

# Xen virtualization with Novell SUSE Linux Enterprise Server 10 on HP ProLiant servers



Executive summary.....	2
Xen virtualization overview .....	2
Business case .....	3
Novell solution.....	3
Why HP and Novell.....	3
Software solution .....	4
Using Xen – A technical use case .....	4
Putting Xen to use.....	5
Problem solving with Xen .....	6
Notes about using Xen.....	6
Upgrading a VM .....	7
For more information.....	9

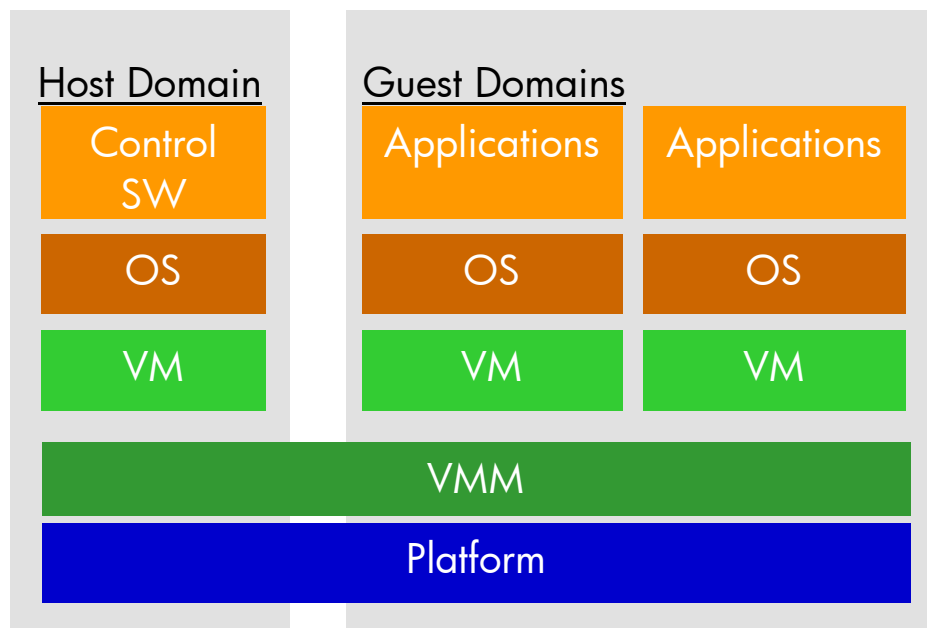
## Executive summary

Virtualization has become the buzz word of the industry, with people looking at consolidation of multiple servers onto a single server. Many reasons for consolidation come to mind, including cost reduction, better utilization of hardware, and better control of resources. This document will look at virtualization using Xen, Novell's implementation of Xen within SUSE Linux Enterprise Server 10, and implementation on HP ProLiant servers.

## Xen virtualization overview

Xen is an open-sourced project to provide a server environment to host virtual machines. The hypervisor or virtual machine monitor (VMM) is the software layer that is initially loaded to provide the virtual machine (VM) server functionality. The VMM runs between the server hardware and Linux operating system and is loaded first at boot. Once the VMM has loaded, the Xen VM Server is loaded to create and control the other VMs and communicate with the server hardware. This Xen VM Server is referred to as Dom0 or domain0 and runs in privileged mode. A VM, also referred to as a guest domain (DomU), is an isolated environment running an operating system and applications. This guest domain runs unprivileged. A guest domain may or may not know it is running in a VM, depending on whether it is a para-virtualized or fully-virtualized VM.

Figure 1. Xen Overview



A para-virtualized VM means that the VMM has APIs (application program interfaces) to assist in accessing the hardware, and the guest operating system has been modified to know it is running in a VM. The VMM emulates the underlying hardware by presenting virtual devices to the guest operating system. A fully-virtualized VM requires no modifications to the guest operating system, but the VMM must trap and emulate all privileged instructions. For fully-virtualized VMs the physical server must have processors that support Virtualization Technology (VT). HP has VT enabled hardware utilizing both Intel® and AMD™ processors. Fully-virtualized VMs perform slower than para-virtualized VMs.

