

HP Virtual Connect and VMware Infrastructure 3



Introduction	2
Virtualization in the enterprise	2
HP Virtual Connect	2
VMware Infrastructure 3	4
HP Virtual Connect and VMware: better together	4
Usage scenarios	4
New server deployment	4
Rapid recovery	6
Functional testing	9
Summary	11
For more information	12

Introduction

In the past, there were only two ways to connect server blades to outside networks and storage—switches and pass-through devices—each with its own advantages and compromises. Pass-through devices were simple, but required many failure-prone cables and complex cable management processes. Blade switches, on the other hand, reduced the number of cables needed, but added components for Local Area Network (LAN) and Storage Area Network (SAN) administrators to manage. In both cases, several people were needed to perform very common and what should be very simple, server tasks. With either pass-through devices or switches, adding or replacing a server required coordination of the storage, network and server administrators to complete the change successfully.

HP worked with users across the world and brought together server, network and storage administrators to find a better answer to these traditional approaches. What would it take to pre-provision your entire infrastructure once, then add or replace servers on-the-fly? What if you could take a server's LAN or SAN connection profile and simply apply it to another server automatically? What if, for every new server you added or replaced, the infrastructure was ready to go when you plugged it in and powered it up? And what if you could do it all with compatibility and familiar industry standards? The answers were unanimous and all pointed to Virtual Connect.

HP Virtual Connect is designed to enable a new way to add, move or replace server blades – without impacting networks or requiring multiple experts at each step. Virtual Connect is a major innovation for customers and their businesses. Virtual Connect provides a better way for IT organizations to work together and provides these benefits to all parties involved without the traditional compromises.

Virtualization-savvy customers could see an immediate compounding of functionality between Virtual Connect and VMware Infrastructure 3. Customers using VMware Virtual Infrastructure products already realize the power of flexible infrastructure; they are already realizing the benefits of shedding the traditional constraints that tie an operating system to a physical system. With the addition of Virtual Connect, administrators and architects could extend this flexibility throughout their entire application environments, from the fabric to the application.

This paper is a collaboration between Hewlett-Packard and VMware to demonstrate common usage scenarios, explaining how these two complementary technologies can deliver the most complete and most agile datacenter available in an industry-standard environment.

Target audience: This paper is intended for IT decision makers and administrators seeking to reduce costs and improve manageability while increasing flexibility within their datacenters. This paper assumes an understanding of the basic principles of VMware virtualization technology and products.

Virtualization in the enterprise

HP Virtual Connect

HP Virtual Connect makes IT change-ready by virtualizing physical connections – both those between the server and the SAN and those between the server and the LAN. Change may be necessitated by system failures, a physical server migration or even the simple addition of a new server. In a traditional IT environment, change takes time and coordination between server, SAN and LAN administrators. The time to make required changes may take hours, days or even weeks to scope the change, coordinate the change plan and implement. Virtual Connect eliminates this barrier to change.

