

Virtualizing Business-Critical Applications with Confidence

TECHNICAL WHITE PAPER

Challenges Using Traditional High-Availability Solutions

Within any business, there are a number of applications that are critical to its success. As a result, these applications and the systems they run on require a higher level of availability. One of the most common methods used to increase the availability of a business-critical application in a physical environment is to deploy a traditional high-availability clustering solution such as Veritas Cluster Server (VCS) or Microsoft Cluster Services (MSCS). These solutions protect against an unplanned failure of a component by providing the ability to restart an application or set of applications on additional servers in the cluster. While typically associated with protection against unplanned server failures, these solutions can also be utilized to decrease the effects of a planned outage by shifting applications to redundant servers, allowing for maintenance of the original server.

The trade-off with increasing application availability through traditional high-availability clustering is an additional cost in terms of redundant hardware, clustering software/support, and additional complexity. Management costs also increase due to the need to maintain multiple systems that are identical in configuration and patch levels. Operationally, it is extremely difficult to deploy a limited number of spare servers to provide redundancy for a larger set of applications due to difficulties with application compatibility, server patch levels, and so on. This typically results in the use of small two-node clusters deployed for only the most critical applications, leaving the majority of applications not clustered at all.

High Availability in VMware Environments

As customers move forward with VMware® virtualization solutions, they recognize an immediate set of benefits far beyond a simple reduction in servers. VMware includes proven and widely deployed business-continuity solutions in VMware vSphere™ 4.1 (“vSphere”) in the form of VMware vMotion™, VMware High Availability (VMware HA) and VMware Fault Tolerance (VMware FT).

Utilizing VMware’s revolutionary vMotion technology, IT administrators are able to move applications for server maintenance with zero downtime and data loss. Coupled with the operating system isolation natively provided by VMware virtualization, it is very simple to provide a small set of highly consolidated servers capable of providing very high uptime with reduced administrative cost.

VMware HA provides a simple, reliable way to increase the availability of virtual machines hosting critical applications. VMware HA is a virtualization-based distributed infrastructure service of VMware vSphere 4.1, which monitors the health of virtual machines and the VMware ESX® hosts upon which they reside. If a fault is detected, the virtual machine is automatically restarted on another ESX host with adequate capacity to host it. VMware HA is included in all vSphere editions and can be enabled on a VMware cluster with a single check box. As VMware HA utilizes the storage and network connectivity already in place to support vMotion, enabling high availability is as simple as ensuring you have adequate server capacity to handle failure of one or more ESX hosts.

VMware FT extends the capabilities of VMware HA to provide even higher levels of availability for mission-critical applications by allowing instantaneous transfer of services to a secondary image of VMware FT-enabled virtual machines. This allows virtual machines to continue operations, even when a server failure occurs, with zero downtime and user interruption.

Challenges in Virtualizing Tier 1 and Tier 2 Applications

VMware HA and VMware FT technology provide increased availability for a large percentage of customers. In fact, more than 80 percent of VMware customers leverage one or both of these technologies to protect most or all of their virtual machines. However, a method to increase the availability at the application layer is often desired, especially for business-critical Tier 1 and Tier 2 applications. Without this protection at the application level, organizations are exposed to application failures that might happen inside the virtual machine. In many cases, organizations have attempted to deploy a traditional application-clustering solution into the virtual machine’s guest operating system for this purpose.

While deploying a traditional application-clustering solution into the virtual machine addresses failures at the application layer, it also creates significant issues with the day-to-day operations of a virtualized environment, as these solutions were designed for physical environments. These issues include the added complexity of maintaining multiple identical virtual machines to properly host failover, additional capacity needed to host spare servers, and difficulty in mapping application location to a specific virtual machine within the VMware management solution. More important, the addition of in-guest clustering significantly impacts the ability to make use of advanced VMware features such as vMotion, VMware HA, VMware Distributed Resource Scheduler (VMware DRS), and VMware Distributed Power Management (VMware DPM).

Extending Application High Availability in vSphere 4.1

With the release of vSphere 4.1, VMware is introducing an application programming interface (API) to allow third-party software vendors to deploy application monitoring components inside a VMware guest OS and inform VMware HA when problems arise. This API will allow application-clustering vendors to develop application monitoring and control solutions that fully complement the virtual machine high availability and management provided by vSphere.

