How the Service Works

1 Overview

This service estimates the impact of one or more host failures in a cluster in which vSphere HA and vSphere DRS are enabled. For each simulated failure, the service reports

- The VMs that could not be restarted due to insufficient capacity to meet VM reservations, and
- For each VM that could be restarted and those that remained powered during the failure, an estimate of how much of its historical CPU and memory demand could be satisfied with the reduced cluster capacity.

Using this output, you can assess whether the estimated resource allocation for each VM would be acceptable should such a failure occur in your actual environment. If the estimated allocations would not be sufficient, VMware recommends removing any unnecessary dependencies in the cluster and/or adding more capacity, and then re-assessing the situation using the tool. For more specific suggestions, see the section below entitled "Taking Action Based on the Results".

2 Accuracy of the Results

The resource allocations reported by the service are computed using the DRS algorithm and are accurate. For example, if you simulate the failure of a specific host and then that host failed under exactly the same resource conditions as you simulated, the resource-allocation estimates would be very similar to what the VMs would experience in an actual failure.

However, keep in mind that the simulation uses a snapshot of the cluster, and at the time of a genuine failure, the cluster conditions are most likely different. So, we recommend reserving some headroom in your cluster to account for increases in CPU or memory demand, changes in VM reservations, etc.

The results reported by the service are for a cluster with both HA and DRS enabled. After determining the VMs that could be restarted following the simulated failure, the simulator invokes the DRS algorithm to load balance the cluster. DRS load balancing is an important post-recovery step because HA does not take into account VM resource demands when placing VMs, and hence, could overcommit the resources on hosts. VMware recommends enabling DRS in a HA enabled cluster to minimize the performance impact of HA restarts.

Finally, the list of VMs that would not be restarted after a simulated failure may not be representative of the set that could not be restarted in an actual failure. The service does not consider HA restart priorities and so it may report that a high priority VM could not be restarted while a low priority one could, whereas in practice, the opposite would have occurred.

3 How To Use

To use the service,

1. Upload a DRS dump file covering a period of interest, which contains a snapshot of your cluster inventory and resource usage for the last hour. Note that the service only supports DRS dump files recorded by vSphere 5.0, 5.1, and 5.5.

2. Specify the failure scenario you wish to simulate using the service's UI
3. View and analyze the results.

The next three subsections discusses these steps in more detail.

3.1 Choosing a DRS Dump File

We recommend using a dump file that was recorded by DRS during a representative heavy load condition in the cluster of interest. Because you can't anticipate when a failure will occur, it is best to simulate a failure using data collected during a heavy-load business critical period. In this way, you can assess your ability to tolerate a failure at the worst possible time. DRS generates a dump files every time it runs to load balance the cluster. By default, the load balancing runs every 5 minutes.

Note: the service only supports DRS dump files recorded by vSphere 5.0, 5.1, and 5.5.Locating the dump file of interest for a particular cluster is a bit tricky due to how DRS saves the dump files for each cluster. VMware plans to make this process easier in future.

We recommend the following process for retrieving the DRS dump file of interest

1 Record the name of the cluster and the datacenter it is in
2 Determine the managed object ID (MoID) for that cluster
3 Locate the DRS dump file for this cluster
4 Choose a representative file

Step 1 – Record the name of the cluster and the datacenter it is in

Using the UI, identify the cluster of interest, and the vCenter Server datacenter in which it is located, and record this information.

Step 2 - Determine the managed object ID (MoID) for that cluster

To find the MoID for the cluster of interest, we recommend that you use the vCenter Server's managed object browser (MOB). Using the MOB, you will navigate to an object that lists cluster names and their MoIDs. The MOB can be accessed using any web browser. The exact navigation steps are as follows:

a) Open the URL https://<server IP address or FQDN>/mob/?moid=group-d1. For example, https://myVC.vmware.com/mob/?moid=group-d1. This URL will retrieve the vCenter Server top-level inventory folder. You will need to log in using a username that has system view privileges.
b) Look on this page for a property "childEntity". It lists the datacenters in the vCenter Server inventory. Click on the datacenter that contains the cluster of interest. For example, referring to the screen shot shown below, if the cluster is in the "Test Datacenter", then you should click on "datacenter-11218".