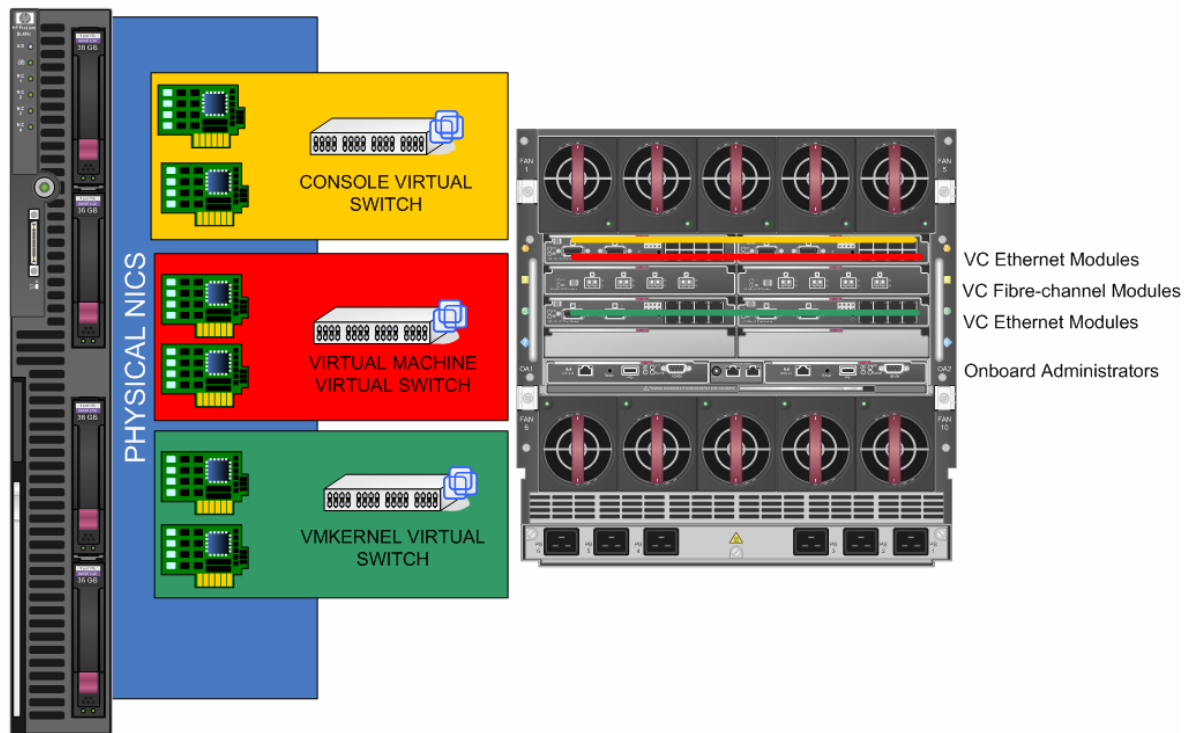
Virtual Connect interconnect options include a 1/10Gb Ethernet module and a 4Gb Fibre Channel module. These modules provide an abstraction layer between the server and both the LAN and SAN fabrics. During initial setup, Virtual Connect allows a server administrator to define a series of profiles

that can be tied to the server bays, instead of the physical servers, within the HP BladeSystem c7000 enclosure. Virtual Connect assigns Ethernet Media Access Control (MAC) addresses and Fibre Channel Worldwide Names (WWN) to these bays. This allows the server administrator to provide LAN and SAN administrators with all potential MAC addresses and WWNs before any servers are even inserted. LAN and SAN administrators can then configure the Ethernet and Fibre Channel access policies and, still, no servers have been inserted. Then, finally, when a server is inserted, the physical MAC addresses and WWNs of the local adapters are replaced with Virtual Connect assigned IDs, defined in the slot profile. The operating system, once booted, sees and utilizes these Virtual Connect assigned IDs. These IDs remain consistent within the enclosure bay, even if a server is replaced – there are no changes required on the LAN or SAN access policies or cabling.

The Virtual Connect modules are not switches; rather the Virtual Connect modules aggregate the LAN and SAN connections running from each of the blade servers in the enclosure. External communications are then routed over a few uplinks on the Virtual Connect Ethernet and Fibre Channel modules to the first layer of switches in the LAN and SAN managed networks. This BladeSystem architecture eliminates at least one required layer of switch management and also simplifies cabling for the enclosure.

With many blade environments, compromise is the name of the game when it comes to virtualization. Best practices suggest numerous network interface cards, while many blades are limited in the number of NICs they can provide. With HP BladeSystem c-Class and Virtual Connect, up to 12 NICs can be utilized within a single server while still maintaining redundant connections to Fibre Channel. Thus, in a VMware Infrastructure 3 environment, full redundancy across the Service Console, Virtual Machine and VMkernel networks can be achieved as outlined in Figure 1.

Figure 1. Network mapping between VMware Infrastructure 3 and HP Virtual Connect.



In a traditional environment, this series of redundant connections per server means numerous switch ports to manage. For example, 16 rack mount servers with 12 network interface cards each would

require 192 network ports. With HP Virtual Connect, an enclosure of eight full-height servers with 12 NICs and a full complement of Virtual Connect modules becomes a single logical device to configure. This environment can be made completely redundant with as few as 8 cables linked to the first layer of managed switches.

VMware Infrastructure 3

VMware Infrastructure 3 is a state-of-the-art server virtualization and virtualization management platform. VMware Infrastructure 3 consists of several components, including VirtualCenter, VMotion, Virtual SMP and ESX Server. Together, the suite delivers comprehensive virtualization management, resource optimization, application availability and operational automation capabilities. However, this paper will focus only on the components that are directly impacted by the presence of Virtual Connect – the operation and configuration of the thin, hypervisor operating system, VMware ESX Server, and the management interface, VirtualCenter.

