Global Knowledge®

Expert Reference Series of White Papers

Virtual Machine Migration Methods in vSphere

Virtual Machine Migration Methods in vSphere

John Hales, Global Knowledge VMware Instructor,
A+, Network+, CTT+, MCSE, MCDBA, MOUS, VCP, VCAP, VCI, EMCSA

Introduction

One of the advantages of vSphere is that you can move a virtual machine from one location to another, across servers, storage locations—even data centers. Physical servers don't have that ability and that can have many implications for disaster recovery, availability, etc.

The goal of this white paper is to explain:

1) Why migrations are useful
2) The methods that vSphere makes available for you to manually move a virtual machine (VM)
3) How vSphere can automate the process for you in various scenarios

Finally, we'll take a look at some of the future directions that are expected in a future version of vSphere (vSphere 6 is expected to be available in the first quarter of 2015, but as it is in beta, this is subject to change).

Note that we will not describe the technical requirements to implement any of the technologies described here.

Use Cases for Migrations

Let's begin with a review of use cases for moving a VM from one location to another. Then as we go through the various migration mechanisms, you can see how they will be helpful in solving various challenges.

- **Performance:** The most common use case for moving VMs is for performance. You may find that one physical server is very busy from a CPU or memory perspective, while another server is not very busy. You can move a VM from one place to another for better load balancing. This can be done with no downtime to the VMs if desired.
- **Maintenance:** If you need to take a physical server down for planned maintenance, whether to do server upgrades like adding RAM or CPUs, to replace failed components, or to patch ESXi itself, you can move the VMs to other servers so that applications can stay online while you are performing maintenance. This allows you to do maintenance during normal business hours instead of overnight hours on weekends, for example.
- **High Availability:** If the server crashes, all of the VMs on that server will crash as well. In the physical world, access to the services provided by that server would be unavailable until the server is repaired and brought back online. With vSphere, however, when a server crashes (unplanned), the VMs can automatically restart on a different physical server. This can be extended even to include not just the loss of a server, but of an entire site with all of its servers and all (or some portion) of those VMs can be restarted on vSphere servers in a different physical site.
- **Physical Server Replacement:** If you need to replace a physical server with a new model, you can migrate your VMs to the new server and then decommission the old server. This can be done with no downtime to the VMs themselves.

Manual Methods

In this section, we will review all of the ways that you can manually move a VM from one location to another. We will list the restrictions around each option and the license level required to use it. All of them can be accomplished via the new Web Client, and all but the shared nothing method can also be done with the old C# (Windows) client. To do any of them, you will need to have the ESXi servers managed by vCenter. A cold migration can be done without vCenter and specific notes are given in that section.

CPU compatibility will not be described in detail in this white paper; but at a high level, either the CPUs are identical (same family, like Intel Broadwell or Haswell—speed, number of cores, etc., can vary without any problems) or you have implemented Enhanced vMotion Compatibility (EVC) in vCenter clusters or configured advanced options on each VM. All of the options below, with the exception of cold migrations, require compatible CPUs on the source and destination physical servers as just described.

Cold (Powered Off)

A cold migration is taking a powered off VM and moving it between servers and/or storage locations. It can be done with any license level. It is possible to do this type of migration without vCenter, but that involves removing the VM from the inventory (***not storage***) of the first server and then adding it to the second server. You can move the VM anywhere within vCenter, even across data centers. It doesn't require anything specific relative to the CPUs on either physical server as the VM will be powered on from scratch, so you could even move a VM running on an AMD-based host to one running Intel CPUs with no issues. This type of migration has almost no restrictions on what can be moved where.

Warm (Suspended)

A warm migration is taking a VM that has been suspended (via vSphere) and changing locations. It has the same capabilities and restrictions as a cold migration, with the exception of the CPU, which needs to be compatible. Thus, you can migrate a VM across hosts in different datacenters and/or across storage locations.