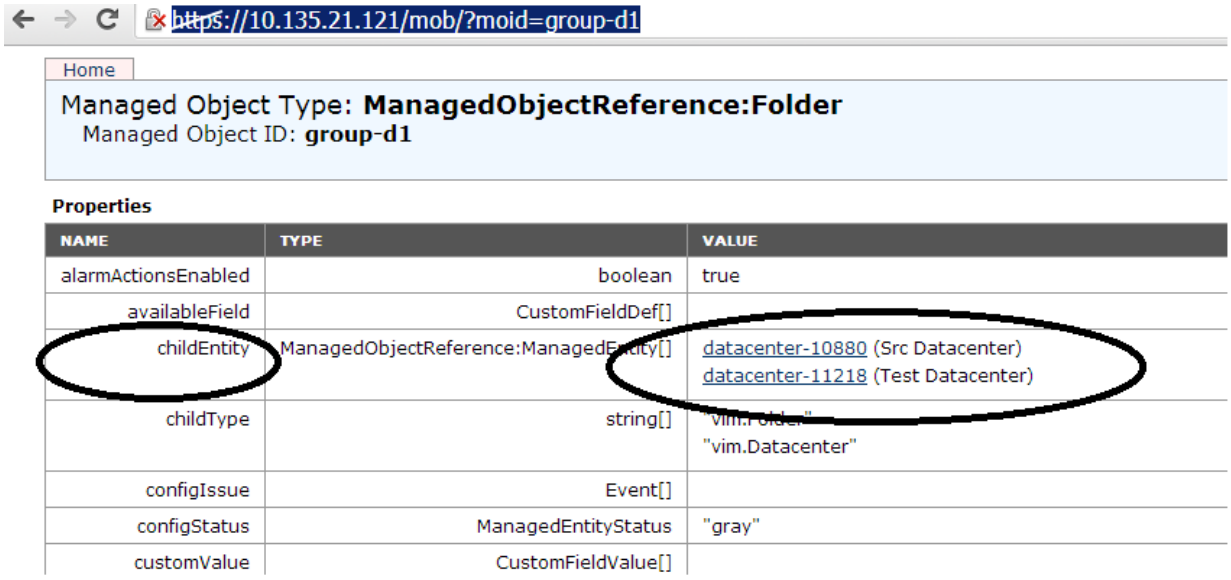

c) Clicking on the datacenter link will take you to a page listing the contents of the corresponding Datacenter object. Find the property "hostFolder" and click on its link. In the example below, click on the link "group-h11220".

