

Cisco Virtualized Workload Mobility Introduction

The ability to move workloads between physical locations within the virtualized Data Center (one or more physical Data Centers used to share IT assets and resources) has been a goal of progressive IT departments since the introduction of virtualized server environments. One of the main drivers for the continued increased virtualized server concept being adopted is really the control it gives the IT department because now the IT department has the flexibility and agility that comes with decoupling the Hardware and Software from one another. Ultimately, this means new use cases can be developed addressing specific requirements and adapted to a specific environment. In this virtualized environment, the ability to dynamically and manually redistribute virtual machines to new locations provides the opportunity for value-added and operationally efficient new use cases.

In the virtualized environment, many different use cases have been discussed at times, including the following:

- Data Center Capacity Expansion and/or Consolidation

Applications need to be migrated from one data center to another without business downtime as part of a data center migration, maintenance or consolidation efforts. Virtual machines need to be migrated to a secondary data center as part of data center expansion to address power, cooling, and space constraints in the primary data center

- Virtualized Server Resource Distribution over Distance

Virtual machines need to be migrated between data centers to provide compute power from data centers closer to the clients (follow the sun) or to load-balance across multiple sites.

- Disaster Planning Strategies, including Disaster Avoidance Scenarios

Data centers in the path of natural calamities (such as hurricanes) needs to proactively migrate the mission-critical application environment to another data center.

Individually, each of these use cases is addressing specific needs of the business. Taken together, these use cases aggregate into one clearly identifiable use case: Virtualized Workload Mobility (VWM).

Virtualized Workload Mobility is a vendor and technically agnostic term that describes the ability to move virtualized compute resources to new, geographically disperse, locations. For example, a computer Operating System (as examples: Red Hat Linux or Windows 7) that has been virtualized on a virtualization host Operating System (as examples: VMware ESXi, Windows Hyper-V, or Red Hat KVM) can be transported over a computer network to a new location that is in a physically different data center.

Additionally, Virtualized Workload Mobility refers to active, hot, or live workloads. That is, these virtualized workloads are providing real-time compute resources before, during and after the virtualized workload have been transported from one physical location to the other. Note that this live migration characteristic is the primary benefit when developing the specific use case and technical solution and therefore must be maintained at all times.

Virtualized Workload Mobility Use Case

The Virtualized Workload Mobility use case requirements include the following:

- **Live Migration:** The live migration consists of migrating the memory content of the virtual machine, maintaining access to the existing storage infrastructure containing the virtual machine's stored content, and providing continuous network connectivity to the existing network domain.
- **Non-disruptive:** The existing virtual machine transactions must be maintained and any new virtual machine transaction must be allowed during any stage of the live migration.
- **Continuous Data Availability:** The virtual machine volatile data will be migrated during the live migration event and must be accessible at all times. The virtual machine non-volatile data, the stored data typically residing on a back-end disk array, must also be accessible at all times during the live migration. The Continuous Data Availability use case requirements have a direct impact on the storage component of Data Center Interconnect for the Virtualized Workload Mobility use case and will be addressed further.

With the Virtualized Workload Mobility use case described and the main use case requirements identified, there is also a need to identify the challenges, concerns and possible constraints that need to be considered and how they impact the solution development and implementation.

Use Case Challenges, Concerns & Constraints

The use case challenges, concerns and constraints can be derived by a systematic review of the use case requirements.

- **Live Migration:** With the Live Migration requirement, a virtualization host operating system must be identified and integrated into the solution. Vendor's that provide virtualization host operating systems include VMware, Red Hat and Microsoft, and they all have their own technical and supportability requirements. This means that the solution will have to consider the selected vendors requirements, which may impact the design and implementation of the Virtualized Workload Mobility solution. Virtualization Host OS selection becomes a critical component in the development.
 - **Criteria 1:** Live Migration 5 millisecond round trip timer requirement



Note

There are additional Live Migration requirements that are addressed inherently in this solution.