The joint solution will include two layers of protection. The first is the in-guest protection provided by the application HA vendor. This application-layer protection can include application-specific capabilities such as component-level monitoring, restarting of failed services, performance monitoring, and so forth. The second layer is VMware HA, which can restart the virtual machine in cases where the in-guest solution cannot resolve the issue.

The user can enable application monitoring as a part of the VM Monitoring Status section of the VMware HA settings. Enabling application monitoring allows the application-monitoring solution to register with the VMware application awareness API and communicate application status with VMware HA. Inside vCenter Server, the user will be able to determine which virtual machines are monitored on the application level and which ones are only monitored for basic virtual machine health. In Figure 1 we see how the user can control the monitoring level in the cluster and per virtual machine.

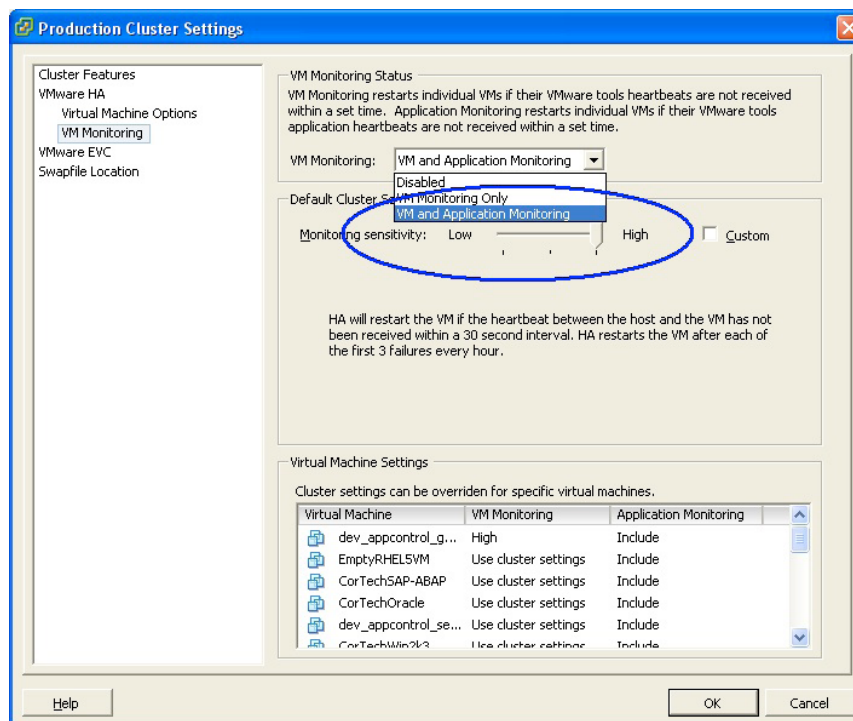


Figure 1. Configuring the VM Monitoring Level to Include Application Monitoring.

Introducing the VMware HA and Symantec ApplicationHA Solution

VMware is pleased to work jointly with Symantec to provide Symantec ApplicationHA for VMware. This solution marries the best-in-class virtualization infrastructure from VMware with the industry-leading application high availability package from Symantec. ApplicationHA leverages more than 12 years of development of Veritas Cluster Server to provide an application-monitoring package that runs inside a VMware guest operating system and fully integrates with VMware HA to provide virtual machine restart as needed to react to any application issue.

Combined together, this solution:

- Enhances the availability of Tier 1 virtualized applications by providing a mechanism to detect and recover from application failures
- Improves the manageability of virtualized applications by providing visibility of the application's status and application control (that is, starting or stopping the application) through a single pane of glass — VMware vSphere Client, the user interface for vCenter Server
- Reduces operational complexity and TCO by eliminating the need for different operating system-based clustering products
- Allows customers to enhance the availability of their applications without sacrificing the use of advanced VMware features such as vMotion and DRS

Symantec ApplicationHA Overview

Symantec ApplicationHA leverages the VMware HA application awareness API to provide comprehensive application availability in VMware environments. It is comprised of two main sets of components, as shown in Figure 2:

- The guest component, which is installed in each ApplicationHA-enabled virtual machine
- A vCenter Server plug-in that provides visibility and management of the applications

