

DRS Deepdive

Last week I mentioned which [metrics DRS used](#) for load balancing VMs across a cluster. Of course the obvious question was when the DRS Deepdive would be posted. I must admit I'm not an expert on this topic as like most of you I always took for granted that it worked out of the box. I can't remember that there ever was the need to troubleshoot DRS related problems, or better said I don't think I've ever seen an issue which was DRS related.

This article will focus on two primary DRS functions:

1. Load balancing VMs due to imbalanced Cluster
2. VMPlacement when booting

I will not be focusing on Resource Pools at all as I feel that there are already more than enough articles which explain these. The [Resource Management Guide](#) also contains a wealth of info on resource pools and this should be your starting place!

Load Balancing

First of all VMware DRS evaluates your cluster every [5 minutes](#). If there's an imbalance in load it will reorganize your cluster, with the help of VMotion, to create an evenly balanced cluster again. So how does it detect an imbalanced Cluster? First of all let's start with a screenshot:


VMware DRS	
Migration Automation Level:	Fully Automated
Power Management Automation Level:	Off
DRS Recommendations:	0
DRS Faults:	0
Migration Threshold:	Apply priority 3 or higher recommendations
Target host load standard deviation:	<= 0.163
Current host load standard deviation:	0.022 ( Load balanced)
View Resource Distribution Chart	

fig 1

There are three major elements here:

1. Migration Threshold
2. Target host load standard deviation
3. Current host load standard deviation

Keep in mind that when you change the "Migration Threshold" the value of the "Target host load standard deviation" will also change. In other words the Migration Threshold dictates how much the cluster can be "imbalanced". There also appears to be a direct relationship between the amount of hosts in a cluster and the "Target host load standard deviation". However, I haven't found any reference to support this observation. (Two host cluster with threshold set to three has a THLSD of [0.2](#), a three host cluster has a THLSD of [0.163](#).) As said every 5 minutes DRS will calculate the sum of the resource entitlements of all virtual machines on a single host and divides that number by the capacity of the host:

$$\text{sum(expected VM loads)} / (\text{capacity of host})$$

The result of all hosts will then be used to compute an average and the standard deviation. (Which effectively is the "Current host load standard deviation" you see in the screenshot([fig1](#)).) I'm not going to explain what a standard deviation is as it's explained extensively on [Wiki](#).

If the environment is imbalanced and the Current host load standard deviation exceeds the value of the "Target host load standard deviation" DRS will either recommend migrations or perform migrations depending on the chosen setting.

The question left is how does DRS decide which VM or set of VMs it will VMotion...

The following procedure is used to form a set of recommendations to correct the imbalanced cluster:

```
While (load imbalance metric > threshold) {
  move = GetBestMove();
  If no good migration is found:
    stop;
```

ABOUT YELLOW-BRICKS.COM

Welcome to Yellow-Bricks.com a blog about virtualization, with a strong focus on VMware, HA and vCloud solutions.

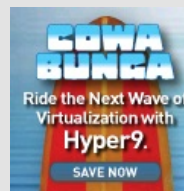
Some facts:

999 Blog Articles

1950 Twitter Followers

5503 RSS Subscribers

SPONSORS



ADVERTISEMENTS

VMware I/O Performance

VM to LUN optimization for your business critical apps
www.virtualinstruments.com

See the industry's best

I/O virtualization solutions Meet us at VMworld 2010, Booth 1413
www.virtensys.com

VM Capacity Management

VM Capacity Management is More Than Real Time Monitoring. Learn Why!
www.VKemel.com/LearnVMhy



RECOMMENDED



vSphere 4.0 Quick Start Guide

Shortcuts down the path of Virtualization

Buy it now just \$15.99.