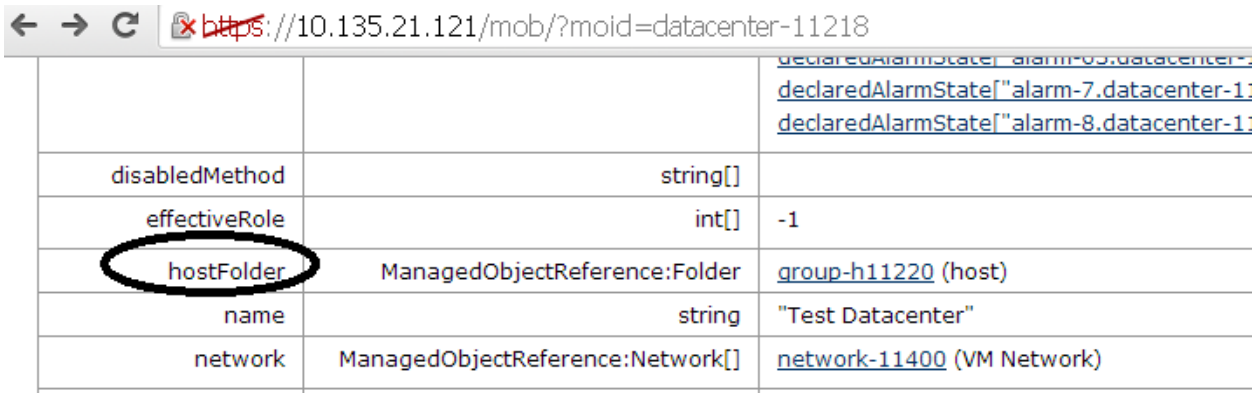

d) The next page will list the clusters within the datacenter. Record the MoID of the cluster of interest. In this example, the MoID for cluster HaCluster2 is domain-c11428. Cluster MoIDs all have the form **domain-cX** where X is a number.

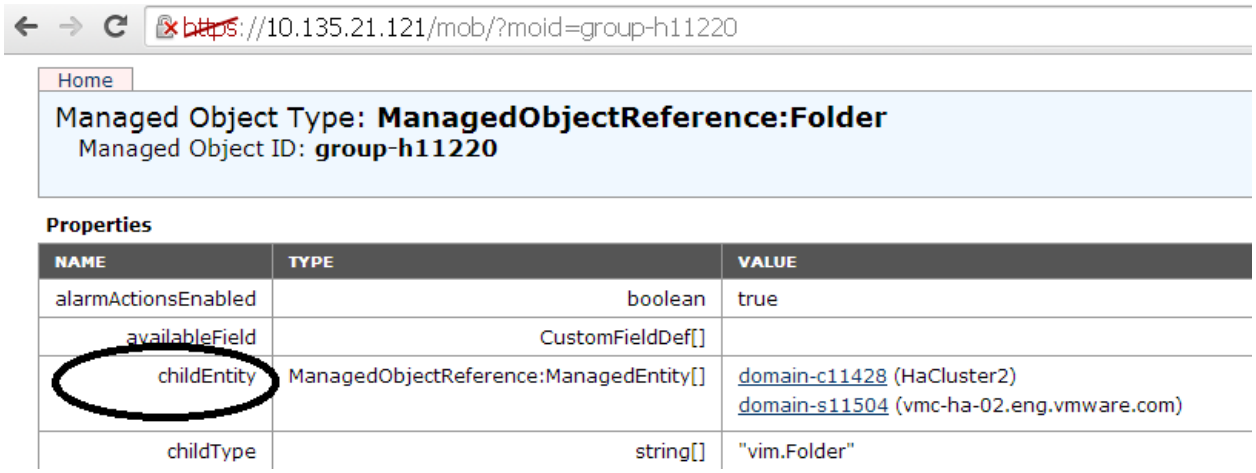

Step 3 Locate the DRS dump file directory for this cluster

DRS stores the dump files for all clusters within a directory called drmdump. Please see KB 1021804 for the location of this directory on the file system where the vCenter Server log files are stored.

Within the drmdump directory, there is a subdirectory for each cluster that DRS is managing. These are named using the cluster MoIDs. Change to the directory that corresponds to your target cluster.

```
/var/log/vmware/vpxd/drmdump/domain-c11428:ls
1667797000-proposeActions.dump.gz  1671397000-proposeActions.dump.gz  1674997000-proposeActions.dump.gz
1668097000-proposeActions.dump.gz  1671697000-proposeActions.dump.gz  1675297000-proposeActions.dump.gz
1668397000-proposeActions.dump.gz  1671997000-proposeActions.dump.gz  1675597000-proposeActions.dump.gz
1668697000-proposeActions.dump.gz  1672297000-proposeActions.dump.gz  1675897000-proposeActions.dump.gz
1668997000-proposeActions.dump.gz  1672597000-proposeActions.dump.gz  1676197000-proposeActions.dump.gz
1669297000-proposeActions.dump.gz  1672897000-proposeActions.dump.gz  1676497000-proposeActions.dump.gz
1669597000-proposeActions.dump.gz  1673197000-proposeActions.dump.gz  1676797000-proposeActions.dump.gz
1669897000-proposeActions.dump.gz  1673497000-proposeActions.dump.gz  1677097000-proposeActions.dump.gz
1670197000-proposeActions.dump.gz  1673797000-proposeActions.dump.gz  1677397000-proposeActions.dump.gz
1670497000-proposeActions.dump.gz  1674097000-proposeActions.dump.gz  1677697000-proposeActions.dump.gz
1670797000-proposeActions.dump.gz  1674397000-proposeActions.dump.gz  1677997000-proposeActions.dump.gz
1671097000-proposeActions.dump.gz  1674697000-proposeActions.dump.gz  1678297000-proposeActions.dump.gz
```