ESX Server allows administrators to create, run and manage multiple virtual machines on a single server host. Within the virtual machines, the operating system remains unchanged but accesses hardware resources through device abstractions that are separate from and independent of the underlying physical hardware. In a virtualized environment, the traditional rigid linkages between operating systems and hardware are replaced by hardware-independent, flexible device abstractions. By replacing these rigid bonds that are device-specific, VMware virtualization technology allows for hardware-independent backup and recovery, datacenter mobility and unprecedented platform flexibility.

Flexibility is the key to efficiency, and VMware Virtual Infrastructure combined with HP Virtual Connect offers this flexibility. The combination of these technologies leverages two powerful layers of abstraction that introduce flexibility at two critical points in an application environment. Virtual Connect adds flexibility between the physical server and the I/O fabrics and VMware ESX Server brings flexibility between the virtual host operating system and the physical server hardware. With abstractions below the server – at the fabric level – and above the server – at the OS level – the physical server components of the datacenter become dynamic resources that can respond to business needs without requiring costly and time-consuming reconfiguration or redeployment operations.

HP Virtual Connect and VMware: better together

To demonstrate how HP Virtual Connect adds value in a VMware Infrastructure 3 environment, HP and VMware selected common usage scenarios from a virtualized datacenter. These usage scenarios included the processes of adding a server with pre-provisioned network and storage connections and rapidly recovering a failed server through a “rip and replace” operation – both without ever changing its physical connections or reconfiguring LAN and SAN switches.

Usage scenarios

New server deployment

In a traditional environment, the addition of a new server requires the coordination of SAN, LAN and server administrators. In order to bring a host online, whether to start new service or expand capacity to an existing service, a server must be ordered and delivered, the WWNs of any Fibre Channel host bus adapters and MAC addresses of any network interface cards must be inventoried and passed to the SAN and LAN administrators, and the administrators must prepare the datacenter infrastructure to accept the new server.

Even after the access policies are in place, adding the new server still must wait on the creation of LUNs and the proper cabling of networks. This process must be repeated for each server that is

added. To make matters worse, as capacity requirements expand, time to fulfill these requirements naturally increases due to the additional complexity.

Figure 2. HP BladeSystem configuration used for the new server deployment usage scenario.



With HP Virtual Connect, planning and deployment are a one-time process for the SAN and LAN administrators. All server profiles are defined at the Virtual Connect Domain level. WWNs and MAC addresses are available to the SAN and LAN administrators, even before the servers arrive. This allows all 16 server bays of an HP BladeSystem c7000 enclosure to be preconfigured, not only for initial deployment, but also for future growth with no further involvement from the SAN and LAN groups.

The new server deployment usage scenario refers to the addition of a new HP ProLiant c-Class blade server to a chassis and the installation and configuration of VMware ESX Server. This is a very common and, often, “first step” in understanding the joint solutions enabled with Virtual Connect and ESX Server. This usage scenario also covers the addition of a server for the expansion or creation of a VirtualCenter-managed resource pool. In both of these cases, the key point to the new server deployment usage scenario is that the profile and configuration of the server being added is both unique and new. This usage scenario does not cover the replacement or migration of an existing physical system.

The specific intent in testing this usage scenario using HP Virtual Connect is to determine if a server administrator can, in a preconfigured Virtual Connect domain, insert and configure a virtualization host without needing to change settings at the SAN and LAN levels. This test was carried out during initial enclosure population. The steps involved for this test were:

1. Create and apply Virtual Connect profile to chassis slot
2. Power-on server and install VMware ESX Server
3. Configure ESX Server host and discover the SAN-based VMware File System (VMFS) volumes

This usage scenario would be deemed successful if the following conditions were met:

- Networks were already available to the host and communication to the service console, virtual machines and VMkernel networks required no administrative assistance or LAN changes

- SAN LUNs were already visible to the hosts post-install without the need to modify zoning, cabling or LUN presentation