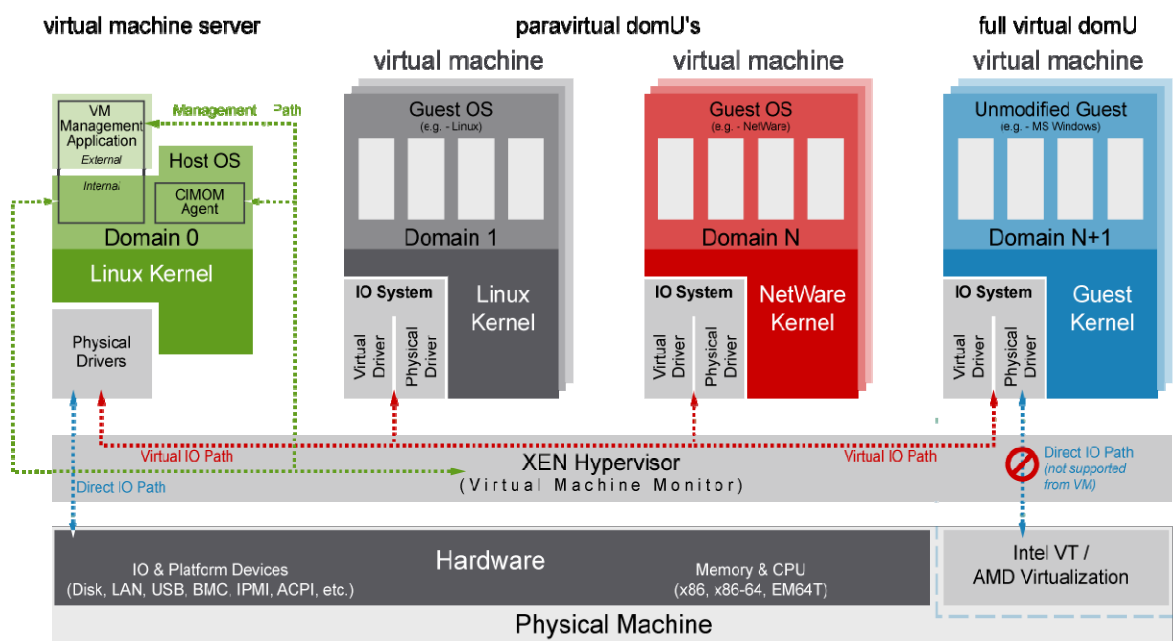
## Business case

As IT centers look to consolidate multiple servers into a virtualized environment to save costs, the cost of virtualization software also becomes a concern. Although in its infancy compared to other virtualization software products, Xen has made great strides in a short period of time, and with the latest Open Source releases of Xen it is now ready to be utilized in the data center.

## Novell solution

Novell is committed to the propagation of Xen virtualization technologies. With the release of SUSE Linux Enterprise Server 10 Service Pack 1 (SP1), Novell has included all of Xen 3.0.4 with some features of Xen 3.1.0 (formerly Xen 3.0.5) as an integrated part of the release. SUSE Linux Enterprise Server 10 simplifies virtualization technology management with YaST, a comprehensive installation, configuration and administration suite that is integrated into the SUSE Linux Enterprise platform. YaST is ideal for creating virtual machine profiles and configuring individual VMs because it takes you through the same series of screens that you use to provision a physical machine. Being integrated into the SUSE Linux Enterprise Server 10 operating systems also means any updates to the Xen software can be obtained from Novell and their Novell Customer Center for operating system updates.

Figure 2. Novell XEN Diagram



Novell has also done extensive work on the ability to mix and match different Dom0 and hypervisor combinations. This work allows 32-bit Dom0 (host domain) to run on a 64-bit hypervisor, allowing more flexibility. For a list of what is supported, visit the Novell website, <http://www.novell.com/documentation>.

