


## **Practical 1: implement VMWare ESXi for server virtualization:**

- VMware ESX is an enterprise-level computer virtualization product offered by VMware, Inc.
- ESX is a component of VMware's larger offering, VMware Infrastructure, which adds management and reliability services to the core server product. VMware is replacing the original ESX with ESXi.
- VMware ESX and VMware ESXi are Type 1 hypervisors that are VMware's enterprise software hypervisors for guest virtual servers that run directly on host server hardware without requiring an additional underlying operating system.
- The basic server requires some form of persistent storage (typically an array of hard disk drives) that store the hypervisor and support files.
- A smaller footprint variant, ESXi, does away with the first requirement by permitting placement of the hypervisor on a dedicated compact storage device. Both variants support the services offered by VMware Infrastructure.

### Architecture:-

- VMware states that the ESX product runs on bare metal. In contrast to other VMware products, it does not run atop a third-party operating system, but instead includes its own kernel.
- A Linux kernel was started first, and is used to load a variety of specialized virtualization components, including VMware's vmkernel component.
- This previously booted Linux kernel then becomes the first running virtual machine and is called the service console.
- Thus, at normal run-time, the vmkernel is running on the bare computer and the Linux-based service console runs as the first virtual machine.
- As of version 4.1, VMware has dropped development of ESX and now focuses exclusively on ESXi, which is devoid of a Linux kernel.
- The vmkernel itself, which VMware says is a microkernel, has three interfaces to the outside world:
  - hardware
  - guest systems
  - service console (Console OS)

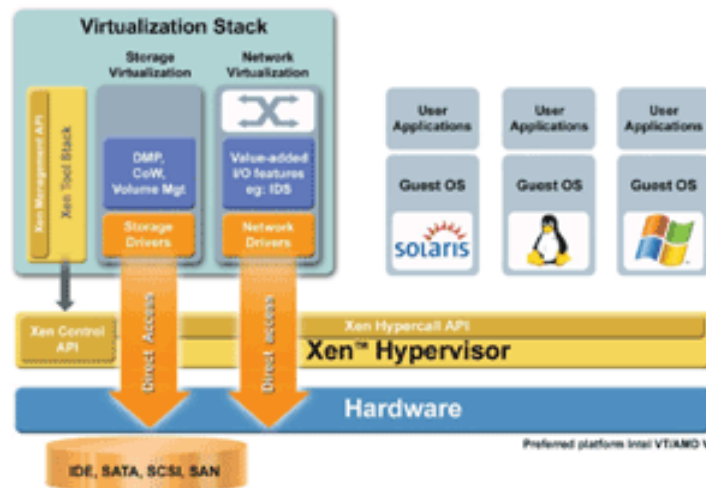


## Practical 2: Implement ZEN for server virtualization:

- Xen is a hypervisor using a microkernel design, providing services that allow multiple computer operating systems to execute on the same computer hardware concurrently.
- The University of Cambridge Computer Laboratory developed the first versions of Xen.
- The Xen community develops and maintains Xen as free and open-source software, subject to the requirements of the GNU General Public License (GPL), version 2.
- Xen is currently available for the IA-32, x86-64 and ARM instruction sets. Xen runs in a more privileged CPU state than any other software on the machine.
- Responsibilities of the hypervisor include memory management and CPU scheduling of all virtual machines ("domains"), and for launching the most privileged domain ("dom0") - the only virtual machine which by default has direct access to hardware.
- From the dom0 the hypervisor can be managed and unprivileged domains ("domU") can be launched.
- The dom0 domain is typically a version of Linux, or BSD. User domains may either be traditional operating systems, such as Microsoft Windows under which privileged instructions are provided by hardware virtualization instructions (if the host processor supports x86 virtualization, e.g., Intel VT-x and AMD-V), or para-virtualized operating system whereby the operating system is aware that it is running inside a virtual machine, and so makes

hypercalls directly, rather than issuing privileged instructions.