Four Virtual Connect modules were inserted into the enclosure. Switch bays 1 and 2 housed Virtual Connect Ethernet modules, while switch bays 3 and 4 housed the Virtual Connect Fibre Channel modules. The Virtual Connect Manager was used to configure a domain using Virtual Connect assigned MAC addresses and WWNs. Server profiles were defined for bays 1-6 and 9-14. Two networks were defined. The first was called *consolevmkernel* and would be reserved for the Service Console and VMkernel networks. The second network was called *vms* and would serve as the network for virtual machine communication. Each network was mapped out of a Virtual Connect module uplink in switch bays 1 and 2. Two SAN fabrics were assigned, and a SAN administrator presented Virtual Connect assigned WWNs to an HP StorageWorks 4000 Enterprise Virtual Array (EVA4000) SAN prior to any servers being inserted into the enclosure. The Fibre Channel modules were connected to an external HP B-Series 2/8 switch with N_port_ID virtualization (NPIV) active on all ports. A single cable per module was connected and all connections from the HP StorageWorks EVA4000 were connected to the same switch.

Two servers were inserted in slots 9 and 11 of the HP BladeSystem c7000 enclosure. The HP Integrated Lights-Out (iLO) of each blade was used to install VMware ESX Server 3.0.1 using iLO virtual media. Once installed, a cluster named "VirtualConnect" was created in VirtualCenter and the two hosts were added to the cluster. Once licensed, networking was configured for each host and, as expected, all networks were available and no LAN administrator assistance was required.

Under Storage Configuration in VirtualCenter, the Add Storage link was clicked and, after selecting to add a new LUN, all LUNs presented during initial domain setup were visible to the hosts and were added rapidly.

In this usage scenario, the amount of time it takes to add an ESX Server host to an environment is dramatically reduced. Virtual Connect allows for the presentation of Ethernet networks and SAN fabrics to hosts before they are ever brought online. This pre-provisioning of WWNs, MAC addresses and associated connectivity ensures that as capacity demand expands, the involvement of the SAN and LAN administrators is not needed after initial configuration.

Rapid recovery

The rapid recovery usage scenario refers to the replacement and reconfiguration of a server blade to replace an existing system that is being re-provisioned or has failed. In this case, the goal is to add the replacement host and restore the physical resources by assuming the identity of the original physical host. In this case, neither the profile nor the configuration is new; both were previously applied to the original host. This is a common usage scenario that demonstrates, very clearly, the power of the abstracted Virtual Connect and ESX Server interfaces.

Figure 3. Physical configuration of the systems used for the rapid recovery usage scenario.



The server profile for this test was configured to present any server in the bay with a network connection to both the Service Console and Virtual Machine networks. Boot from SAN parameters were entered during profile creation using the Array Node WWN provided by the SAN administrator during initial setup. Figure 4 details the SAN parameters used in the Virtual Connect profile.

Figure 4. FC SAN connections defined in the test server profile.

FC SAN Connections (Physical ports)

Port	Connected To	FC SAN Name	Status	Port Speed	WWPN
1	Bay 3	SAN_3	OK	Auto	50:06:0B:00:00:C2:62:08
2	Bay 4	SAN_4	OK	Auto	50:06:0B:00:00:C2:62:0A

Show Fibre Channel Boot Parameters

Port	SAN Boot	Target Port Name (WWPN)	LUN
1	Primary	50001fe15007dad8	1
2	Use BIOS		0

Assign Profile to Server Bay

Power	Server Bay Assignment	SN	Model	Status	UID
On	AD-U1100-BLENC: Bay 3 (ProLiant BL465c G1) Select Location	USM64903C1	ProLiant BL465c G1	OK	

The specific intent in testing this usage scenario using HP Virtual Connect is to determine if a server administrator can, in a preconfigured Virtual Connect domain, replace an ESX Server host without needing to change settings at the SAN and LAN levels and have the replacement host assume the role and identity of the server that was replaced. The steps involved for this test were:

1. Create and apply Virtual Connect profile to chassis slot.
2. Invoke a simulated total system failure by removing power to the server.
3. VMware HA (High Availability) should redistribute the virtual machines from the failed host to the other host in the cluster.
4. The failed blade is replaced with an un-configured blade.
5. The system should, with no administrative involvement, boot to the Virtual Connect assigned LUN, assume the identity and configuration of the previous host and appear in the cluster as the failed server. Once booted and rejoined into the cluster, VMware DRS (Distributed Resource Scheduler) should redistribute some virtual machines to the new, replacement host.

