

DRS Interview Questions

Q1: what is the default invocation period for DRS. Can we change this. If yes then how?

Ans: The default invocation period is 300 seconds (5 minutes). But this can be changed via the configuration file vpxd.cfg. We have to change the value of <pollperiodsec> as shown below:

```
<config>
```

```
<drm>
```

```
<pollperiodsec>
```

```
300
```

```
</pollperiodsec>
```

```
</drm>
```

```
</config>
```

Just change the value 300 to a custom value defined by you. The range of supported value is 60 secs to 3600 secs.

Q2: What is the role of VPXA in DRS?

Ans: VPXA is the vCenter agent that runs inside Esxi hosts and it enables a 2 way communication between Esxi hosts and vCenter Server. VPXA is responsible for:

- 1) Keeping the status of Esxi and VM's in sync
- 2) It sends info to vCenter server when a VM's power state is changed or a VM is vMotioned from one host to other.

DRS uses this information which is presented by Esxi hosts to vCenter server for calculating the load balance and proposed migrations in case of cluster imbalance.

Q3: Will DRS work if vCenter server is down? If no then explain why DRS is dependent on vCenter server.

Ans: No DRS will not work if vCenter Server is down.

DRS depends upon vCenter server for information like current power state of virtual machines, change in power state of any VM, number of datastores to which Esxi hosts are connected and the memory and cpu configurations of a VM.

DRS will use all these information while calculating the load on the cluster and proposing migration recommendations when a cluster needs to be balanced.

Q4: How DRS calculates there are imbalances in cluster? What are the things that DRS takes into account for determining this?

Ans: To calculate the cluster imbalance, DRS compares the Current Hosts Load Standard Deviation (CHLSD) to Target Hosts Load Standard Deviation (THLSD) and if

$CHLSD < THLSD$

Cluster is considered as imbalanced.

CHLSD Calculation

DRS computes Normalized Entitlement (NE) of each Esxi hosts and the standard deviation associated with it. NE is nothing but calculation of how much resources are currently utilized out of total resources. NE is calculated by summing up dynamic entitlements (usage) of all VM's that are running on an Esxi host and diving this by Esxi host capacity.

So, $NE = \text{Dynamic Usage of all VM's} / \text{Total host capacity}$

THLSD Calculation

THLSD is derived from DRS migration threshold which is defined at the time of configuring DRS. Each threshold level sets different imbalance tolerance margin. The aggressive threshold sets a tight margin allowing for little imbalance, while conservative thresholds tolerates bigger imbalances.

Q6: What is DRS cost-benefit-risk approach?

Ans: DRS uses a cost benefit risk approach to load balance a cluster. Before presenting the migration recommendations for a VM to load balance a cluster, DRS calculates 3 things:

- 1) Cost: What will be the cost of migrating a VM from source to destination host? Cost here refers to the CPU and memory.

When a vMotion process is invoked on a VM, it reserves 30% of the CPU core (For 1 GB NIC) of that VM and 100% CPU core (For 10 GB NIC). This reservation is created on both source and destination host. These reserved resources can't be allocated to any other VM while vMotion is in progress. This can put some pressure on an ESXi host when it is heavily loaded.

- 2) Benefits: What will be the resource benefit that an ESXi host will get after migrating a virtual machine to another ESXi host. If after the migration CHLSD value of the ESXi host comes down then DRS will consider that migration as benefit.
- 3) Risk: Risk accounts for possibility of irregular loads. Suppose a VM has inconsistent and spiky demands of resources, then migrating such VM's is not a good move because maybe at the time of migration VM resource demand was low but after completion of migration, VM's resource demands suddenly increased. In this case it

can cause the increase in destination Esxi hosts CHLSD and again DRS has to perform migration of that VM to bring down CHLSD of the Esxi host where that VM was migrated.

Note: DRS recommends migrations if *Benefit obtained due to a migration < cost associated with that migration.*

Q8: What are the factors that affect DRS recommendations?

Ans: Following are the factors which affect the DRS recommendations:

1- VM size and Initial Placement: When a new VM is created or a VM is powered on, DRS selects a host where this VM should be initially placed. DRS prefers the registered host as long as placement of that VM on this host will not cause cluster imbalance.

During placement of such VM's DRS uses a worst case scenario because it doesn't have historical data for that VM. DRS assumes both CPU and memory demand of this VM is equal to its configured size.

Oversized VM's can temporarily cause cluster imbalance and can cause unnecessary migrations of active VM's.

2- PollPeriodSec (Length of DRS Invocation): The default value of PollPeriodSec is 300 seconds. Range of PollPeriodSec is 60 sec to 3600 sec. shortening this period will cause increase in vCenter overhead as cluster imbalance will be calculated frequently.

Increasing this value decreases the frequency of cluster balance calculation and can leave your cluster imbalanced for longer period of times but allows for larger number of vMotions due to long invocation interval.