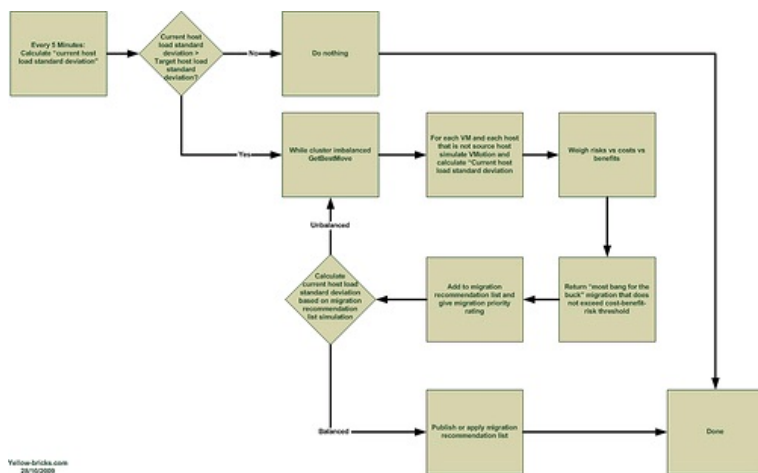
ADVERTISEMENTS



Watch this video on www.youtube.com

SQLserver on VMware

Use Veloxum to improve performance immediately, Free Trial



VM Placement

The placement of a VM when being powered on is as you know part of DRS. DRS analyzes the cluster based on the algorithm described in "Load Balancing". The question of course is for the VM which is being powered on what kind of values does DRS work with? Here's the catch, DRS assumes that 100% of the provisioned resources for this VM will be used. DRS does not take limits or reservations into account. Just like HA, DRS has got "admission control". If DRS can't guarantee the full 100% of the resources provisioned for this VM can be used it will VMotion VMs away so that it can power on this single VM. If however there are not enough resources available it will not power on this VM.

That's it for now... Like I said earlier, if you have more in-depth details feel free to chip in as this is a grey area for most people.

5 Responses to "DRS Deepdive"



chuckgman says:

Thursday, October 29, 2009 at 03:05

excellent. nice work.



Anthony says:

Monday, November 23, 2009 at 06:37

This is a fine piece of work! There are still administrators out there that are concerned about automating DRS. This article does a great job of addressing those concerns by describing the algorithms used to mitigate risk before VM workloads are re-balanced.



mnbng says:

Thursday, April 29, 2010 at 13:54

Thank you Duncan – a very useful document. However, I wonder if you (or readers of your blog) are able to clarify a couple of points for me:

I cannot find information on how the 'load' on each Host is calculated and therefore the deviation. Does it relate to the VMs' CPU or memory usage or both?

Also, in your last section 'VM Placement' you state that when a VM is powered on "DRS assumes that 100% of the provisioned resources for this VM will be used." Am I right in thinking that if the VM subsequently does not use all the resources that it has been provisioned, the next time DRS evaluates the Cluster it will calculate a different Host load (than the expected value DRS calculated when the VM requested being powered on) assuming all other VMs remain the same?

I am presuming that DRS works by comparing actual loads used by VMs once they are powered on and only uses the load from 100% of provisioned resources when calculating which Host to use when Powering on a VM. Is this correct?

Thanks



iwan says:

Sunday, August 1, 2010 at 11:29

Thanks Duncan. Always an in-depth article.

Just wondering if DRS in vSphere 4.1 now take into account HA into account. For example:
I have 20 very small VM and 2 very huge VM, and I only have 2 hosts.

Assuming all other factors are identical, will DRS spread the VM "nicely"? That means each host will have 1 big VM and 10 small VM?

Thanks from Singapore. And hopefully I get to meet you in VMworld SFO.
e1



Duncan Epping says:

Sunday, August 1, 2010 at 11:57

<http://www.yellow-bricks.com/2010/07/14/vsphere-4-1-vmware-ha-new-maximums-and-drs-integration-will-make-our-life-easier/>

Leave a Reply

Name (required)

Mail (will not be published) (required)

Website

Submit Comment

Copyright © Yellow Bricks - Building blocks for virtualization...

Powered by WordPress | Modern Style theme by FlexiThemes