The test environment for the rapid recovery usage scenario differed slightly from the server deployment scenario. While switch bays 1 and 2 continued to house Virtual Connect Ethernet modules and switch bays 3 and 4 housed the Virtual Connect Fibre Channel modules, three servers were used – one placed in server bay 3, another in server bay 11 and the third was held aside as the standby replacement server. Prior to installing the servers, the Virtual Connect Manager was used to configure a domain using Virtual Connect assigned MAC addresses and WWNs. Virtual Connect server profiles were defined for bays 3 and 11. In the profile, two networks were defined; the first was called *consolevmkernel*, which would be reserved for the Service Console and VMkernel networks, and the second network, called *vms*, would serve as the network for virtual machine communication. Each network was mapped out of a Virtual Connect module uplink in switch bays 1 and 2. Two SAN fabrics were assigned and a SAN administrator presented Virtual Connect assigned WWNs to an HP StorageWorks EVA4000 SAN prior to any servers being inserted into the enclosure. The Fibre Channel modules were connected to an external HP B-Series 2/8 switch with NPIV active on all ports. A single cable per module was connected and all connections from the HP StorageWorks EVA4000 were connected to the same switch.

Two servers were inserted into the HP BladeSystem c7000 enclosure – one in slot 3 and the other in slot 11 – the additional, spare blade was prepared as a replacement, but not inserted into the enclosure. As mentioned, the Virtual Connect profile for bay 3 configured the server for boot from SAN and presented parameters for a boot LUN. Using iLO Virtual Media, VMware ESX Server 3.0.1 was installed on the SAN-based boot LUN of the host. After ESX Server installation, the hosts were added to VirtualCenter and a cluster named “VirtualConnect” was created. In the host configuration, the existing VMFS volume was added to the storage configuration of both hosts through the Storage Configuration option in the Virtual Infrastructure Client user interface. Clicking the Add Storage link, all LUNs presented during the definition of SAN access policies should be visible to both hosts. The VMFS volumes were added to the hosts, allowing both hosts to access the virtual machines stored on the volume. Then, HP added both the powered-on hosts and 6 Microsoft® Windows® Server 2003 Enterprise Edition virtual machines to the cluster and powered on both the hosts and the virtual machines. VMware DRS was allowed to place each of the virtual machines (VMs) between the hosts, as the VMs were powered on.

After a steady state was reached, the server in bay 3 was pulled from the enclosure to simulate a catastrophic failure. The test engineers waited for VMware HA to recover the virtual machines on the remaining host. Then, the host that was removed from the c-Class enclosure was replaced with the spare, un-configured server.

Because the Virtual Connect modules were managing the physical Fibre Channel and Ethernet addresses, the host was able to see the boot LUN exposed to the WWN of the enclosure bay and boot with the configuration of the previous host. As the host boots, the network connectivity is already in place to allow the ESX Server configuration stored on the boot LUN to restore network connectivity

to VirtualCenter and re-associate the ESX Server host with VirtualCenter and its associated VMware HA and VMware DRS clusters. After the system has reestablished its role within the cluster, DRS automatically distributes the load to the replacement host, as though the replacement were the original system.

This usage scenario is an ideal example of the level of flexibility created in the datacenter by combining VMware Infrastructure 3 and HP Virtual Connect. Without any reconfiguration whatsoever, test engineers were able to replace a physical system and have that system completely restore the resources of the failed host. After the initial configuration, there were no changes required in VirtualCenter, no changes in the SAN fabric, no changes in the network configuration or connectivity and no changes required in the applications. The pre-provisioning of SAN volumes and Ethernet segments, based on Virtual Connect managed profile addresses, combined with the hardware independence and automation of Virtual Infrastructure allowed the system to recover and restore without any administrative changes.

Functional testing

In addition, general testing and validation of core VMware Infrastructure 3 functionality was conducted. This testing included ensuring the expected behavior of:

- VMware VMotion
- VMware Distribute Resource Scheduler (DRS)
- VMware High Availability (HA)