Step 4 - Choose a representative file

Select a “propose actions” dump file for the time range of interest. The service only works with dump files with names of the form X-proposeActions.dump, where X is a number. For example, a file named 1166797000-proposeActions.dump could be used. Use the file-system provided timestamps to select a file for a given time. Finally, use the service’s UI to upload this file.

3.2 Select a Failure Scenario

After the DRS dumpfile is uploaded, you need to specify the type of failure you wish to simulate. The service offers two choices – simulate the failure of any one host, or simulate the failure of one or more specific hosts. The first choice can be used to estimate the worst-case performance of any given VM if any one host failed in the cluster. In practice, all VMs would not experience the worst-case values at once since the worst-case

estimates could correspond to the failure of different hosts. To estimate the impact of a specific host or set of hosts failing, use the second option.

3.3 Interpreting the Results

After simulating a failure, the service reports a summary for all VMs in the cluster, and details on each VM. These reports are based on two key metrics.

1 **Overall resource allocation reduction percentage**: to summarize the per-VM data, the service computes, for each VM, a value summarizing the overall resource allocation reduction for the VM. This value is calculated by first computing the resource allocation of each VM for each simulated resource (CPU and memory) before the simulated failure, and then after the simulated failure. Then, the before-failure and after-failure values are combined into a single number for each case, and finally, the ratio of after-failure to before failure is computed.

   As a simple example, suppose a VM had a memory demand of >= 500 MB and that its memory allocation before the failure was 500 MB. Further, it had no CPU demand. If, after the simulated failure, its memory allocation was 250 MB, then its memory allocation reduction percentage is 50% (250 MB/ 500 MB).

   See the section below for a sample report.

2 **Degree of performance impact**: to call attention to the VMs that were significantly impacted by the simulated failure, the failure groups VMs by their resource allocation reduction percentage. VMs with a percentage above a threshold are considered as being minimally affected.

   The default value is 60%. So, in the example mentioned above, the VM in question would be marked as impacted by the failure. This threshold value can be changed on the VM details page. Select the VMs for which you wish to change the threshold, and click “edit threshold”.

Example Reports

The figure below presents a sample per-VM report. Column A reports the change in the CPU allocation of each VM after the simulated host failure. Column B reports the overall resource allocation reduction percentage and visually. The header “C” reports the number of VMs in the cluster, the number that failed to restart, the number with resource reduction percentages below the threshold, and the number with percentages

above the threshold.

The figure below presents the summary view for the same experiment. Section “A” in the summary displays the histogram of reduction percentage for all VMs in the simulated cluster.