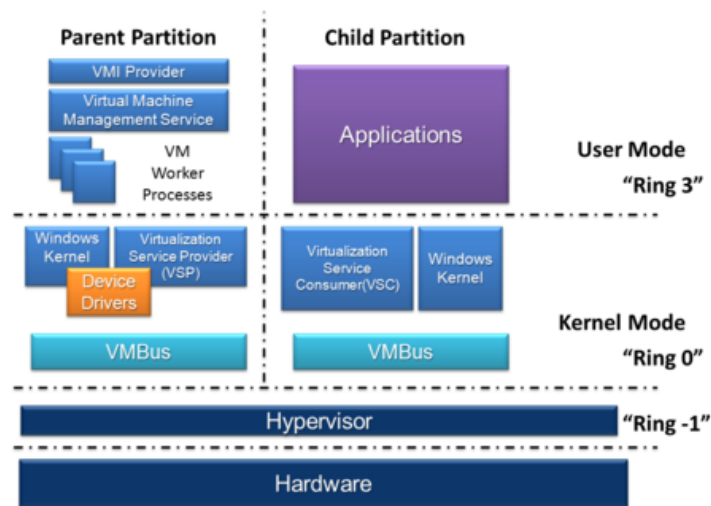
- Xen boots from a bootloader such as GNU GRUB, and then usually loads a paravirtualized host operating system into the host domain (dom0).



### Practical 3: Implement Hyper-v server virtualization:

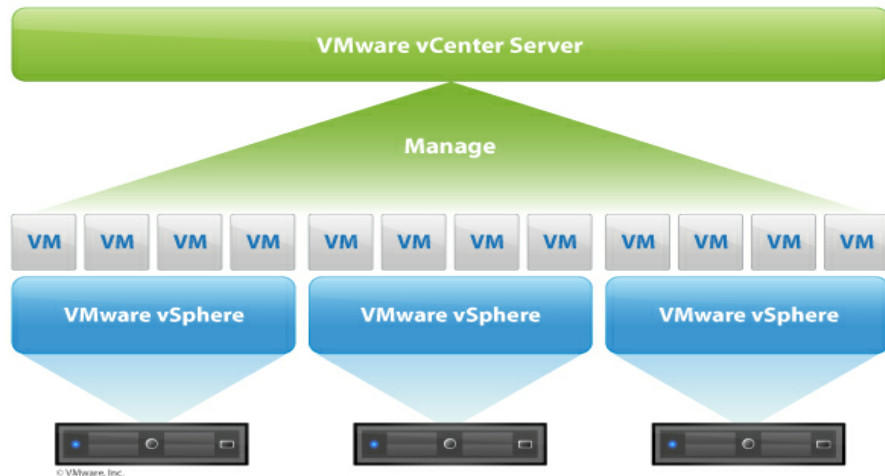
- Hyper-V, codenamed Viridian and formerly known as Windows Server Virtualization, is a native hypervisor that enables platform virtualization on x86-64 systems.
- A beta version of Hyper-V was shipped with certain x86-64 editions of Windows Server 2008, and a finalized version (automatically updated through Windows Update) was released on June 26, 2008.
- Hyper-V has since been released in a free stand-alone version, and has been upgraded to Release 2 (R2) status. It was updated in Windows Server 2012.
- A host server with Hyper-V capability could be accessed remotely by multiple guest computers.
- Each guest computer could perform as if they are using the host server directly.
- Users on the guest computers could run applications in the host server remotely, from his computer, even though that application is not available on the guest computer.
- Hyper-V exists in two variants:
  1. as a stand-alone product called Hyper-V Server: Four major versions have so far been released: Hyper-V Server 2012 R2 (containing the current release of Hyper-V), Hyper-V Server 2012, Hyper-V Server 2008 R2 and Hyper-V Server 2008.

2. As an installable role in Windows Server 2012 R2, Windows Server 2012, Windows Server 2008 R2, Windows Server 2008 and the x64 edition of Windows 8 Pro.



#### Practical 4: Manage vmwareESXi with vCentre server.

- VMware vCenter Server provides centralized visibility, proactive management and extensibility for VMware vSphere from a single console.
- Centralized Control and Visibility
  - **vSphere web client:** Manage the essential functions of vSphere from any browser anywhere in the world.
  - **vCenter single sign-on:** Allow users to log in once and access all instances of vCenter Server, without the need for further authentication.
  - **Custom roles and permissions:** Restrict access to the entire inventory of virtual machines, resource pools and servers by assigning users to custom roles. Users with appropriate privileges can create these custom roles, such as night-shift operator or backup administrator.
  - **Inventory search:** Explore the entire vCenter inventory—including virtual machines, hosts, datastores and networks—from anywhere within vCenter.