## Why HP and Novell

HP and Novell enjoy one of the longest partnerships in the industry, dating from Novell's earliest NetWare offerings to today's cutting-edge solution products and SUSE Linux Enterprise operating systems. HP systems and Novell software are literally designed for one another. Close engineering

and development collaboration between the two companies means that Novell's software is optimized and certified to run smoothly on HP systems. But the strength of the partnership does not end there. HP and Novell join forces to ensure that their joint customers receive world-class customer support and issue resolution through a single point of contact.

## Software solution

With the release of SUSE Linux Enterprise Server 10 Service Pack 1 Novell has advanced the capabilities of Xen and establishes the technology as a viable and responsive alternative for virtualization. Out of the box, SUSE Linux Enterprise Server 10 SP1 Xen can create virtual machines running modified, highly tuned, para-virtualized guest operating systems for optimal performance and leverages both Intel VT and AMD Virtualization (AMD-V™) processor capabilities to fully virtualize Microsoft® Windows® and legacy operating systems.

With Novell SUSE Linux Enterprise 10 SP1 Novell now supports:

- Para-virtualized as modified virtual machine operating system on top of Xen-enabled SUSE Linux Enterprise 10:
  - SUSE Linux Enterprise Server 10 SP1
  - NetWare 6.5 SP7 (as part of OpenEnterprise Server (OES) 2)
- Fully virtualized as unmodified virtual operating system using Intel VT or AMD-V chipsets
  - SUSE Linux Enterprise Server 9 SP3
  - Red Hat Enterprise Linux 4 and 5
  - Microsoft Windows XP, 2003, and 2000

HP has been working closely with Novell in testing and certifying SUSE Linux Enterprise Server 10 SP1 on HP ProLiant servers. A list of all certified and supported servers for SUSE Linux Enterprise Server 10 and Xen can be found at:

<http://h18004.www1.hp.com/products/servers/linux/hpLinuxcert-novell.html>.

## Using Xen – A technical use case

Xen is an emerging technology in open source that brings a new dimension to the field of virtualization technology. As with all open source projects, Xen will benefit by contributions from the Linux and Open-Source communities and development will progress rapidly. As a result end-customers will benefit when companies like Novell integrate Xen into their distributions thereby fostering faster acceptance and growth of Xen virtualization.

Virtualization in general continues to grow in acceptance as people see the cost saving and simplification benefits of consolidating several servers into one while maintaining server individuality and control. With the current performance and load capabilities of the HP ProLiant servers, certain tasks like DHCP/DNS servers, LDAP servers, or NFS/TFTP boot servers rarely generate enough load on a current ProLiant server to warrant dedicating a single system. However running multiple services or tasks on a single server can be complex. For example, it is not desirable to address an NFS server issue by shutting down a DHCP or an LDAP server. Nor should modifications be made to a DHCP server that affect NFS users. Administrators also have the need to balance server usage over time. For example DHCP, TFTP, and NFS server functionality may be required for a short period of time while the balance of the time the server only needs to be an NFS server. Each of these examples suggests the need for better server utilization through virtualization.

Virtualization is also beneficial for legacy applications and operating systems. Older versions of Linux with specific software applications that are currently deployed on older hardware become a problem

when a new hardware purchase occurs. The new hardware will likely not support the old version of Linux however the application is required and cannot be ported to a new version of Linux at the time. Virtualizing the old environment on the new hardware platform is a possible solution to this scenario.