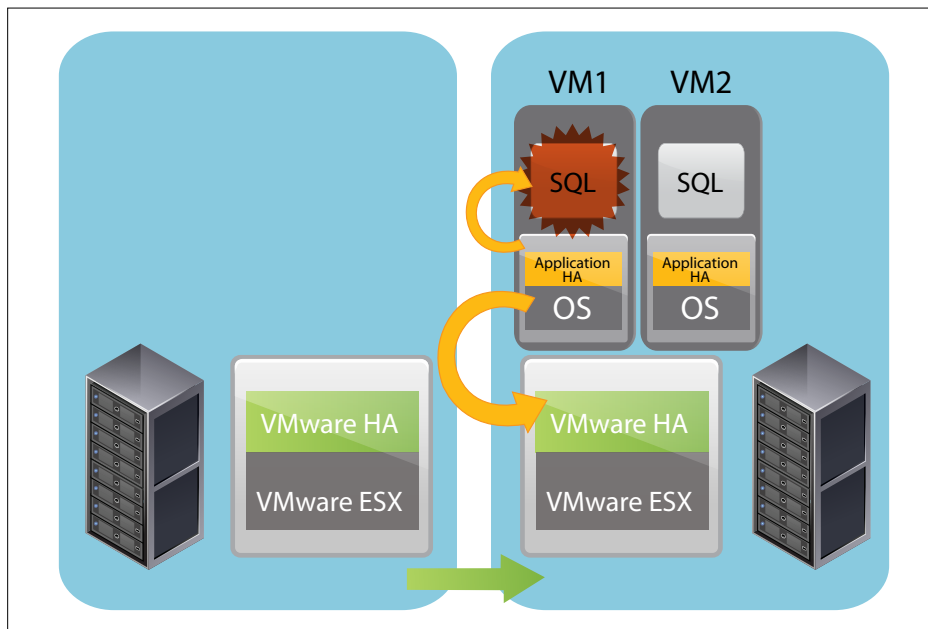


Figure 2. Symantec ApplicationHA Overview.

The guest component encompasses an application agent framework and various application agents. The application agent framework provides the infrastructure that is utilized by the application agents in their execution. The application agents are responsible for the starting, stopping, and monitoring of a given application resource or instance.

By defining the resources that comprise an application, the application agents are able to monitor, start, and stop the application instance and any related resources. For example, if a Microsoft SQL database were to be put under ApplicationHA control, several resources may be required in order to support the database. These might include mount points to make the storage available, the Microsoft SQL instance, a Microsoft SQL Agent to perform scheduled administrative tasks, and a Microsoft SQL Online Analytical Processing (OLAP) service for multidimensional analysis. These all would be individual resources that would need to be monitored to ensure proper operation of the Microsoft SQL database application.

Additionally, in order to bring the Microsoft SQL application online or offline, these resources would require a specific order of operations. To support this, resources can be made dependent on each other, similar to the depiction in Figure 3.

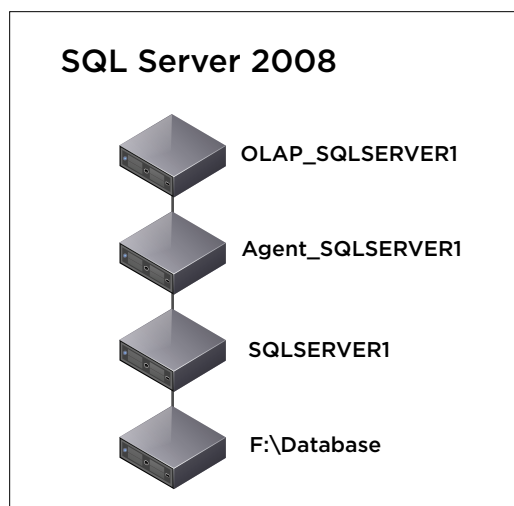


Figure 3. Microsoft SQL 2008 Application Resource Dependency.

Through the use of this dependency model, when an administrator starts or stops the application they can be assured that the application resources are handled by the guest component in the correct sequence.

The resources that comprise an application are continuously monitored at a given interval to ensure proper operation. If the monitoring of a resource detects a failure, the guest component takes action:

1. The guest components attempt to restart the application within the virtual machine. The number of attempts that will be made to restart an application is configurable by the user.
2. If the application does not restart successfully, the guest components communicate to VMware HA through VMware's application awareness API in order to trigger a reboot of the virtual machine by VMware HA. The application is restarted as part of this reboot process.

The Symantec ApplicationHA vCenter Server plug-in displays the status of the application (Offline/Online/Faulted/Partial), as well as the status of the individual resources comprising the application. ApplicationHA state changes are also reflected under the vSphere Alarms tab, enabling administrators to configure notifications as desired.