3- Simultaneous vMotion: vSphere 5.1 allows 8 concurrent vMotions on a single host with 10GbE capabilities. For 1GbE, 4 concurrent vMotion can takes place. Also multi-NIC vMotion is supported in vSphere 5.1 so multiple active NICs and their combined

bandwidth can be used for migration of a VM. In such environment VM's will be migrated quickly and cluster can be balanced in less time.

4- Estimated Total Migration Time: The migration time depends on variables like source and destination host load, active memory usage of VM, link speed and available bandwidth+ latency of the physical network used by vMotion Portgroup.

Q9: what are the use cases for VM-VM affinity rules and VM-VM anti affinity rules?

Ans: VM-VM affinity: This is useful when you require that 2 of your VM's should always run together on an Esxi host. For E.g. Keeping front-end and back-end server of an application on same ESXi host to reduce network latency between the 2 VM's.

Another use case will be running together same types of VM's which are having same type of applications so as to get max benefits of transparent page sharing (TPS)

VM-VM anti-affinity: This is useful when you don't want that 2 of your VM should run together. Keeping servers providing same kind of services on different host will provide resiliency

For e.g. You will not want your DC and ADC run together on same Esxi host because if that Esxi host goes down it can severely impact your environment as both DC and ADC server has gone down together.

Another use case will be running web-server farms or clustered DB-servers in a virtualized environment.

Also, keeping away 2 VM's from each other which are very resource intensive to stop them from monopolizing resource usage.

Q10: What are the use cases of VM-Host affinity rules and VM-Host anti-affinity rules?

Ans: VM-Host affinity: This is useful when you want that your VM should run on a particular Esxi host only.

For e.g. running and oracle DB server which has socket based license. If your environment is having heterogeneous hosts than migrating such VM to a host which has different CPU configuration can violates your license and can cause trouble.

VM-Host anti affinity: This is useful when you want that a particular VM should not run on some particular Esxi hosts. For e.g. your environment has heterogeneous hosts and all the hosts don't have Numa architecture and you want to get benefits of the vNuma inside your VM. In this case you would want your VM to run only on those servers which supports Numa.

Q11: Can DRS override “preferential or should rules”. If yes then how and if no then why?

Ans: Yes DRS can override should rules. When rules are configured inside DRS then DRS creates a rule list and provide migrations recommendations in accordance with the rules defined in rule list. But if the cluster imbalance cannot be solved even after running these migrations then DRS drops the rule list and re-run the load balance algorithm and those migrations also which can break the should rules, in order to load balance a cluster in a better way.

Q12: CAN HA, DRS and DPM override “must or mandatory rules”?

Ans: No HA, DRS and DPM can't override must rules.

Q13: What impact does must rule places on DRS, HA and DPM operations?

Ans: If a migration will cause violation of must rule then that migration will be cancelled by DRS.

IF DPM is trying to put a host in sleep mode for power saving but migration of VM's running on this host can cause a must rule violation, it will prevent DPM to put that host in sleep mode.

If HA is trying to restart VM's after a host failure but if restart of some VM's on a particular host can cause must rule violation then HA will either restart those VM's on some different host or could not restart them at all if no suitable host is available for failover.

Q14: If we have configured some rules (affinity or anti-affinity) in DRS the will those rules work if we disable DRS on a cluster?

Ans: Yes rules will be in affect even if we disable DRS without deleting the rules first.

Q15: If on a VM, VM-Host affinity should rule is configured then can we migrate that VM on an Esxi host that is not part of the DRS cluster?

Ans: No a VM can't be migrated to an Esxi host that is not part of the DRS cluster.

Q16: What are the best practices for disabling DRS?

Ans: Before disabling DRS it is recommended to delete all affinity and anti-affinity rules and then proceed. Because if rules are not deleted and DRS is disabled, rules will be still in affect and can affect cluster operations.

Q17: What are the limitation that are put by DRS mandatory or must rules on a cluster?

Ans: If a mandatory rule is configured on a cluster then it can put following limitations:

- 1) Limit DRS to select hosts for load balancing
- 2) Limit HA to select host for failover
- 3) Limit DPM to select host to power off

- 4) It can affect ability of DRS to defragment the cluster resources. At the time of failover HA can seek assistance for DRS and can ask to defragment resources if a single host is not able to provide adequate resources for failover.

Q18: If a new DRS rule is created but that rule is conflicting with any existing rule then which rule DRS will respect, old rule or new rule while performing DRS actions?

Ans: If a new rule is conflicting with an old rule then new rule will be disabled automatically. DRS will prefer respecting the old rule.

Q19: How many automation levels are there for a VM in respect to DRS? Can VM automation level override cluster automation level?

Ans: VM automation level can override cluster automation level. From a VM prospective there are 5 automation levels. These are:

- 1) Fully Automated: Load balance and Initial placement will be done by DRS automatically
- 2) Partially Automated: Load balance of the VM will be done manually but initial placement will be done automatically
- 3) Manual: VM migration as part of Load balancing and initial placement will be both manual. DRS will only generate recommendation for that VM and administrator has to manually approve this recommendation.
- 4) Default: VM will inherit the DRS automation level as defined at the cluster level.
- 5) Disabled: DRS will not perform any actions on that VM.