## Putting Xen to use

In looking at a real world scenario, our lab receives new HP ProLiant servers for testing and development, often with new or un-released versions of operating systems. Most commercial deployment applications cannot support the new software. To address this problem we built a server using tftpbboot, PXE, DHCP, DNS, and the HP SmartStart Scripting Toolkit to do deployments. This was a working solution until there were issues such as needing the deployment server hardware for specific testing, or needing a deployment server in a different location without having to rebuild everything. The solution was to build a Xen VM running SUSE Linux Enterprise Server 10 paravirtualized that can be a DHCP and DNS server, handle tftpbboot and PXE requests, and run the SmartStart Scripting Toolkit for deployments, plus run NFS and HTTP to allow other servers access to the distribution bits. This deployment VM is a 4 GB sparse file, utilizing 1 GB of memory, 1 virtual CPU and 2 NICs by default. The first NIC allows access to the external network and Internet, the second NIC has DHCP running and is the NIC that responds to PXE and NFS/HTTP requests for installations. ISO images are supplied to the VM by either using iSCSI connections to a server holding the ISO images, or by attaching additional sparse files containing the necessary ISO images to the VM. To create a new deployment server we copy the VM to a server and setup the necessary access to the ISO images. If the server the VM is running on needs to be upgraded we migrate the VM to another server, do the necessary upgrades, and migrate it back with no loss to the DHCP/DNS server or the installation bits.

The deployment VM was best utilized in building the “Enterprise in a Rack”, a demo rack we ship to different trade shows like Brainshare, LinuxWorld, and HP Technology forum. The rack contains ProLiant c-Class blades and DL servers running multiple applications. Because things change it often becomes necessary to rebuild servers and upgrade operating systems, including the deployment server. By utilizing the deployment VM and moving it from server to server as needed we were able to leverage existing servers within the rack to support our deployment VM without lost time of having to rebuild the deployment server due to hardware changes.

The HP BladeSystem c-Class enclosure within the rack contained four c-Class blades running SUSE Linux Enterprise Server 10 and Xen, but in actuality there were fifteen servers running. The first c-Class blade, an HP ProLiant BL480c with 8 GB of memory was running the deployment VM, as well as being the repository for VM images. The deployment VM was a 4 GB sparse file containing the OS, and four 8 GB sparse files containing the installation ISO images. A second c-Class blade, a BL460c with 10 GB of memory and an HP StorageWorks SB40c storage blade, was the Windows host, running two instances of Windows Server 2003 and four instances of Windows XP fully virtualized. Access to the Windows VMs was from an HP 5135 Thin Client. A third c-Class blade, a BL460c with 12 GB of memory and 2 dual-core CPUs was utilized to run six virtual machines, each configured with 1 GB of memory, 1 CPU and 1 NIC. Three of the VMs were running a mySQL stress test, and the other three were running a JBoss demo, all were running SUSE Linux Enterprise Server 10 paravirtualized. The power of Xen was highlighted when, through an administrator error, the storage array containing the VMs was wiped clean, everything deleted. Since all six VMs had been copied from a base VM, it took less than 10 minutes to have all six VMs back up and running. Setting up mySQL and JBoss was all that was necessary to get everything back in working order. The fourth c-Class blade was a BL465c running a Windows XP VM on a management network to allow certain tools requiring a Windows environment and a second VM running SUSE Linux Enterprise Server 10 to allow access to a Linux desktop.

## Problem solving with Xen

Another example of our using Xen was to solve the problem of setting up a training session that was to be repeated multiple times over several days. During the lab, each student needed to configure a server for a specific task. Once the class was over, we needed to quickly reset all the servers to their original state for the next class. Deleting the changes made by each student or doing a complete re-install on each server were both too time consuming and too error-prone, especially on 12 different servers. The solution was to build a VM to an initial state and save the image. During the class, each student operated on their own copy of the VM image, and then once the class was done the initial state image was quickly restored to each server. This also greatly reduced the amount of hardware required for the class. A single BL480c was configured with fourteen NICs, and ran all of the student VMs. Each VM was assigned its own network and NIC to isolate network conflicts.

Although no serious stress testing was done on any of the VMs in either example, the demo system has been up and running for weeks at a time with no failures or issues seen from either the paravirtualized SUSE Linux Enterprise Server 10 VMs or the Windows VMs. The next section discusses a couple of configuration issues to consider when deploying VMs.

