

Technology Backgrounder

Technology Backgrounder Live Migration and its Importance to VMware Cloud Partners



When moving VMware virtual machines to the cloud, the sure-fire way to migrate the VMs and their data completely is to simply stop the VMs, copy their components (OVF files) and assemble them into an Open Virtual Appliance (OVA). You transport the OVA package over the network or via a physical device to the cloud destination, unpack the files and restart the VMs. If you've done things right, the VMs pick up right where they left off.

Of course, it's not quite that easy. The physical resources and configurations in the cloud data center should be comparable to what you were running on in your on-premises data center. Network addresses and access permissions have to be properly set up in the cloud environment as well. Fortunately, VMware provides useful tools like vRealize to address that part of the cloud migration challenge.

The problem...

The problem, of course, is that in the scenario described above, the VMs and their applications are offline while they're being moved to their new home in the cloud. For some types of applications, that downtime may be perfectly acceptable. For other, however, that amount of downtime is simply unacceptable.

Enter live migration, again

The alternative is 'Live Migration,' which we might describe as any migration method that allows you to migrate your VMs' data to the destination in a background process while they're still running in the on-premises location. Done right, live migration can virtually eliminate the disruption to operations, as the VM isn't taken offline for the purpose of transferring its data.

• "Done right, live migration can virtually eliminate the disruption to operations..."

If you do a Google search on 'VMware cloud live migration,' you'll find plenty of good news. 'Live Migration' was a solved problem as long ago as 2014. Then it was solved again in 2015, and again in 2017, and... well, you get the picture.

Why is this? Well, 'live migration' is actually a general concept that encompasses a number of different approaches, by numerous vendors, to keeping an application running as its data is being copied to a cloud destination. And there's still a lot of room for improvement.

Last year's enterprise tools, this year's cloud challenges

There are a considerable number of tools available for live migration. Among the most commonly deployed data replication tools are software solutions based on legacy disaster recovery (DR) frameworks. It seems logical: Since the first step in providing DR for a virtual

machine is to create a remote copy of the VM's data, why not use that capability to replicate the VM's data for a live migration?

Because last year's enterprise tools are challenged by this year's cloud dynamics. For the most part, system administrators think that DR tools are overkill for cloud migration, and I have to agree. Here are some of their collective reasons:

- Complicated and expensive: DR tools were originally designed to be set up very precisely, with specifically defined policies, then left in operation for years. Used for live migration, they require pretty much the same amount of work to set up (usually by a well-paid professional services team), but then they're used for a one-off migration.
- Frustrating to plan and manage: As one admin said, "The number one question I'm asked is 'what's the timeframe for the migration project?' and I have to say, the tools I use don't really answer that question with any precision."

"If you take VMs offline, you can ship their data on a physical device, but DR tools used for live migration don't support device-based data transport."

- Only as good as the network they run on: Tools designed for DR assume a highly available network, running 24x7. In real-world cloud migrations, network connections may be slow or get the hiccups, or servers might occasionally restart. If the data replication process isn't fault tolerant, any interruption means starting data replication over again from the beginning.
- Can't be used with a data transport device: If you take VMs offline, you can ship their data on a physical device, but DR tools used for live migration don't support device-based data transport.
- Disruptive to performance: DR snapshots negatively impact application performance. If this is the expected natural overhead for a DR solution in regular use, administrators will have allocated compute and storage resources to offset the impact. However, if the data replication process for a cloud migration project suddenly introduces the performance hit, the customer may be very unpleasantly surprised.

Not every solution for cloud migration has every one of these drawbacks, but these are among the most common complaints from cloud service providers.

A 2018 approach to live migration

So let's set out objectives for 'live migration done right.' A manifesto for a new solution that leaves legacy architectures behind might look like this:

- Data replication does not disrupt or significantly impact the VMs and applications at the source. By definition, 'live migration' means that VMs remain in normal operation while data replication is in progress. The challenge is the definition of 'normal operation.' As we noted earlier, a newly introduced solution that begins taking snapshots will detrimentally impact application performance. Also, it might not be compatible with the IT organization's existing backup and DR tools and protocols. The answer is to replicate the data without requiring DR-style snapshots.
- Performance impact can be monitored and managed. Even without snapshots, replicating data in a background process will consume some system and network resources, potentially impacting application performance. An ideal live migration solution would give the administrator the ability to adjust the rate of data replication to keep performance impact below defined thresholds. The balance between replication duration and VM performance impact could be managed hands-on or automated with thresholds defined as a policy.

"An ideal live migration solution would give the administrator the ability to adjust the rate of data replication to keep performance impact below defined thresholds."

- The data replication mechanism behaves like a good guest should. Along the same lines of being non-disruptive, it's important that the deployment, configuration and uninstallation of any software be non-disruptive to ongoing operations and tools in both the on-premises data center and the cloud destination. Standard vSphere functions like vMotion, DRS, etc. should be available while data is replicated. Similarly, the use of third-party tools should not be compromised by the replication solution.
- Live migration works even if data is transferred via physical device. For larger projects, the migration solution should allow replication via a physical data transport device, with updated data transferred to the cloud over the network. Today, replicating data via device implies taking the VMs and their applications offline. Device-assisted data transfer should be an option for live migration as well.
- Data replication is scalable and fault-tolerant. Migration projects are a mix of hands-on oversight and automation. The more the administrator can automate data replication and define policy-based responses, the more efficient the migration process will be. Similarly, if

the replication process is interrupted – due to a break in network connectivity or a server restart, etc. – the replication should cleanly resume where it left off rather than restarting from the beginning.

• The data replication process is predictable and manageable. Before beginning data replication, the administrator will typically scope and blueprint the migration task. The data replication solution should be a part of the planning process as well, to assess the factors that will affect the replication process: the amount of data to be replicated, whether sets of VMs must migrate in groups, the amount of activity in the on-premises VMs, the network bandwidth and reliability, etc. Combined with the migration blueprint, this information will help the administrator plan the data replication process and forecast the time and resources required. The forecasting activity can even help the administrator weigh the pros and cons of device-based transport for data replication. If groups of VMs should migrate together, the solution should account for that in the planning process as well as the replication and change-over.

A good first impression

As we've seen, there is still frustration among cloud service providers who wish to move their enterprise customers' virtual machines to their cloud services without disruption, uncertainty and administrative overhead. There is a real need to take the friction out of cloud migration in terms of technology, operations and business relationships. In an ideal world, the migration to the cloud will be the cloud services provider's chance to make a good first impression.

"The solutions that can enable live migration to the cloud with the greatest ease and least disruption will prove to be among the most useful tools in the cloud provider's toolbox."

With the right tools, people and processes, they can start their relationship with their enterprise customers on the right foot. Live migration is essential for most enterprise workloads to move to the cloud, and cloud service providers see live migration as a key requirement to ease their enterprise customers' shift to their cloud services. The solutions that can enable live migration to the cloud with the greatest ease and least disruption will prove to be among the most useful tools in the cloud provider's toolbox.