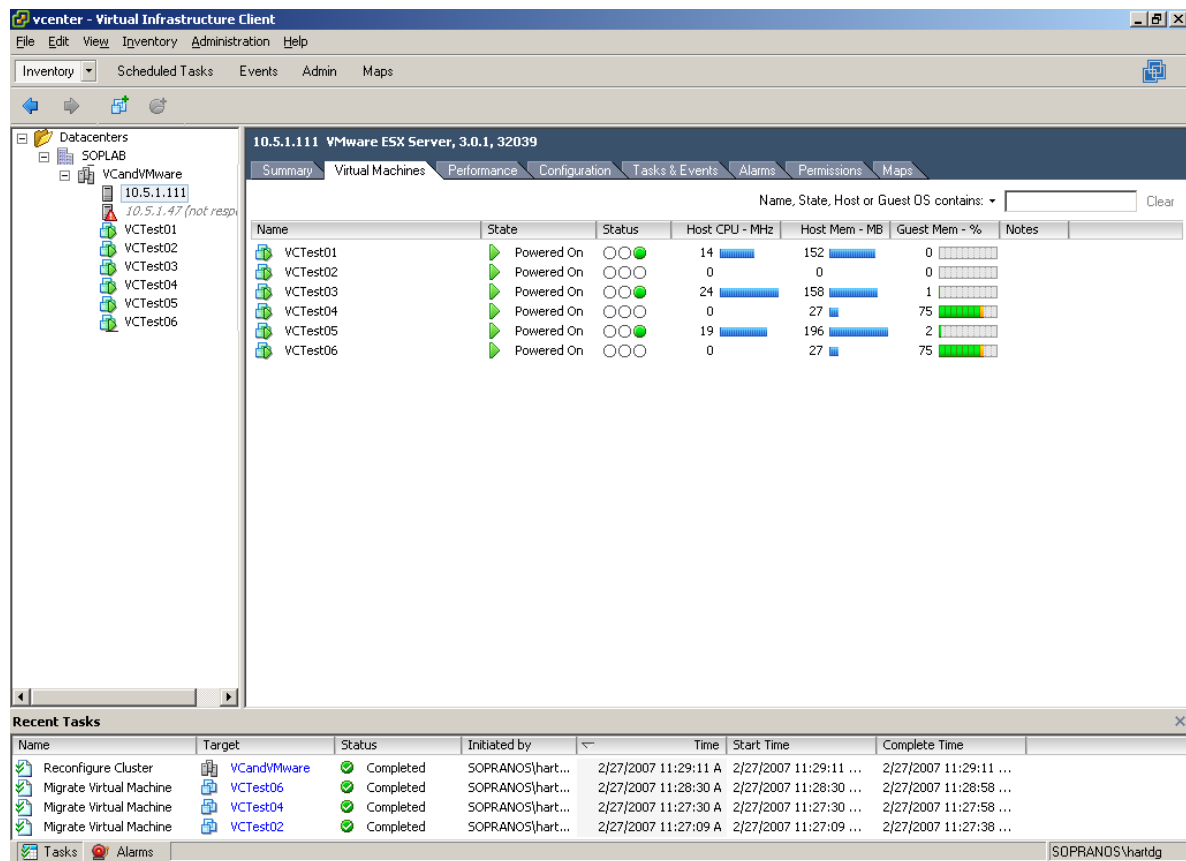
To test these core features, a cluster was configured within VirtualCenter with two HP ProLiant BL465c servers. DRS and HA were not enabled during creation. Ten virtual machines were created on two different datastores. Five virtual machines were assigned to each server.

To test VMotion, 1 virtual machine was migrated between host 1 and host 2. After a brief resting period, the server was migrated back. Both migrations completed successfully.

In order to test Distributed Resource Scheduler (DRS), the team used a CPU load tool to create load within two VMs on ESX Server host number 1 (VCTest1-2K3 and VCTest3-2K3). DRS was activated within the cluster, set to Fully Automatic and the utility was launched from the command line of both virtual machines. The expected behavior would be for the scripts to generate sufficient load to trigger a DRS rebalancing of load across hosts and for VirtualCenter to use VMotion and separate these virtual machines onto separate hosts. This behavior was confirmed with both the Virtual Connect Ethernet and Fibre Channel modules in place and managing the ESX Server physical addresses.