In addition to providing visibility of an application's state, the Symantec ApplicationHA vCenter Server plug-in also allows for the management of the application and ApplicationHA features. This includes the ability to start and stop the application, the ability to enable or disable ApplicationHA functionality and the ability to disable the communication to VMware vCenter Server in order to allow a user to troubleshoot a problem application without triggering a restart of the virtual machine.

Application and Platform Support

Symantec ApplicationHA provides application availability in Windows and Linux VMware guests. It provides extensive support to common off-the-shelf Tier 1 applications such as Microsoft Exchange, Microsoft SQL, IIS, WebLogic, Oracle, and SAP, as well as the ability to monitor and manage custom applications.

Please refer to the Symantec ApplicationHA documentation for the latest information about the supported applications and platforms.

Installation and Configuration

Symantec ApplicationHA provides users with a simplified wizard-based installation and configuration process. The installation wizard allows for the installation and registration of the VMware vCenter Server plug-in as well as the installation of the guest components into a virtual machine.

The installation of the guest components can be remotely pushed out to one or more virtual machines at a time. A view similar to a VMware datacenter tree view in vCenter Server is provided in the installation wizard, in which an administrator can select a group of virtual machines as a target for installation. This view provides several levels of granularity of selection, including that of a VMware datacenter, VMware HA cluster, or an individual virtual machine.

For example, if an administrator wants to provide ApplicationHA to all virtual machines in a Microsoft SQL cluster, they can select those virtual machines, as shown in Figure 4.

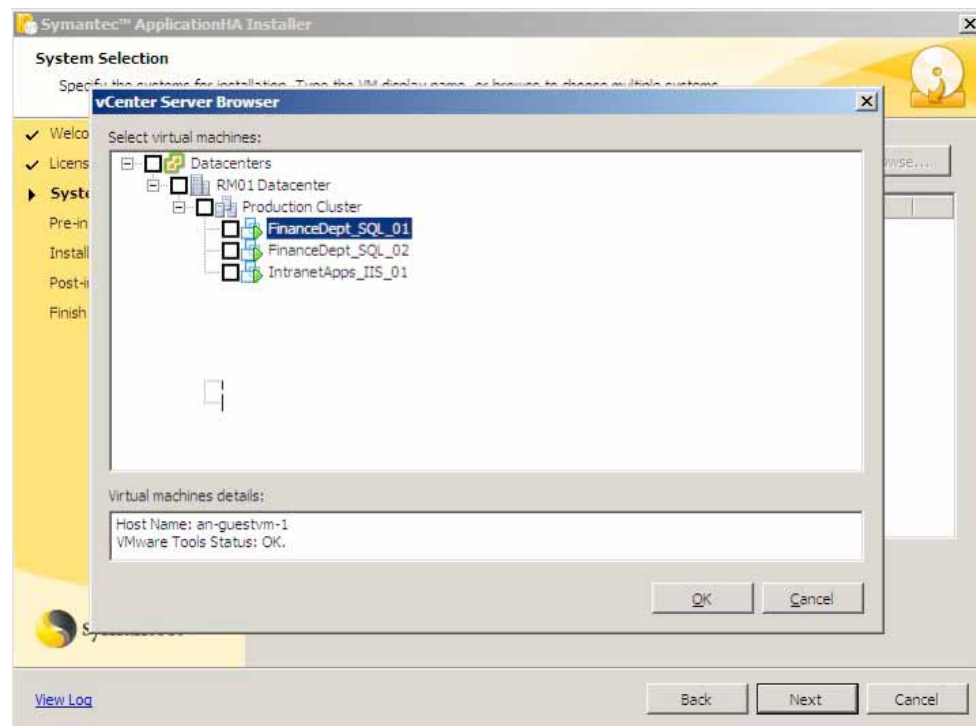


Figure 4. Symantec ApplicationHA Installation Wizard.

A simple wizard-driven process is also provided by Symantec ApplicationHA to assist in configuring and monitoring an application. Off-the-shelf applications, such as Microsoft SQL or Microsoft Exchange, can be configured by this wizard using default parameters common to these applications.

Administrators can also protect non off-the-shelf, or custom, applications. The configuration process for custom applications is also wizard based, as shown in Figure 5, making it easy for users to deploy both packaged and custom applications. By selecting different services, processes and resources that need to be monitored, an administrator can provide enhanced availability to a practically limitless set of applications.