### Practical 5: managing ZEN server with zencenter.

- Virtual Machine Manager (VMM) enables you to deploy and manage virtual machines and services across multiple hypervisors, including Citrix XenServer hosts.
- Through VMM, you can manage the day-to-day operations of XenServer hosts and XenServer pools.
- These operations include the discovery and management of XenServer hosts and pools, and the ability to create, manage, store, place and deploy virtual machines and services on XenServer hosts.
- Managing XenServer hosts through VMM also gives you more choice with regard to Linux-based guest operating systems than if you were only managing Hyper-V.

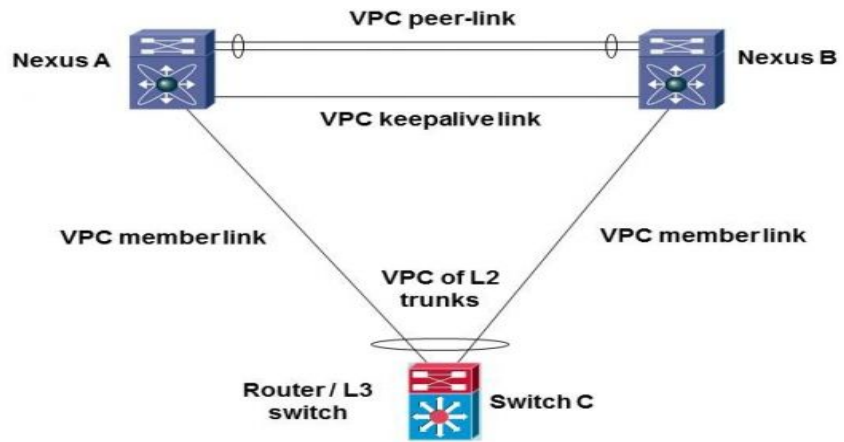
### Practical 6: Understanding Blade server with Cisco UCS/HP simulator:

- The Cisco UCS 5100 Series Blade Server Chassis is a crucial building block of the Cisco Unified Computing System, delivering a scalable and flexible architecture for current and future data center needs, while helping reduce total cost of ownership.
- Cisco's first blade-server chassis offering, the Cisco UCS 5108 Blade Server Chassis, is six rack units (6RU) high, can mount in an industry-standard 19-inch rack, and uses standard front-to-back cooling.
- A chassis can accommodate up to eight half-width, or four full-width Cisco UCS B-Series Blade Servers form factors within the same chassis.
- The Cisco UCS 5108 Blade Server Chassis revolutionizes the use and deployment of blade-based systems.
- By incorporating unified fabric and fabric-extender technology, the Cisco Unified Computing System enables the chassis to:
  - o Have fewer physical components

- o Require no independent management
- o Be more energy efficient than traditional blade-server chassis
- This simplicity eliminates the need for dedicated chassis management and blade switches, reduces cabling, and allowing scalability to 20 chassis without adding complexity.
- The Cisco UCS 5108 Blade Server Chassis is a critical component in delivering the simplicity and IT responsiveness for the data center as part of the Cisco Unified Computing System.

### **Practical 7: implement VLAN with L2/L3 switches or nexus virtual switching**

- Virtual Port Channel (VPC) ties two identical Nexus 9K, 7K, 6K, 5K, or 3K switches together by doing a great job of spoofing Layer 2 (L2), including STP BPDU and FHRp behavior.
- At L2, the switches behave like one switch, while retaining their distinct identities, unlike the Cisco VSS technique used with the 6500 switches.
- For VPC, you would likely connect the two switches together by 10 G links, preferably one on each of two different line cards for robustness, and declaring that link to be a port-channel and a VPC peer-link. It is a special link for carrying L2 traffic between the VPC peer switches when there is a link failure.
- Once the VPC pair is set up (configured), you can then start adding multi-chassis port-channels that connect to the VPC pair.
- When you connect up another device, and tell it two or more links are in a port-channel, and they're spread across the two VPC peers, that's a VPC port-channel. The links that make it up are referred to as member links.
- When setting up member links, you do have to configure the VPC peer end of things as well. Both ends of the member links have to be configured for port-channel. The VPC switch ends get a special "vpc number" command as well.



**Practical 8: Simulating SAN with navisphere/ netapps:**