

# 16 of the biggest HA and DRS configuration mistakes

01<sup>Sep</sup>

Justin Vashisht [VMWare](#)

@concentratdgreg (Greg Shields), is my new favorite technology speaker! His session on avoiding 16 of the biggest HA and DRS mistakes on vSphere was incredible. Greg has a way of making the audience stay focused and engaged throughout the entire session and keeping things super fun. His analogies throughout the session helped clear up some complicated topics. I recorded the session and you can listen to it below. The hour will fly by! Oh and the "single leg table" scenario was awesome when Greg does that bounce lol.

<http://www.zshare.net/download/94285600e73c2b5b/>

For people who want a quick read, Derek Seaman did a [nice writeup](#) on the session. No sense in me redoing it.

- 1. Not planning for hardware change.** The solution is to enable EVC mode on clusters, and try to use hardware that has very similar processors. Do not mix Intel and AMD processors. And even within manufactures, newer processors will support additional instructions not available on older models. Pay attention to the CPUs you buy.
- 2. Not planning for svMotion (storage vMotion).** Snapshots are evil, and you should NOT use them except in rare situations. Make sure your VMDKs are in persistent mode or use RDMs. The servers must see both the source and target datastores, and the cluster must have enough resources to briefly have two copies of the VM concurrently running.
- 3. Not enough cluster hosts.** Plan for adequate cluster resources and build-in a reserve factor. Typically one full server's worth of resources. Solution is to use an admission control policy and set the host failures tolerate to "1".
- 4. Setting the host failures tolerate to "1".** Not all VMs are tier-1 and deserve the same restart priority. Setting aside a full host can be wasteful. Use the percent of cluster resource option and configure a percentage that is less than a single host's contribution to the cluster based on the number of hosts. For example, in a four node cluster each server contributes 25% so set the percentage to something like 20 or 15%.
- 5. Not prioritizing VM restarts.** If you use the suggestion in mistake #4, you must properly configure the VM restart priority since you won't be reserving cluster resources to restart ALL of your VMs. Set your normal VMs to low restart priority then elevate special VMs to medium or high as needed.
- 6. Disabling admission control.** Bad idea! Never, ever, ever do this! Enable the do not power on if inadequate cluster resources are available option.
- 7. Not updating the % policy.** As you add more hosts to a cluster you need to recalculate the host failures percent number, otherwise the system will get out of wack.
- 8. Buying dissimilar servers.** The hosts failure cluster tolerates option bases its calculation on the biggest server in the cluster. If you have six servers with 96GB of RAM then add a server with 384GB of RAM to the cluster, it will really throw out of wack the calculations and you will be left with a lot of unused resources.
- 9. Host isolation response.** This is a confusing subject for many and prior to v5.0, contained some bugs or behaved in a way that customers did not always expect. In many environments you can configure the response to shutdown the guest VMs, but on a per-VM basis change the settings if you have critical apps that you want to ensure don't go down on accident. In v5.0 the datastore heartbeat feature is a welcomed change and is only used if the management network goes down. The servers in the cluster need one datastore in common.
- 10. Overdoing reservations, limits and affinities.** Use shares over reservations and limit the use of affinities (or anti-affinities). These restrictions impact the DRS calculations and can impact performance. Use sparingly!
- 11. Doing memory limits at all.** Don't ever do this, ever, ever! Limit memory usage as close to the app as possible. For example, you can configure SQL to limit the amount of memory it will use within the guest VM.
- 12. Thinking you are smarter than DRS.** No human can calculate all of the variables and come up with the right answer. Let the software do its job.

13. **Not understanding the DRS rebalancing equation.** Far too complex to repeat here, so do some googling for this one.

14. **Being too liberal.** Migrations take resources, be they network bandwidth or CPU time. Don't have DRS continually moving workloads between servers. Configure thresholds to do sensible migrations when resources really are out of balance. vMotion was cool 5 years ago, but no need to have DRS continually move workloads just to be cool.

15. **Too many cluster hosts.** Although the technical limit is 32 hosts per cluster, the sweet spot is 16-24 hosts. Any larger and the calculations DRS does every five minutes become very complex and consume more and more resources.

16. **Creating big VMs.** This has new meaning with the 5.0 vTax licensing scheme. Assign the right amount of memory and vCPUs to a VM. Don't be too liberal. Right size the VM, don't supersize them.