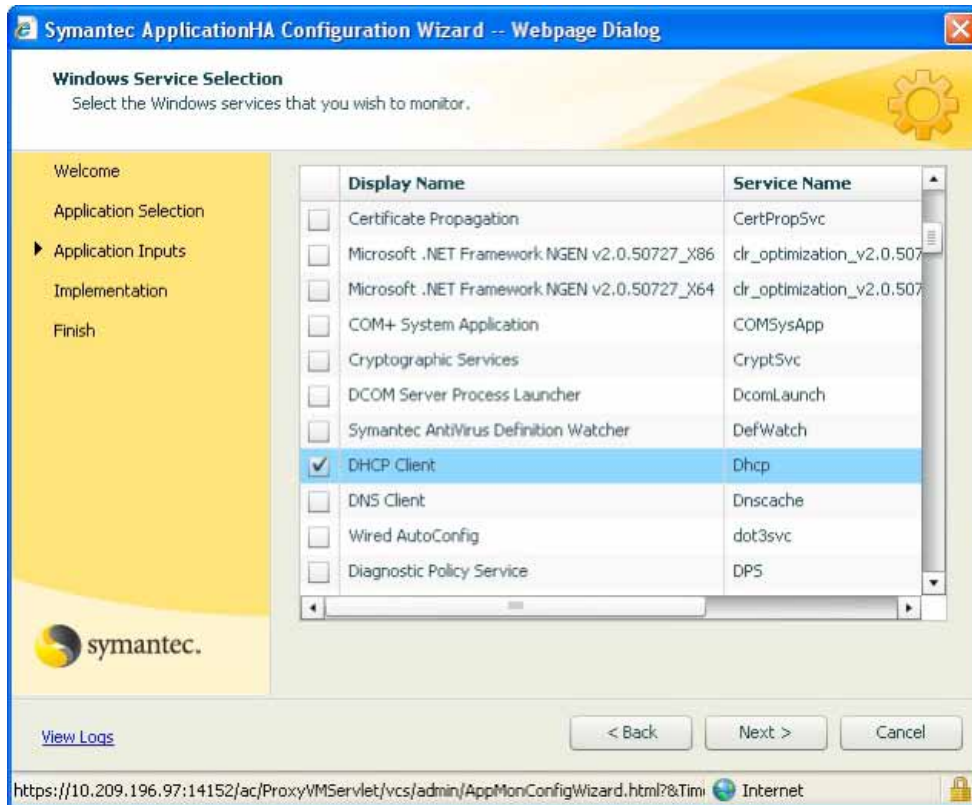


Figure 5. Symantec ApplicationHA Custom Application Configuration Wizard.

Management of Symantec ApplicationHA

Symantec ApplicationHA is managed through VMware vCenter Server through its provided plug-in. Using an already familiar interface, VMware administrators can view the status of an application and the resources that are associated with it.

Additionally, it provides the ability to perform operations specific to Symantec ApplicationHA, such as:

- Start or stop an application
- Enable or disable the communication between VMware HA and Symantec ApplicationHA
- Configure or unconfigure Symantec ApplicationHA

For example, Figure 6 shows the view of a Microsoft SQL virtual machine that has been enabled for ApplicationHA. The plug-in appears in the context of the virtual machine in vCenter Server's Datacenter view on the left.