Hot (Powered On)

While you can move VMs that are powered off or suspended, the main advantage of VM migrations is that VMs can be moved while powered on with no downtime. This allows you to handle all of the use cases (other than the unplanned crash) with no interruption to the services provided by the VMs being moved. Virtually all of the migrations you are likely to do fall into this category of hot migrations.

vMotion

vMotion takes a live, running VM and moves it from one host to another. vMotion can only be used to change hosts, not storage locations, thus the VM must be running on shared storage accessible to both source and destination hosts. vMotion can only be used within the same vCenter Datacenter as well. Of all the migration types, vMotion is the most restrictive in terms of requirements, but also the most often used. It requires a VMkernel Port configured on each host with the vMotion capability enabled. Also note that the VM memory will be moved across this network unencrypted, so this should be a secure, closed network for the hosts to use or across a network encrypted link, such as across a VPN.

To use this feature, you must have vSphere Essentials Plus or higher and vCenter.

Storage vMotion

Storage vMotion takes a live, running VM and moves it from one storage location to another, all the while the VM stays running on the same host, thus no downtime. Because the VM is moving between storage locations on the same host, the host must have access to both storage locations.

From a licensing perspective, it requires vSphere Standard or higher and vCenter (though Essentials Plus does allow you to change both host and datastore at the same time, it does not support standard sVmotion).

vSphere can leverage capabilities native to the array to move a VM within a single array without copying all the data to the server and then writing it back to the array. This requires Enterprise or Enterprise Plus licensing and a storage array (either block storage or NAS, it does not matter with ESXi 5.0 and later) that supports the VAAI (vSphere APIs for Array Integration) standard. Instead, vSphere tells the array to move it from the source location to the destination and let it know when it is finished with the move. It is much faster, more efficient, and lowers the load on both the host and the storage network.

Shared Nothing

With Shared Nothing vMotion (also called Enhanced vMotion, Cross Host vMotion, x-vMotion, or Unified vMotion—it doesn't technically have an official name), you are combining both a Storage vMotion and a vMotion in one migration. This allows you to do a vMotion migration without shared storage—in fact, you could even migrate from Host A on local storage to Host B on local storage. Note that the CPU compatibility requirements of vMotion are still required.

Note that this must always be done by an administrator. None of the automated solutions described below will ever use this mechanism.

From a licensing perspective, if you have a license that supports vMotion as described previously, you can also do this type of vMotion.

Automatic Methods within a Site

While it is great that VMs can be manually moved from one location to another, the great power that VMware provides is in the automation of these migrations. It would be expensive and time consuming to have an administrator constantly monitor for any condition that could be improved through the use of one of the migration methods as well as to handle various server failure scenarios. This is something that computers are great at—gathering basic data points and taking automated action when one or more are outside defined limits.

In this section, we'll look at all of the fully automated ways to move a VM. The VMs need to be in the same vCenter Datacenter for these technologies to work, as clusters are limited to a single datacenter. We'll look at handling physical server failures (or component failures that may make the server unavailable) as well as load balancing mechanisms available. Both High Availability (HA) and Distributed Resource Scheduler (DRS) require clusters to be created. Both require shared storage to function.

High Availability (HA) and Fault Tolerance (FT)

High Availability's goal is to minimize downtime, not prevent it. This feature is available in all editions of vSphere except Essentials. It is designed to handle the failure of any or all of the following:

- Loss of a physical ESXi server
- Loss of a VM
- Loss of an application within a VM

In the first case, the loss of a server, when the server fails all VMs on that server fail immediately. Within a few seconds, one or more servers in the cluster will know that the server has failed and the master node, which coordinates all of the HA activities within the cluster in conjunction with vCenter, will assign the failed VMs to the surviving nodes in the cluster. In this case, no vMotion was performed—the VM failed and was then restarted on a different node. This can be done even if vCenter is unavailable—in fact, HA can even restart a failed vCenter server when the host it was running on failed. This is the original and primary use case of HA. It is also the only use case that will move a VM to a different physical server.

From a configuration perspective, you can mix and match any kind of server in the cluster as the VM will be restarted from scratch. Thus you could (but really shouldn't) mix and match Intel and AMD in the same cluster.

In the remaining two use cases, the VM is restarted on the same physical host, and thus it is not really a migration scenario. We will not describe them further here as a migration is not involved.

What if downtime is unacceptable? HA is not sufficient. This is where Fault Tolerance (FT) comes in. FT is a feature of HA; in other words, you first configure HA and then layer on the FT capability.

