



Proactive High Availability with OMIVV

Tech Note by:

Pandiyar Varadharajan
Jonas Werner
Tad Walsh

SUMMARY

The newly introduced VMware Proactive HA functionality, in combination with OpenManage Integration for VMware vCenter (OMIVV), can detect potential hardware issues and act upon proactively – before they result in actual downtime.

Proactive HA and OMIVV enable policy-driven preventive action in case an issue in the underlying hardware layer is detected.

Via the policy settings, automated actions can be taken to move VMs off an affected host or prevent any new VMs to be started there.

The ability to do this automatically, without human intervention, helps to improve uptime for services deployed in the virtual environment while simultaneously enabling IT Administrators to be more efficient and productive.

Virtualization has come to permeate all layers of the modern IT infrastructure, from virtual machines to virtualized storage and networking. The benefits of virtualization, in the shape of increased efficiency, faster service delivery, cost savings, ease of management, and ensuring higher levels of availability, have been well-documented. Customers acknowledge these benefits but also point out that as increasing number of applications, services and business-critical systems are virtualized, they come to rely upon the underlying hardware layer to ensure redundancy in case of failure. Moreover, the more powerful the server, the larger the number of virtual machines supported and the greater the potential impact to the loss of a given system. Thus, along with the benefits of virtualization comes the risk of a hardware failure bringing down or delaying a significant number of a company's services.

The solution to this is closer interaction between the underlying hardware and the virtualization platform, combined with a way to communicate potential hardware issues proactively, before they become fatal. With the OpenManage Integration for VMware vCenter (OMIVV) version 4.0.1 and the new VMware Proactive HA provider (vCenter version 6.5), Dell EMC have provided another layer of reliability to resolve this challenge.

From reactive to proactive

VMware's High Availability (HA) functionality to protect application uptime has been available on most virtualized platforms for years, and is implemented pervasively in customer IT environments. Also available for some time, Dell EMC PowerEdge servers have several innovations, such as the redundant SD card system for hosting the ESXi hypervisor, and Fault-Resilient memory to provide an extra layer of memory protection for the hypervisor itself, that help to ensure uptime of the hypervisor and the workloads that it supports.

Customers indicate however, that VMware HA is reactive: It responds only once a host has completely failed and the consequences are being felt: HA ensures that virtual machines will be rebooted, but the impact of a non-graceful shutdown and the ensuing wait-time for them to reboot after a failure can have potentially serious ramifications. To overcome this exposure, the newly introduced VMware Proactive HA functionality, in combination with the OMIVV virtual appliance, can detect potential hardware issues and act upon them proactively – before they result in actual downtime.

Policy-driven preventive action

Proactive HA allows for policy-driven preventive action in case an issue in the underlying hardware layer is detected. Via the policy settings, automated actions can be taken to move VMs off an affected host or prevent any new VMs to be started there. The ability to do this automatically, without human intervention, helps to improve uptime for services deployed in the virtual environment while simultaneously enabling IT Administrators to be more efficient and productive.

VMware vSphere 6.5 with Proactive HA

With the release of vSphere 6.5, one of the new features introduced is **Proactive HA** functionality, enabling hardware health information to be communicated to VMware vCenter. Proactive HA allows a customer to set policies on what to do when a server state has degraded, including automated actions such as putting the degraded host into a specific state (e.g. move running workloads to other hosts in the cluster) or simply keep new virtual machines from being moved to the impacted host. This is actually a feature of the Distributed Resource Scheduler (DRS) feature in vCenter, which springs into action prior to the need to activate the standard HA functionality.

Dell EMC OpenManage Integration for VMware vCenter (OMIVV)

OMIVV is a Dell EMC plugin to VMware vCenter to bridge the gap between VMware virtualization platform and the underlying Dell EMC servers. It is easily deployed as a virtual appliance on an existing VMware cluster and hooks into the Dell EMC PowerEdge server iDRAC as well as the ESXi hypervisor deployed on each host.

With the introduction of VMware vSphere 6.5 and OMIVV 4.0 it became possible to leverage Proactive HA functionality to communicate hardware state information to IT Admins. The communications indicate Moderately Degraded (warning) and Severely Degraded (critical) hardware states, giving IT Admins deeper understanding of the health and behavior of their systems and helping to ensure increased uptime.

Proactive HA compared to normal HA functionality

In a normal clustered VMware environment, the HA functionality springs to action when a failure of a host occurs. Any virtual machines (VMs) that were running on the affected host will fail simultaneously as the host becomes unavailable, but will be automatically restarted on the remaining hosts in the cluster. Since this results in a non-graceful shutdown of the VMs as well as incurring a waiting time while they boot, this is a disruptive event which in worst case can result in data loss.

As the name suggests, Proactive HA takes action prior to failure of a host due to hardware outage. Downtime and potential data loss can potentially be prevented thanks to Proactive HA's ability to move VMs off the affected host and onto healthy servers without them having to shut down and restart. This is depicted in Figure 1 below, which illustrates a Proactive HA event flow in the case of a power supply unit (PSU) failure:

