Saving Your Servers from Disaster Using VMware Virtualization

In this article, we'll take a look at some of the key features of virtualization—specifically VMware virtualization solutions—that help defend your datacenter against disaster. Virtualization offers hardware independence—a huge time saver in recovering from a disaster. In addition to this core capability, virtualization provides several other ways to protect your business from server downtime.

Disasters come in all shapes and sizes

Disasters that can wreak havoc on your infrastructure include:

• Total server failure (for example, local disk failure or motherboard failure).

• Pre-failure warning on a shared disk array.

• A runaway process on a server that causes high CPU utilization.

• An entire datacenter that is hit by a major disaster, like a tornado or fire.

Downtime is a serious risk to businesses of all sizes. Even a minute of downtime on a single server can cost millions in lost revenue. Many companies never recover from a service outage. Fortunately, virtualization enables hardware independence that prevents the effects of server downtime caused by any type of disaster in your datacenter. Virtualized servers are entirely encapsulated in a virtual machine and stored as a set of files in shared storage. The source and the target servers can access these files concurrently. The active memory and precise execution state of a virtual machine can then be rapidly transmitted over a high-speed network.

In other words, virtual machines are completely portable. Their state can be mirrored from ESX server to ESX server, using Fault Tolerance (FT), in such a way that if a server fails, no end user applications go down.

You can also use VMware® VMotion and VMware® Storage VMotion technologies, to perform live migrations and hardware maintenance with zero downtime and proactively move virtual machines away from failing or underperforming servers, thereby protecting them from unplanned downtime, like a disaster.

To illustrate what I mean, consider the following situation. Imagine you have an email server running Microsoft Exchange and providing email services for 500 users, and you suddenly have a serious hardware failure. You find out that it could take over 24 hours to locate and install the parts necessary to repair the server.

However, you happen to have an extra server available, but it's a different model from a different vendor. What would it take to get the email server running on the extra server? With a physical infrastructure, it could take a lot of work over many hours or days, depending on the complexity of the failure. However, if that server was virtualized, you could have it running in minutes. All you would need to do is restore the virtual machine from shared storage to the new host.

In another example of a server hardware failure, let's say that the same Exchange server was virtualized but the ESX server that it was running on has failed. Fortunately, that Exchange virtual machine was protected by VMware's new Fault Tolerance (FT). In this case, the end users never even noticed a second of downtime in their mail client. How is that possible with a total server hardware failure? With FT, the system memory of the running Exchange server was sent to another ESX server that provides the secondary version of that Exchange server. Once this is synchronized, all changes made on the primary Exchange server are sent over the wire to the secondary. When that primary ESX server with the Exchange VM failed, the secondary VM just took over. Selfishly you could be amazed by the time saved and headaches prevented to you, as a VMware Admin, but the benefits to the end users and business are far greater.

Here's another example. Say you have a shared iSCSI storage array that alerts you with a pre-failure disk warning. Or, perhaps the array needs a firmware upgrade that requires a restart. Either of these scenarios requires you to move running servers off of the shared storage onto the local storage of your server. With VMware® Storage VMotion, you can move that running storage onto your local VMware ESX host, and bring down the storage array without shutting down the guest virtual machines.
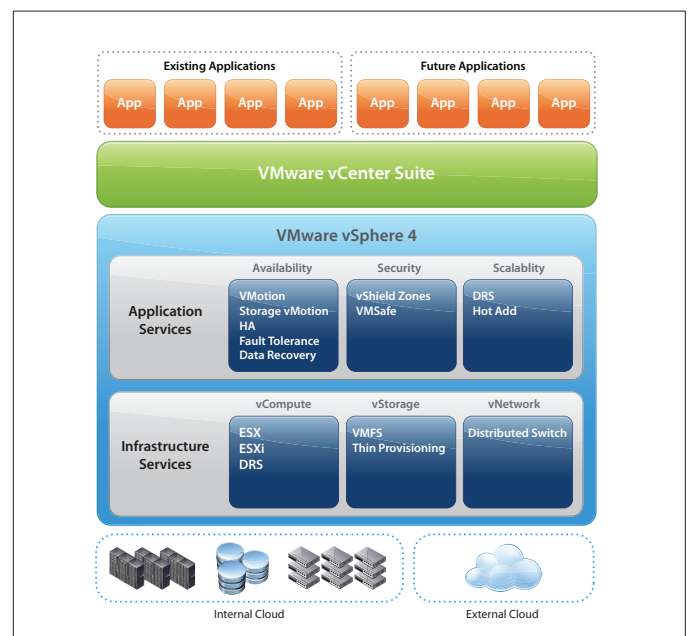
DR features that save the day