- **Non-Disruptive:** The Non-Disruptive requirement means that we have to consider both existing server transactions and new server transactions when designing the architecture. Virtual machine transactions must be maintained so the solution must anticipate both obvious and unique scenarios that may occur.
 - **Criteria 2:** Maintain virtual machine sessions before, during and after the live migration.
- **Continuous Data Availability:** The non-volatile memory (storage) must be accessible wherever the virtual machine is located in the virtualized data center. In the event of a virtual machine migration, the stored content must continuously be accessible and this impacts the Virtualized Workload Mobility solution with respect to how the storage infrastructure is specifically designed and integrated into the solution. Additionally, there are many different storage topologies and products to choose from, so the solution must be flexible and adaptable to different storage scenarios.
 - **Criteria 3:** Maintain storage availability and accessibility with synchronous replication
 - **Criteria 4:** Develop flexible solution around diverse storage topologies

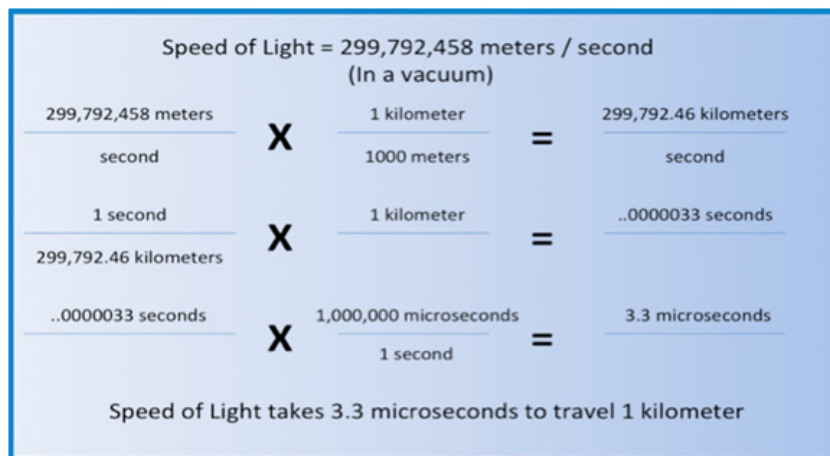
These 4 specific items help more narrowly define the Virtualized Workload Mobility use case to be specific in the design and deployability. We'll review each criterion in detail and how exactly it provides clarity into the use case.

Criteria 1: Live Migration 5 millisecond round trip timer requirement

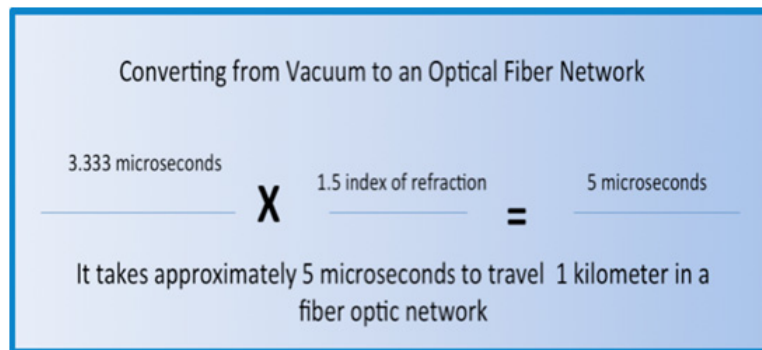
In this solution, we have selected the VMware ESXi 4.1 host operating system, due to its vMotion technology and widespread industry adoption. Please note that new and improved ESXi versions, when available, as well as additional Vendor Host Operating Systems, will be integrated into the solution. One of VMware's vMotion requirements includes that the ESXi source and target host OS's must be within 5 milliseconds, round-trip, from one another. This 5 ms requirement has practical implications in defining the Virtualized Workload Mobility use case with respect to how far the physical data centers may be located from one to the other and we must translate the 5 milliseconds Round-Trip response time requirement into distance between locations.