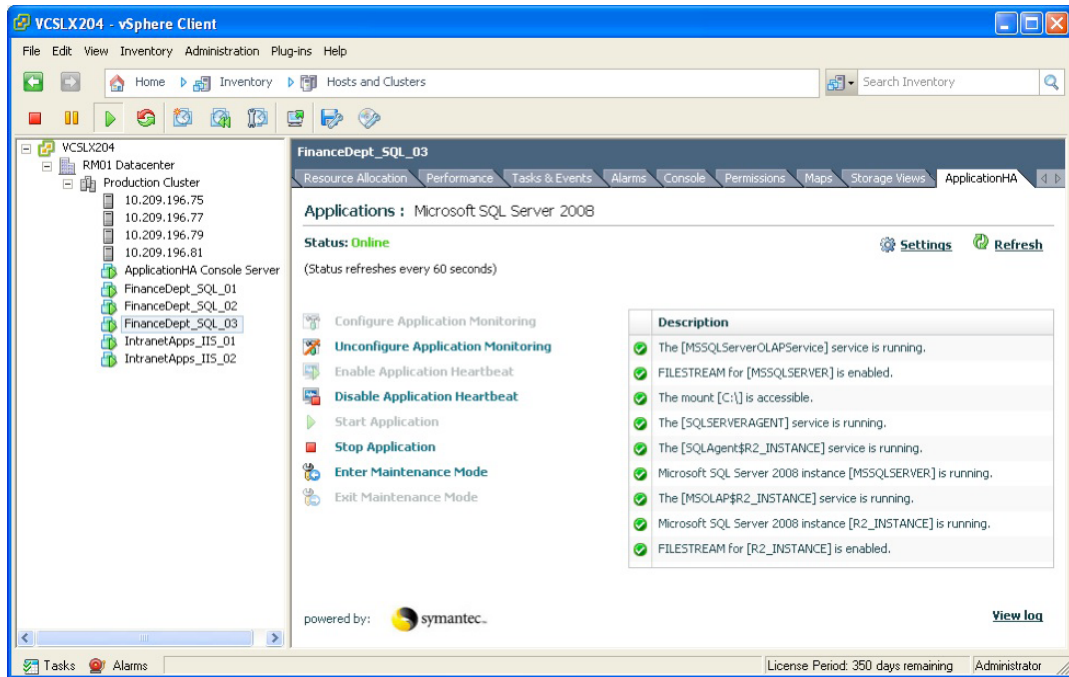


Figure 6. SQL Application Status in Symantec ApplicationHA Tab in vSphere Client.

User access control is provided through the role-based access control mechanism provided by VMware vCenter. This allows for three types of privileges that can be assigned to the vCenter Server user roles in order to provide the appropriate level of access. These three privileges are defined as:

- Guest: Has visibility into status of the application
- Operator: Has visibility into the status of the application and the ability to start or stop the application
- Administrator: Has visibility into the status of the application, the ability to start or stop the application, and the ability to configure ApplicationHA

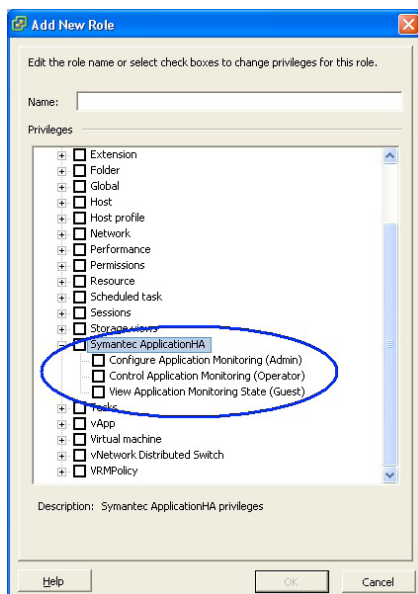


Figure 7. Privilege Model with Symantec ApplicationHA.

ApplicationHA Best Practices

As with any solution, the key to a successful ApplicationHA deployment is to follow best practices. The following list highlights common ApplicationHA best practices:

- Prior to installing ApplicationHA, read the release notes and check the Symantec Veritas Operation Services (VOS) Web site (vos.symantec.com) for any software updates and late-breaking news.
- Configure the VMware HA clusters in vCenter Server prior to installing ApplicationHA. As ApplicationHA leverages the VMware HA cluster configuration details to help simplify deployments, doing this first will simplify the installation process.
- Install the ApplicationHA client inside all virtual machines configured in a VMware HA cluster. Enable “VM and Application Monitoring” for all virtual machines in a VMware HA cluster. Monitoring all applications in the cluster ensures the highest levels of HA.
- When running vCenter Server inside a virtual machine, install the ApplicationHA client components inside the virtual machine and use ApplicationHA to monitor the database and vCenter Server services.
- Prior to configuring application monitoring, ensure the applications are fully installed, configured, and running. ApplicationHA discovery is able to detect installed applications and automatically set up application monitoring.
- Use Veritas Operations Manager (VOM) for managing applications across physical and virtual environments from a single pane of glass or to visualize and protect multi-tier applications or for enabling users who do not have access to vCenter Server but need to visualize and control applications in VMware virtual machines.

Conclusion

VMware’s application awareness API provides a foundation that third-party vendors can build upon in order to provide application awareness within a virtualized environment. Symantec ApplicationHA utilizes this API to develop a robust solution that allows for increased levels of availability for critical applications, avoiding many of the issues encountered when implementing a traditional clustering solution in a virtualized environment. VMware administrators will also appreciate the VMware vSphere interface that they are familiar with to visualize and administer the solution.

Overall, this solution represents another leap forward in providing support for the most critical applications within a VMware virtualized environment.