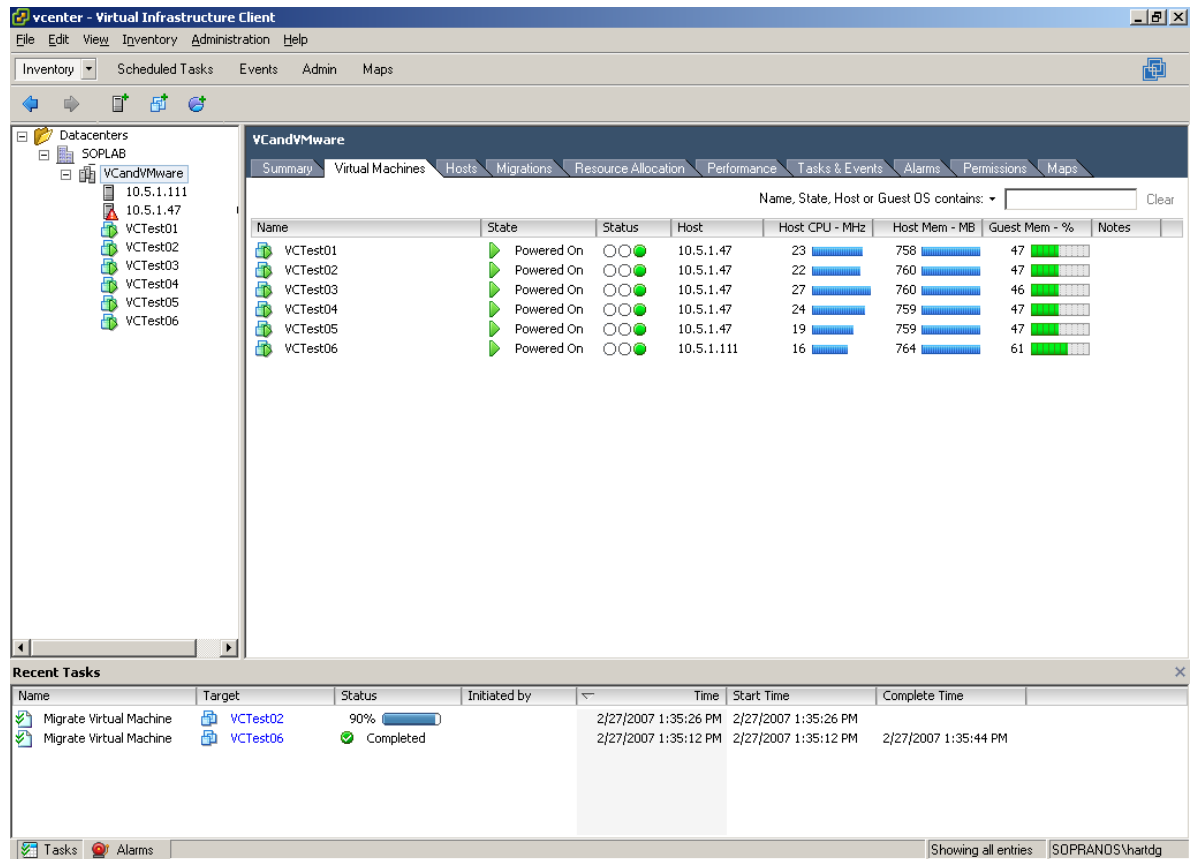
To test VMware HA, the feature was activated at the cluster level. The iLO of ESX Server host 2 was used to power the server down abruptly. The virtual machines from host number 2 were successfully restarted on host 1 yielding six active virtual machines as shown in Figure 5. Again, for this test, both Virtual Connect modules were present and managing the physical addresses.

Figure 5. Six hosts powered on host 1 after host 2 simulated failure.



As an extended test of functionality, it was expected that VMware DRS would automatically redistribute the virtual machines between the two hosts when the failed host was brought back online without interruption of service at the virtual machine level. This behavior was observed and confirmed as shown in Figure 6.

Figure 6. Recent tasks showing successful VMotions carried out between hosts 1 and 2.



All testing for core functionality of VMware ESX Server 3.0 and VMware Infrastructure 3 passed with servers configured in a Virtual Connect domain without issue or specific configuration within either ESX Server or the virtual machines.

Summary

HP Virtual Connect and VMware Infrastructure 3 combine to bring the ultimate in flexibility to enterprise computing. When combined, customers achieve new levels of change readiness and simplification of processes.

For more information

HP and VMware on HP.com, <http://www.hp.com/go/vmware>

HP BladeSystem, <http://www.hp.com/go/blades>

HP Virtual Connect Primer, <http://h71028.www7.hp.com/ERC/downloads/4AA0-5821ENW.pdf>

HP Virtual Connect for BladeSystem c-Class User Guide,
<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00865618/c00865618.pdf>

To help us improve our documents, please provide feedback at www.hp.com/solutions/feedback

© 2007 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA1-1145ENW, Rev. 1, March 2007