First thing to note is that the support requirement is a 5 ms round-trip response time so we calculate the distance between the physical locations using 2.5 milliseconds as the one way support requirement. In a vacuum, the speed of light travels 299,792,458 meters per second and we use this to convert to the time it takes to travel 1 kilometer.

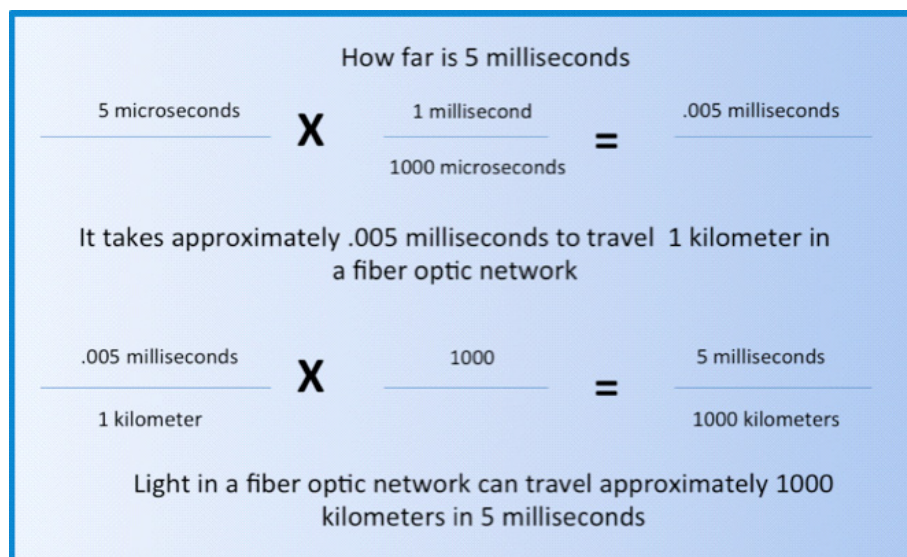
Figure 1-1 Light Speed Conversion to Latency



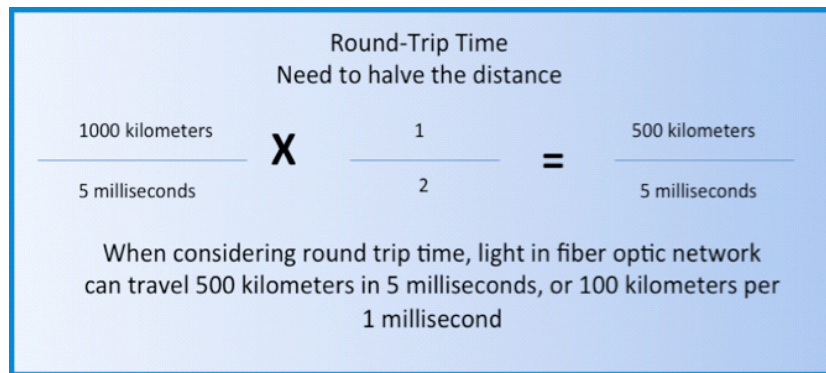
The speed of light, when light is travelling in a vacuum, takes approximately 3.3 microseconds to travel 1 kilometer. However, the signal transmission medium on computer network systems is typically a fiber optic network and there is an interaction between electrons bound to the atoms of the optic materials that impede the signal, thereby increasing the time it takes for the signal to travel 1 kilometer. This interaction is called the refractive index.

Figure 1-2 Practical Latency in a Fiber Optic Network

Knowing that it takes approximately 5 microseconds per kilometer in a fiber network, the calculation to determine the distance of 5 milliseconds is straightforward.

Figure 1-3 How Far is 5 Milliseconds

Finally, we must consider that the Criteria 1 requirement is a “round trip” timer. That is, the time between sites includes both the time to the remote data center and the return time from the remote data center. In short, the round-trip distance that the speed of light over a fiber optic network can traverse is actually half of the one-way direction. Therefore, the round trip distance for 5 milliseconds is approximately 500 kilometers in a fiber optic network, which can be reduced to 100 kilometers per 1 millisecond.