With FT, when the VM is powered on or the feature is added to a running VM, FT will basically do a vMotion to a second server in the cluster, but instead of completing the move like standard vMotion does, FT will keep the two servers in lock step (via vLockstep Technology). If the primary copy fails, within a couple of seconds, the secondary will become the new primary and keep processing and will in turn spin up a new secondary server. On the other hand, if the secondary fails, the primary will create a new secondary and keep running. In either case, the goal of FT is no downtime and no data loss for applications.

From a configuration perspective, since the VM is running in parallel on two different servers, the servers must be vMotion compatible (among other configuration requirements).

HA is included with vSphere Essentials Plus and higher editions.

Distributed Resource Scheduler (DRS)

While you could manually vMotion VMs around to handle all the use cases (other than host or VM failure described in the HA section), doing so would be time consuming for an admin, especially if trying to load balance the VMs. To this end, the DRS feature, available in Enterprise and higher editions of vSphere, will automatically take care of that for you.

It will move VMs from one host to another via vMotion to load balance CPU and memory utilization across a group of hosts that are linked into a cluster. It can also migrate all the VMs off a host that are in need either physical maintenance (such as fixing broken components or upgrading RAM, CPUs, etc.) or software maintenance (such as upgrading the version of vSphere or patching drivers).

In addition, rules can be specified that keep a group of VMs together on the same host (known as an affinity rule) to reduce the physical bandwidth required for VMs that frequently communicate between each other.

Likewise, rules can be created that keep a group of VMs on separate hosts (known as anti-affinity) rules. This is often because they consume a lot of resources and would not necessarily play well together on the same host from a resource consumption point (think several large database servers). Another reason to use this type of rule is for high availability reasons. You wouldn't want all of your Active Directory domain controllers on a single host, because if it were to fail, you'd have no running domain controllers, at least until HA got one or more restarted.

The third type of rule is a VM-Host affinity rule. With this type of rule, you can group VMs together and hosts together and then specify that a group of VMs should/should not or must/must not run on a group of hosts. Should rules are just that—suggestions that can be overridden if necessary. Must rules, however, can never be overridden for any reason—not by HA, DRS, or even manually via vMotion. Should rules are often created for high availability reasons or to keep location-critical VMs (such as vCenter) on just a few hosts so they can be found if vCenter is unavailable for any reason. Must rules are typically created for software licensing reasons to keep VMs on certain hosts in a cluster that have been licensed for certain applications and off the rest of the hosts in the cluster.

Just a quick tip on rules: while you can create many rules and DRS will enforce them, the more rules you create, the fewer options DRS has to load balance, thus potentially giving sub-optimal load-balancing solutions to meet all the rules.

DRS comes with vSphere Enterprise Edition and higher and has been available for many years.

Storage DRS

New in vSphere 5, Storage DRS aims to do for storage what DRS does for compute (i.e., CPU and memory), namely to load balance. Storage DRS (often called simple SDRS) allows a group of datastores to be combined into a cluster, much like a group of hosts can be combined into a host cluster. Just like a DRS cluster can't (for all practical purposes) contain both AMD and Intel CPUs, so an SDRS cluster can't contain NFS and VMFS volumes in the same cluster.

An SDRS cluster is designed to load balance storage capacity and optionally I/O latency across all of the nodes in the cluster. Parameters can be specified to determine the threshold when it begins load balancing (for example not until it becomes half-full) and how big the difference between two datastores needs to be before moving VMDKs between them.

An SDRS cluster has both intra-VM affinity and anti-affinity rules, similar to a DRS cluster. These rules keep all of a VM's virtual disks on the same datastore within the cluster (typically to ensure that the VM will have all the components it needs to run) or they can separate them onto different datastores within the cluster, usually for performance reasons. Rules can also be created to keep a group of VMs separated onto different datastores, often for high availability reasons. An example of this is a rule that separates the Active Directory domain controllers onto different datastores within the cluster so that a disk fault that causes VMs to be inaccessible doesn't take out all of an organization's domain controllers.

SDRS checks disk space utilization every five minutes by default to ensure that a datastore doesn't run out of space, impacting all the VMs on that datastore. On the other hand, in balancing latency, it looks at an eight-hour average by default, as many VMs couldn't even be completely moved in five minutes and it would be a huge waste of resources to try to do so. Rather, the goal is to try to achieve the best long-term latency possible. So how are the short-term spikes in demand handled? How does vSphere ensure that the right VMs have the right access to storage? Storage I/O Control (SIOC). It uses shares and optionally input/output operations per second (IOPS) limits to handle the minute-to-minute fluctuations in demand and does so on a vCenter-wide basis.