In any disaster, fast recovery and continuous service are critical requirements. VMware High Availability (HA) allows you to automatically migrate and restart all virtual machines from a failed server to any other server in a high availability resource pool, immediately, to accommodate virtual guests.

When I say "immediately," I mean that guests are moved in just a few seconds and become available as soon as you reboot the servers. VMware HA is made affordable for smaller IT environments as part of the VMware Essentials Plus Edition.

Also included in the Essentials Plus Edition is Data Recovery — VMware's new Backup and Recovery application for vSphere. As Data Recovery understands your VMware virtual infrastructure it allows you to more quickly and easily backup and restore virtual machines. With Data Recovery, you can backup VMs with no downtime to the end users and restore VMs so easily that any junior admin can do it.

VMware DRS is available in VMware vSphere Enterprise and Enterprise Plus editions.

Moving on to some of the more advanced disaster recovery features, we have FT and VMotion. As we discussed above, FT allows you to synchronize between a primary and secondary VM running on two ESX servers. If the primary ESX server is lost, the secondary VM on the secondary ESX server just takes over, with no downtime for the end user applications. While we have many examples of disaster recovery features in this article, FT is likely the most technologically amazing feature that VMware has to offer today.

VMotion allows you to move a running virtual machine from one ESX server to another, again, with zero downtime to the end user applications. By using VMotion, virtual machines can move off of an ESX server that needs to go down for hardware maintenance or ESX server patches.

Both VMware FT and VMotion are included starting in the vSphere Advanced edition. Another feature, VMware Distributed Resource Scheduler (DRS), provides for server load balancing across multiple VMware ESX hosts. In terms of disaster recovery,

VMware DRS can "save the day" by moving virtual guest servers from one VMware ESX host to another, in the event that the primary host is overloaded. Consider a situation in which one of your less critical servers has a run-away process and is taking all the CPU resources of one of your VMware ESX hosts. If that server also supports another workload, such as the database server for your e-commerce site, all e-commerce transactions would be affected. VMware DRS will move the database server to another VMware ESX host in the resource pool that has free resources, without downtime. DRS utilizes VMotion to balance the load of VMs by migrating them across the vSphere virtual datacenter.

Finally, a very valuable disaster recovery feature is VMware's Site Recovery Manager (SRM). This product is dedicated to bringing up the virtual infrastructure in an orderly way, as you previously defined and tested. You can read more about SRM and download a free trial at the Site Recovery Manager product page.

Pick your flavor

Disasters come in many flavors, and VMware virtualization solutions can help you address them quickly and easily. With VMware vSphere Essentials Plus Edition and the vSphere Acceleration Kits, even smaller IT environments can implement enterprise-class virtualization and disaster recovery solutions, affordably. To learn more, visit www.vmware.com.

About David Davis

David Davis is Director of Infrastructure at TrainSignal.com – the global leader in video training for IT Pros. He has a number of certifications including vExpert, VCP, CISSP, and CCIE #9369. Additionally, David has authored hundreds of articles and six different video training courses at Train Signal with his most popular course being VMware vSphere. His personal website is VMwareVideos.com. You can follow David on Twitter or connect with David on LinkedIn.