Figure 1-4 Round-Trip Distance

However, and this is critical, there are many additional factors that add latency, time, into a computer systems network and this increased time must be considered. Unfortunately, this added systems latency is unique to each environment and is dependent of the signal degradation, quality of the fiber or layer 1 connection, and the different elements, that are intermediate between the source and destination. In practice this distance is somewhere between 100 and 400 kilometers. The Virtualized Workload Mobility use case considers the practical range of 5 milliseconds to be between 100 and 400 kilometers and we highly recommend that you evaluate your environment to define this further.

Criteria 2: Maintain virtual machine performance before, during and after the live migration.

The objective with Criteria 2 is to that ensure service levels and performance are not impacted when executing a mobility event when monitoring traffic flows. Note that virtual machine and application performance are dependent on many factors, including the resources (CPU and Memory) of the host (physical server) and the specific application requirements and dependencies. It should not be expected that all virtual machines and applications would maintain the exact same performance level during and after the mobility event due to the variable considerations of the physical server (CPU and Memory), unique application requirements, and the additional virtual machines utilization of that same physical server. However, given the same properties in both data centers, it is reasonable to expect similar performance during and after the mobility event as before the mobility event.

Criteria 3: Maintain storage availability and accessibility with synchronous replication

One of the key requirements as expressed in field research is the ability to have data accessible in all location in the virtualized data center at all times. This means that stored data, storage, must be replicated between locations, and further, the replication must be synchronous. Assuming two data centers, synchronous replication requires a “write” to be acknowledged by both the local data center storage and remote data center storage for the input/output (I/O) operation to be considered complete (and possibly preventing data inconsistency or corruption from occurring). Because the write from the remote storage must be acknowledged, the distance to the remote data center becomes a bottleneck and factor when determining the distance between data center locations when using synchronous replication.

Other factors that affect synchronous replication distances include the application Input/Output (I/O) requirements and the read/write ratio of the Fibre Channel Protocol transactions. Based on these variables, synchronous replication distances range from 50-300 kilometers. This solution uses the synchronous replication distance at approximately 100 kilometers.

Criteria 4: Develop flexible solution around diverse storage topologies

One of the key drivers in developing a deployable solution is to ensure that it can be inserted into both greenfield and brownfield deployments. That is, the solution is flexible enough that it can be deployed into systems that have unique or different characteristics. Specifically, the solution must be adaptable to different Storage Networking systems and deployments. Two different storage topology considerations include storage located only in the primary data center and intelligent storage delivery systems.

- **Local Storage** refers to the storage content located at the primary data center. After a virtual machine is migrated to the new data center, the stored content remains in the primary data center. Subsequent read/writes to the storage array will have to traverse the Data Center Interconnect and / or Fibre Channel network. This not an optimal scenario, but one to consider and understand when developing the Virtualized Workload Mobility use case.
- **Intelligent Storage** over distance refers to using specialized hardware and software to provide available and reliable storage content. Storage vendors provide both optimized industry protocol enhancements as well as value added features on their hardware and software. These storage solutions increase the application performance and metrics when used in a Virtualized Workload Mobility environment.

The solution will introduce both a local storage and an intelligent storage solution so a baseline metric can be analyzed against a metric when using the intelligent storage solution.

Use Case Criteria Summary


To summarize, the use case requirements and criteria provide the boundary when defining Virtualized Workload Mobility. Below is a list of the main points:

- 100 kilometers between data centers
- Synchronous replication for continuous data availability
- Two different storage solutions
 - Local Storage solution
 - Intelligent Storage solution
- Maintain Application Performance before, during and after Live Migration

Data Center Interconnect Solution Overview

The term DCI (Data Center Interconnect) is relevant in all scenarios where different levels of connectivity are required between two or more data center locations in order to provide flexibility for deploying applications and resiliency schemes.

[Figure 1-5](#) summarizes the three general types of connectivity required for a DCI solution.