The goal of an SDRS cluster, like a DRS cluster, is to take the manual process of load balancing, watching for overloaded datastores, etc. and automate (subject to the policies specified by the administrator) the entire process under the direction of the vCenter server.

Storage DRS is only available in the Enterprise Plus edition of vSphere.

Semi-Automatic Methods between Sites

The methods described in the Automatic Methods section work great within a site and are more or less fully automated based on the policies set by the administrator. They work well and should be implemented.

But what about Business Continuity (BC) and Disaster Recovery (DR)? You probably don't want them fully automated. You don't want VMs to move on their own, take outages, etc. without someone making a conscious choice to do so. That's where vSphere Replication (VR) and Site Recovery Manager (SRM) come in. They require some administrator action to initiate.

vSphere Replication (VR)

vSphere Replication (VR) has been included with vSphere since 5.0, though the interface to manage it outside of SRM didn't get introduced until 5.1. vSphere 5.5 added many new features, including:

- The ability to make Multiple Point-In-Time (MPIT) copies of a VM, i.e., to keep copies of up to 24 previous points in times that can be reverted to if the most recent one has an issue, such as a virus or data corruption. You specify how many copies to keep per day and how many days to keep them.
- Network improvements that make the copies faster (from two to 100 times faster).
- Support for VSAN as the datastore on either side.

VR is included in all editions of vSphere (Essentials Plus and higher) for free.

VR is configured on a per VM basis, specifying where to replicate each VM (site, cluster, and datastore) and the Recovery Point Objective (RPO) (worst case, how much data can be lost). The RPO can be specified from fifteen minutes to twenty-four hours. Note that you don't tell VR when to replicate—it decides based on volume of data, other VMs that need to be replicated, available bandwidth, etc., how often to replicate to meet the configured RPO. In addition, you can have VR use VSS in Windows to quiesce the applications to get application consistent points in time, not just file system consistent points.

After the initial full replication (over the network or by restoring a backup for example), only the changes since the last replication will be replicated to minimize both the time and bandwidth required to replicate the VM. The system saves the changes on the destination side in a redo log, which is automatically applied once the replication successfully completes (or can be rolled back in the event of a failure), ensuring that data is always in a consistent state.

Multiple replication topologies are possible including unidirectional from a protected to a DR site, bidirectional between two sites (for different VMs, each VM is unidirectional), distributed sites to a central site (e.g., branch office to headquarters), and central site to branch offices.

To failover a VM to the remote site, you run a wizard and it will bring the VM online; this needs to be done individually for each VM to relocate. Once online, to revert to a previous point in time, use the standard Snapshot Manager to select the desired point.

Site Recovery Manager (SRM)

VR is a great product, but it requires administrator time and effort to setup and to failover each VM. The administrator needs to know what VMs to failover in what order to ensure applications start and run correctly. After a failover, an administrator must stop replication, set it up in the opposite direction, failover back to the original site, stop replication again, and finally reconfigure replication for the next disaster. All of this can be time consuming and doesn't work well if administrators are not available at the remote site to bring everything online again.

Site Recovery Manager (SRM) solves those problems by automating the entire process. This allows groups of VMs to be protected together, failing them over as a group in just a few clicks, and automating the process of reversing the replication direction after the disaster has passed.

SRM can use VR as previously described, but can also leverage array-based replication if desired, allowing for lower RPOs (potentially down to synchronous replication). Both can even be combined if desired, allowing critical VMs to be replicated via array-based replication and less critical VMs to leverage VR for example.

One of the main advantages of SRM, however, is the ability to test failovers to ensure the process will work as expected. This can be done without any disruption, allowing production to continue during the test. The copy on the remote site can be tested to ensure it works, but it can also be used to test upgrades and patches before they are applied in production to ensure they work properly or to try new versions of code on a copy of production data before applying them. In addition, all of these tests are fully logged and can be reported on to show compliance with company policies or for regulatory compliance reasons.