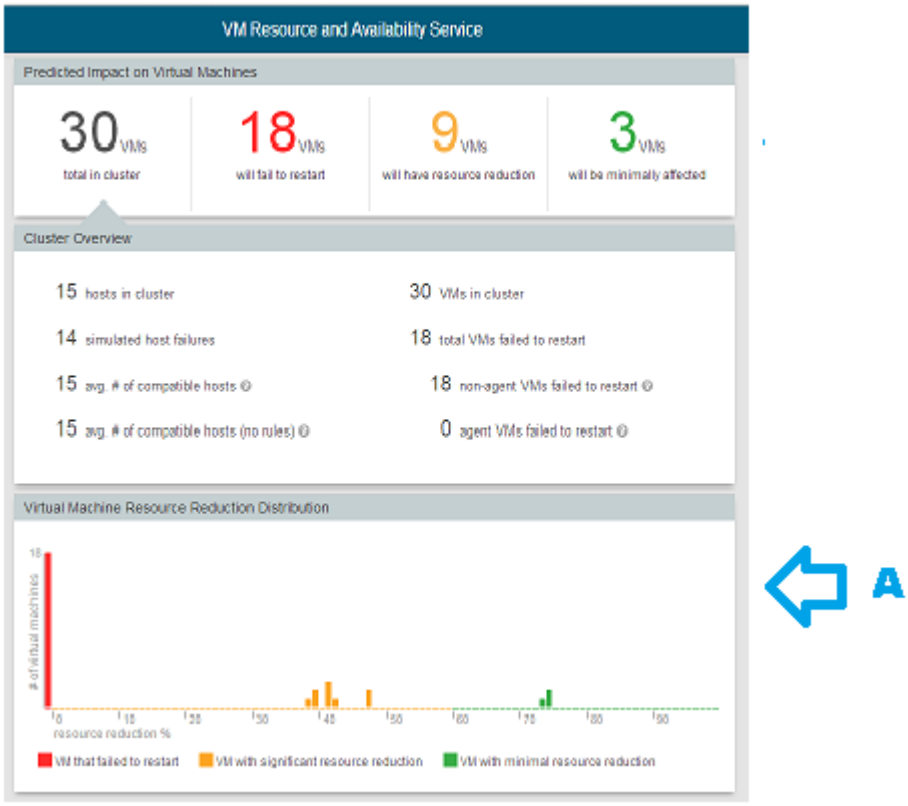

A

3.4 Taking Action Based on the Results

The service allows you to estimate the resource allocations for each VM after a simulated failure. If the reductions in resource allocation are not acceptable or some VMs were not restarted, you can mitigate this situation through the following steps.

1 If some VMs could not be restarted, configure HA to reserve additional capacity in case of failure. See the availability guide (LINK) for information on the specific policies.

2 Maximize the number of hosts on which each VM can run on for as many VMs as possible. Constraining a VM to run on a subset of the hosts limits impedes HA's and DRS's ability to restart the VM and move VMs around to take maximum advantage of the cluster capacity that remains after a failure. A number of settings can limit the hosts on which a VM runs. For example, you have defined required VM to host rules, or the network a VM depends on is not available on all hosts. Limit the use of such constraints as much as possible.

3 Increase cluster capacity by moving some VMs to another cluster, adding more resources (e.g., memory) to the cluster hosts, or adding additional hosts to the cluster.

4 Use HA and Resource settings to ensure your critical VMs get the resources they need. Specifically, if a VM could not restart, increase its HA restart priority or decrease the HA restart priority of less important VMs. These changes will cause HA to allocate available capacity to the critical VMs before considering others. If a VM restarted but obtained insufficient resources, consider changing VM share values and/or

VM reservations. VMware recommends using reservations if the preferred resource allocation could not be achieved with share value changes. While reservations provide guaranteed resource allocations, they can prevent this VM or other VMs from being restarted at all.

4 Final Considerations

The simulation results assume that DRS is enabled in the target HA cluster. If, in practice, you don't have DRS enabled, consider doing so. DRS is an important component in post-failure recovery. DRS provides two functions: if HA cannot place a VM due to insufficient unreserved capacity on any one host, DRS will try to make capacity available for the VM; and after HA has restarted the VMs, DRS will load balance the cluster to minimize resource contention and hence maximize VM performance.

To assess the impact of failures on other resources aside from CPU and memory, consider simulating the impact of a failure by placing a representative number of hosts in maintenance mode. Of course, doing so can be very disruptive to the cluster VMs, so this approach should be used with care.