datacenter, because that is the vMotion boundary. You are able to place a virtual machine on any ESX host inside your vCenter datacenter. But let's focus on VM-host affinity rules. If a host is in a VM-host must affinity rule, they are considered compatible hosts, all the others are tagged as incompatible hosts. DRS, DPM and HA are unable to place the virtual machines on incompatible hosts. Due to the incompatible host designation, the mandatory VM-Host is a feature what can be (undeniably) described as hard partitioning. You cannot place and run a VM on incompatible host.



doni says:
July 16th, 2010 at 10:59

Good read, thanks. Looks to me this would also be useful for those of us with limited budgets & EVC clusters who'd like to more-easily keep certain VMs on the more powerful CPUs.



Tomi Hakala says:
July 16th, 2010 at 11:21

Frank, thanks, there might be some hope in this.



Greg says:
July 16th, 2010 at 11:35

Also good for multi tenant clusters. This will solve a few of my issues from the past, another great new feature of vSphere 4.1



Nick says:
July 16th, 2010 at 17:59

Frank,

With this new feature, where does that put the use of stretched clusters(i.e. having one cluster stretched across multiple sites)? It was my understanding that this was previously unsupported as DRS/HA was not 'site aware'...now you can just restrict groups of VMs to hosts in one datacenter(or more if necessary) so they're not vMotioning across the WAN during the day.

If stretched clusters are now supported, where do you think that puts SRM?



Jeff Snively says:
July 16th, 2010 at 19:14

Thanks for the overview.

Say you have a 5 server cluster with 80% utilization. One server fails and HA restarts the guests, unaware of affinity rules. I can imagine a situation in which DRS would be unable to recover to your preferred affinity state.

latoga labs | vSphere 4.1 Hidden Gem: Host Affinity Rules says:
July 17th, 2010 at 02:03

[...] on providing a quick overview of the Host Affinity Rules in this post, but Frank Denneman already provided a great overview that I can't really improve on. Start with his post to understand the basics of the new [...]



latoga says:
July 17th, 2010 at 02:09

Great Post Frank! Was thinking of writing up the same thing, but you beat me to the punch. I've been talking to my enterprise clients for months now about how they could use Host Affinity Rules. I shared [a few of these use cases](#) over on my blog as well.

Disable DRS and VM-Host rules | frankdenneman.nl says:
July 22nd, 2010 at 06:49

[...] 4.1 introduces DRS VM-Host Affinity rules and offer two types of rules, Mandatory (must run on /must not run on) and preferential (should run [...])

vSphere 4.1 and its new CPU scheduler | P2V It! says:
July 22nd, 2010 at 12:11

[...] Frank Denneman: Load Based Teaming☐, DPM scheduled tasks☐ and VM to Hosts affinity rule☐ [...]

DRS-FT integration | frankdenneman.nl says:
July 22nd, 2010 at 15:00

[...] chassis if the design requires this. For more information about VM-Host affinity rules please visit this [...]

Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Submit Comment