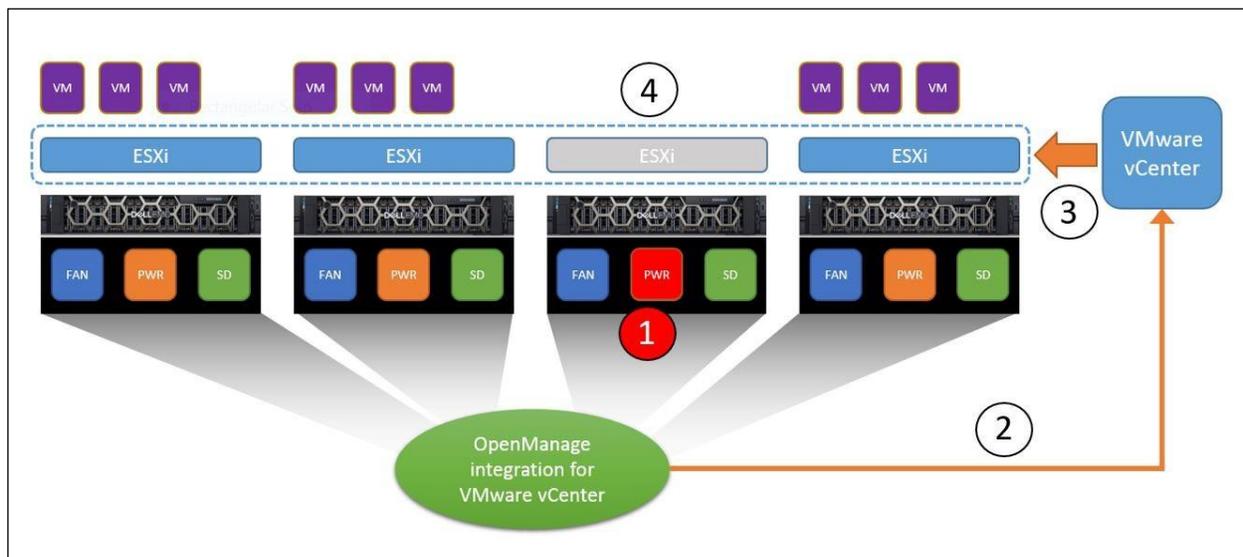


Figure 1: Proactive HA event flow in the case of PSU failure

The sequence of events depicted in Figure 1 above is as follows:

1. A server experiences a Power Supply Unit (PSU) failure.
2. The OpenManage Integration for VMware vCenter (OMIVV) plugin detects the change in server health status and informs VMware vCenter, flagging the server as Moderately Degraded (Warning) or Severely Degraded (Critical), depending on the alert settings. For this scenario, assume that the condition is Severely Degraded (Critical).

3. VMware vCenter performs the pre-set Proactive HA action associated with a Severely Degraded (Critical) health problem. In this case the action is to enable Maintenance mode.
4. Maintenance mode is enabled on the affected server and, if automated response is enabled, the Distributed Resource Scheduler (DRS) feature in vCenter will automatically migrate any VMs to healthy servers in the same cluster.

Actions and responses to health status changes

Actions will be taken depending on the type and severity of the status reported to VMware vCenter by OMIVV. Two basic types of actions are available and they can be implemented in three different ways:

- The first and most basic action is to enter **Maintenance** mode and automatically have all VM's on an affected host be live-migrated to other hosts within the cluster. If a fully-automatic DRS is configured on the cluster, this will happen without administrator intervention.
- The second option is called **Quarantine** and was introduced as a new status type in vSphere version 6.5 together with Proactive HA. In quarantine status, vSphere will prevent additional VMs from being scheduled to boot on or be migrated to a particular server. VMs currently running on the server will be unaffected and remain where they are.
- The third option is **Mixed Mode**, a mix of Maintenance and Quarantine modes. Using this option, users can instruct the system to e.g. go into quarantine mode for moderate failures and into maintenance mode for severe failures, as shown in Table 1 below:

Maintenance only	Maintenance mode for all failures
Quarantine only	Quarantine for all failures
Mixed mode	Quarantine mode for moderate failures and Maintenance mode for severe failures

Table 1: OMIVV enables Proactive HA responses in three different modes

Licensing

Both VMware vSphere and OMIVV are licensed products. The VMware vSphere licensing requirement to use Proactive HA starts with vSphere Enterprise Plus. OMIVV is licensed per PowerEdge ESXi host on a three or five year subscription model. Please contact your Dell EMC sales representative for a quotation.

A 90-day evaluation license of OMIVV covering five ESXi hosts is available for registration and download here:

<https://marketing.dell.com/software-download>

Summary

Using the OpenManage Integration for VMware vCenter, IT Administrators can benefit from having VMware vCenter obtain the hardware health status of PowerEdge servers and take automated action based on hardware status changes. Supporting Proactive HA within the OpenManage Integration for VMware vCenter further enhances the functionality of VMware vSphere, the manageability of PowerEdge servers and provides another powerful tool to help IT Administrators save time, save money and reduce the potential for downtime.

Links for further information:

- For an expanded discussion of topics in this Tech Note, please see the 25-page white paper “**Proactive HA with Dell EMC OpenManage Integration for VMware vCenter (OMIVV)**” at: http://en.community.dell.com/techcenter/extras/m/white_papers/20444225.aspx. This white paper provides a deeper discussion of the topics, set-up and enablement guidance, troubleshooting tips, and many descriptive graphics and screen shots.
- The OMIVV virtual appliance can be downloaded directly from the Dell EMC Support webpage: <http://www.dell.com/support/home/sg/en/sghsd1/Drivers/DriversDetails?driverid=DYKN6>.
- A 90-day evaluation license of OMIVV covering five ESXi hosts is available for customers who register here: <https://marketing.dell.com/software-download>
- Additional information pertaining to OMIVV, including videos and whitepapers, can be found on the OMIVV section of the Dell EMC TechCenter webpage: <http://en.community.dell.com/techcenter/systems-management/w/wiki/1961.openmanage-integration-for-vmware-vcenter>