SRM can also be used for planned failovers to avoid a coming disaster, like in advance of a hurricane or flood. It can also be used to migrate VMs between datacenters, such as for datacenter consolidation projects. This allows administrators to test and verify things will work as expected, while ensuring no data is lost when it is actually failed over.

SRM is available as a standalone package or as part of the vCloud Suite. If purchased as a standalone package, it is available in two editions, Standard and Enterprise. The only difference is that Standard is limited to a maximum of 75 protected VMs per site, while there are no limitations with Enterprise. In both cases, it is licensed per protected VM. Licenses are sold in twenty-five-pack bundles. If purchased as part of the vCloud Suite, on the other hand, you need to purchase the Enterprise Edition of the suite (Standard and Advanced don't include SRM). In this case, it is licensed per socket, not per VM, and has no specific limits other than those imposed by SRM in general. One other note on licensing—it is only needed for the protected VMs on the protected site. No SRM licensing is required on the DR site (unless there are VMs on that site that need to be protected as well).

Future Directions

All of the content of this white paper so far has described features and functionality available today in vSphere 5.5. What is coming in the next version? Much of this is confidential (although anyone can sign up for the beta program and get the details by going to https://communities.vmware.com/community/vmtn/vsphere-beta and signing up for it). The information below was discussed at VMworld and is thus publicly available. Details are minimal as that is all that is currently available. Also note that this is about beta software, and features may be added, removed, or substantially changed when the product is actually released.

Long Distance vMotion

vMotion will allow for replication of much longer distances than today, potentially even cross country or across oceans. Specific bandwidth (probably 250 Mbps or more available) and round trip network latency (probably up to around 100 milliseconds) requirements must still be met, but the standards are much lower than that available today. Both networks will need to be in the same Layer 2 domain (i.e., IP subnet) as the IP address won't change.

In addition, Layer 3 networks will be supported for vMotion, allowing vMotion to be properly routed, whereas today the source and destination VMkernel ports used in vMotion must be in the same Layer 2 domain. This feature should also allow VMs to be moved from on-premises instances of vCenter into VMware-based clouds, such as the new vCloud Air (formerly vCloud Hybrid Service, or vCHS).

Cross-vCenter Migration

vCenter has always been a boundary for vMotion (in fact it is a boundary for everything in vSphere except SRM and VR). This restriction will be removed, allowing VMs to be replicated across any vCenter, though probably within the same Single Sign On (SSO) Domain only. Both networks (i.e., the network in each vCenter) will need to be in the same Layer 2 domain (i.e., IP subnet) as the IP address won't change.

Multi vCPU Fault Tolerance

Since FT was first introduced, it has been limited to single vCPU VMs. Thus, what is billed as a solution for your critical VMs is limited to VMs that only have one vCPU, which excludes the majority of VMs that most organizations consider critical. The next version of vSphere will support up to four vCPUs per FT VM. This has

been one of the most awaited new features for several versions of vSphere. Note that for this to work, the FT network must be 10 GB or faster.

Conclusion

This white paper has reviewed the various mechanisms to move a VM from one location to another, for normal purposes as well as for HA and DR purposes. It has also reviewed the licensing required for each method and some use cases and limitations of each.

Learn More

Learn more about how you can improve productivity, enhance efficiency, and sharpen your competitive edge through training.

- VMware vSphere: VCP5-DCV Prep Kit
- VMware vSphere: Install, Configure, Manage [V5.5]
- VMware vSphere: Fast Track [V5.5]
- VMware vSphere: Optimize and Scale [V5.5]
- VMware vSphere: Troubleshooting Workshop [V5.5]

Visit **www.globalknowledge.com** or call **1-800-COURSES (1-800-268-7737)** to speak with a Global Knowledge training advisor.

About the Author

John Hales (A+, Network+, CTT+, MCSE, MCDBA, MOUS, MCT, VCA-DCV, VCA-Cloud, VCA-Workforce Mobility, VCP, VCP-DT, VCAP-DCA, VCI, EMCSA) is a VMware instructor at Global Knowledge, teaching all of the vSphere and View classes that Global Knowledge offers. John has written a book called *Administering vSphere 5: Planning, Implementing, and Troubleshooting* published by Cengage, as well as other technical books—from exam-preparation books to quick-reference guides, as well as custom courseware for individual customers. John lives with his wife and children in Sunrise, Florida.