Figure 1-5 DCI Connectivity Overview

DCI Components

Four critical functional components are addressed in this DCI solution, as follows:

- **LAN Extension:** Provides a single Layer 2 domain across data centers. The data center applications are often legacy or use embedded IP addressing that drives Layer 2 expansion across data centers. Layer 2 Extension provides a transparent mechanism to distribute the physical resources required by some application frameworks such as the mobility of the active machine (virtual or physical).
- **Path Optimization:** deals with the fact that every time a specific VLAN (subnet) is stretched between two (or more) locations that are geographically remote, specific considerations need to be made regarding the routing path between client devices that need to access application servers located on that subnet. Same challenges and considerations also apply to server-to-server communication, especially for multi-tier application deployments. Path Optimization includes various technologies that allow optimizing the communication path in these different scenarios.
- **Layer 3 Extension:** Provides routed connectivity between data centers used for segmentation/virtualization and file server backup applications. This may be Layer 3 VPN-based connectivity, and may require bandwidth and QoS considerations.
- **SAN Extension:** This presents different types of challenges and considerations because of the requirements in terms of distance and latency and the fact that Fibre Channel cannot natively be transported over an IP network.


In addition to the four functional components listed, a holistic DCI solution usually offers a couple more of functions. First, there are Layer 3 routing capabilities between Data Center sites. This is nothing new, since routed connectivity between data centers used for segmentation/virtualization and file server backup applications has been developed for quite some time. Some of the routed communication may happen leveraging the physical infrastructure already used also to provide LAN extension services; at the same time, it is common practice to have also L3 connectivity between data centers through the MAN/WAN routed network.

The second aspect deals with the integration of network services (as FW, load-balancers). This also represents an important design aspect of a DCI solution, given the challenges brought up by the usual requirement of maintaining stateful services access while moving workloads between data center sites.

Virtualized Workload Mobility Solution Overview

DCI enables the Virtualized Workload Mobility use case. The DCI system has four main components that leverage existing products and technologies and is modular in design, allowing for DCI system to support new and evolving use cases. For the Virtualized Workload Mobility use case, the solution components are listed below as represented in [Figure 1-6](#).

Figure 1-6 Virtualized Workload Mobility Solution Components



LAN Extension: The LAN Extension technology extends the Layer 2 system beyond the single data center and connects 2 or more data centers via a single Layer 2 domain. This use case highlighted two different LAN Extension technologies:

- Overlay Transport Virtualization (OTV) using the Nexus 7000
- Virtual Port Channels (vPC) using the Nexus 7000

Path Optimization: As virtual machines are migrated between data centers, the traffic flow for the client-server transaction may become suboptimal. Suboptimal traffic flow may lead to application performance degradation.

- Ingress Path Optimization (traffic entering the data center): Application Control Engine (ACE) and the Global Site Selector (GSS) with workflow integration with VMware vCenter
- Egress Path Optimization (traffic leaving the data center): First Hop Redundancy Protocol (FHRP) using HSRP Localization on the Nexus 7000

Workload Mobility: Workload Mobility requires a virtualized infrastructure to support migration of virtual machines.


- Nexus 1000v (N1KV) distributed virtual switch in a stretched cluster mode. Please see [Figure 1-7](#)

**Note**

The mobility solution was also designed with consideration to live migration between two different clusters.

- Virtual Security Gateway (VSG) to provide protection and isolation of virtual machines before, during and after the migration event
- Virtual Network Management Center (VNMC) for VSG provisioning
- VMware ESXi Host Operating Systems to enable and support virtual machine instances

Figure 1-7 *Stretched Cluster - A Single ESXi Cluster Stretched Across 100 kilometers*



Storage Elasticity: This use case requires synchronous replication with flexibility to integrate intelligent storage solutions.

- MDS 9000 for Fibre Channel VSAN connectivity
- Intelligent Storage Solutions from NetApp

A key point to recognize is that the solution components are designed to be independent from one another. That is, new technologies and products can replace existing products and technologies within a specified component.