## Notes about using Xen

When copying image files the expectation is that the defined NICs in the image file will use the information supplied in the configuration file and have the correct IP information. In SUSE Linux Enterprise Server, NIC configuration files use the MAC address as part of the filename, then associate that configuration file with eth0, eth1, or whatever was defined for that NIC. For example, a NIC with a MAC address of 00:00:00:11:11:11 would create a configuration file in `/etc/sysconfig/networks` called `ifcfg-eth-id-00:00:00:11:11:11` and this may be associated with eth0. If you copy the image file to a new file, and create a Xen configuration file containing the MAC address of 00:00:00:22:22:22, when you start the new VM it will have two Ethernet entries in the `/etc/sysconfig/networks` directory: `ifcfg-eth-id-00:00:00:11:11:11` tied to eth0 and `ifcfg-eth-id-00:00:00:22:22:22` tied to eth1, which has not been configured yet. Because SUSE Linux Enterprise Server 10 forces persistent names by default, each new NIC will be assigned a new eth association. There are two approaches to solving this problem. The first requires editing the `/etc/sysconfig/networks/config` file, locate the variable `FORCE_PERSISTENT_NAMES` and change the value to no. Then you need to change the names of the `ifcfg` files to represent the necessary Ethernet association, e.g., change `ifcfg-eth-id-00:00:00:11:11:11` to `ifcfg-eth0`. Once this has been done, copying and starting the image file will use the MAC address from the Xen configuration file for eth0 and obtain the expected IP information. The second approach requires modifying the `/etc/udev/rules.d/30-net_persistent_names.rules` and assigning specific MAC addresses to specific NICs.

Secondly, by default SUSE Linux Enterprise Server 10 Xen creates one network bridge at startup – `xenbr0`. This bridge will be tied to a physical NIC in `domain0`. If you have more than one physical NIC in the server and wish to have a second NIC defined in the VM, you must create the additional bridges manually. SUSE Linux Enterprise Server 10 has scripts in `/etc/xen/scripts` to start/stop bridges, setup NAT, and other tasks. For this we use the `/etc/xen/scripts/network-bridge` script to create and start a second bridge. The basics required to start a new bridge are: start option, vif number, bridge name, netdevice (eth1 for example) and hardware device configuration file. For example:

```
/etc/xen/scripts/network-bridge start vifnum=1 bridge=xenbr1 netdev=eth1 hwddev=ifcfg-eth1
```

would start a new bridge named `xenbr1` on eth1 using the configuration information from the `ifcfg-eth1` file. The line:

```
vif = [ 'bridge=xenbr0, mac=00:00:00:11:11:11', 'bridge-xenbr1, mac=00:00:00:22:22:22' ]
```

would pass two NICs to the VM with different MAC addresses, eth0 pointing to xenbr0 and eth1 pointing to xenbr1.

To start multiple bridges look at: <http://wiki.xensource.com/xenwiki/XenNetworking>. This shows how to modify the networking scripts in Xen to take advantage of a custom start script. Another approach to starting multiple bridges can be found in the `/etc/xen/xend-config.sxd` file.

## Upgrading a VM

As stated earlier, this is a lab, and so multiple installations and upgrades happen to both the software and the hardware. At some point, while upgrading from beta versions of SP1 to newer versions of SP1 some of the VMs running SUSE Linux Enterprise Server 10 did not get updated. This created a scenario where the VMs would not boot, completely understandable when dealing with beta code. However, this was an easy fix for the VMs. To solve this problem, and upgrade the VM, we added SP1 DVD ISO as a CD-ROM to the VM configuration file:

```
disk=['file:/var/lib/xen/images/srv1/xvda,xvda,w', 'file:/root/SLES-10-SP1-DVD1-i386.iso, xvdb, r',]
```

The next issue was to have a kernel and ramdisk from SP1 that could be used to boot the VM. This was created by starting to do an installation of a SUSE Linux Enterprise Server 10 SP1 VM through the `virt-manager` tool, and make copies of the kernel and ramdisk in `/tmp`. The VM configuration file for the VM to be updated was modified to point to the kernel and ramdisk:

```
kernel = "/tmp/kernel.save"  
ramdisk = '/tmp/initrd.save'
```

It is also necessary to specify the installation device:

```
extra=" TERM=xterm install=hd:/dev/xvdb"
```

and comment out the bootloader and bootargs lines in the current configuration file. Once the VM is started using an `xm create <config file>` command, the VM will boot from the SP1 CD allowing you to do an upgrade. It is important that you select Upgrade and not New Installation. When the upgrade finishes loading software it will try to boot but fail because it is not running the correct kernel. Terminate the VM using the `xm destroy <VMid>` command and edit the configuration file for the VM. Delete the kernel and ramdisk lines, and remove the comments from the bootloader and bootargs lines, then restart the VM. The VM will now finish the upgrade. During this process the graphics may not function correctly and you may have to finish the upgrade in TEXT mode. Once the upgrade has finished, if you run your VM at init level 5 you will need to run `sax2` to reset the graphics configuration. It should be noted however that VMs are rarely run at init level 5 because the X windows graphics are memory and CPU hogs. Most graphics access to the VMs is done through `vnc` (virtual network computing), which does not require the X windows to be running. As a performance recommendation, run the Linux VMs at run level 3, and use `vnc` to access the graphics on the VM. It should also be noted that the upgrade problem was self-induced due to a lab environment, and in the real world proper steps should be taken to prevent this from happening.

## Summary

Stress testing of Xen in SUSE Linux Enterprise Server 10 Service Pack 1 is ongoing. Novell is releasing para-virtualized drivers for SUSE Linux Enterprise Server 9 and Windows to improve performance of both. Novell is also releasing support for Open Enterprise Server NetWare running para-virtualized on SUSE Linux Enterprise Server 10 Xen. In the future there will be support for fully-virtualized NetWare, this will allow migration of NetWare servers to virtualized environments with Xen.

With the release of SUSE Linux Enterprise Server 10 SP1 Novell has brought more power and support to Xen, and made it ready to utilize in the data center. Novell ZENworks Linux Management (ZLM) complements YaST by enabling administrators to centrally control how to deploy and update systems—or VMs—inside the firewall. By using YaST and ZENworks Linux Management, data center administrators can easily install, configure, update, secure and manage SUSE Linux Enterprise Server 10 and the virtual machines it hosts. They can also start up, shut down or suspend virtual machine operating system instances and migrate VMs from one physical server to another in order to relocate a load or prepare for maintenance.

Customers considering using virtualization within the data center will find Novell SUSE Linux Enterprise Server 10 and Xen a responsive choice. HP fully supports Novell SUSE Linux Enterprise Server 10 on HP ProLiant servers, and with the release of Service Pack 1 for SUSE Linux Enterprise Server 10 embraces Xen on HP ProLiant servers

## For more information

HP ProLiant Essentials software, [www.hp.com/go/proliantessentials](http://www.hp.com/go/proliantessentials)

HP ProLiant servers, [www.hp.com/go/proliant](http://www.hp.com/go/proliant)

HP Systems Insight Manager (HP SIM), <http://www.hp.com/go/hpsim>

HP SmartStart Scripting Toolkit,

[http://h18013.www1.hp.com/products/servers/management/toolkit/downloadtkpage.html?jumpid=reg\\_R1002\\_USEN](http://h18013.www1.hp.com/products/servers/management/toolkit/downloadtkpage.html?jumpid=reg_R1002_USEN)

Novell SUSE Linux Enterprise Server 10 documentation,

<http://www.novell.com/documentation/sles10/>

To help us improve our documents, please provide feedback at [www.hp.com/solutions/feedback](http://www.hp.com/solutions/feedback).

© 2006, 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.

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. AMD and AMD-V are trademarks of Advanced Micro Devices, Inc. Microsoft and Windows are U.S. registered trademarks of Microsoft Corporation.

4AA0-9251ENW, Revision 3, June 2007

