

# Xen Expert Days

## Virtualization with Xen

### SUSE Linux Enterprise Server 10

**Ralf Dannert**  
Technology Specialist  
[rdannert@novell.com](mailto:rdannert@novell.com)

**Novell®**


# Agenda

- Use cases
- Terminology and Architecture
- VM installation
- Using Xen
- Case Studies
- Roadmap
- Novell offerings
- Helpful Links



# Server Virtualization: Analyst's View

Server Virtualization Shipments, 2004–2009

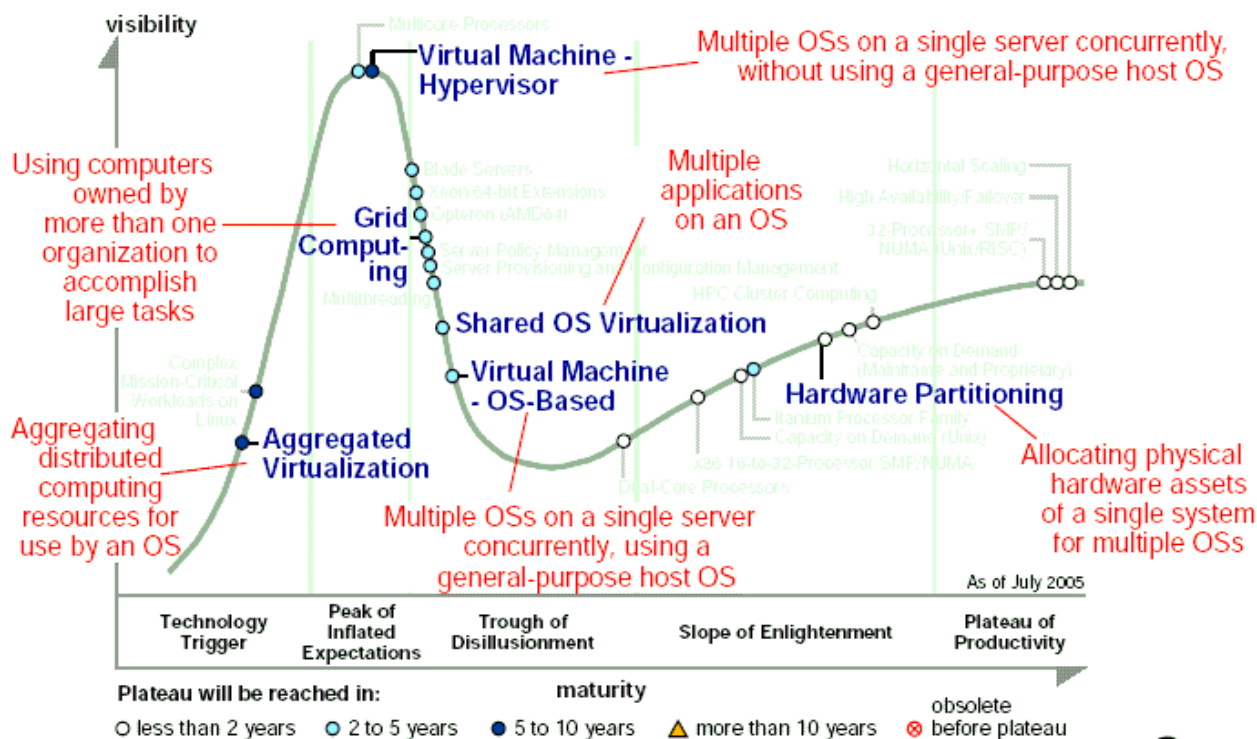


Source: IDC, 2006

Strategic Planning Assumption: Aggregated virtualization technologies will reach the peak of the Server Hype Cycle by 2008 (0.7 probability).

Virtual machine hypervisors will mature quickly, and adoption will grow so that they will reach the Plateau of Productivity by 2009 (0.8 probability).

## Server Virtualization on the Hype Cycle for Server Technologies




# Executive Summary


- **Novell product portfolio offers choice**
  - Customers can choose to deploy virtualization technologies provided by Novell and other VT vendors
- **Novell virtualization strategy is focused on Xen**
  - Customer demand for server consolidation and price / performance will foster rapid acceptance
- **Novell supports customers**
  - Virtual Machine Server Hardware from partners and Virtual Machine configurations are listed in YES certification bulletins
  - Novell Technical Services supports installation and operation

# Use Cases


# Uses of Virtualization




Examples: VMs, LPARs, Virtual Disks, VLANs



Examples: Virtual Disks, Virtual Storage Pools




Examples: iSCSI, Architecture Emulators



Examples: Spare CPU Substitution, CUoD

# Use cases

- Virtualization allows for more flexibility
  - Virtual Machines (VM) isolate hardware differences due to a abstracted resource layer between hardware and OS
  - Decouples software stacks from hardware life cycles
  - Dynamic provisioning reduces time to operation: pre-configured application stacks are faster to deploy
  - Integrated high availability increases reliability

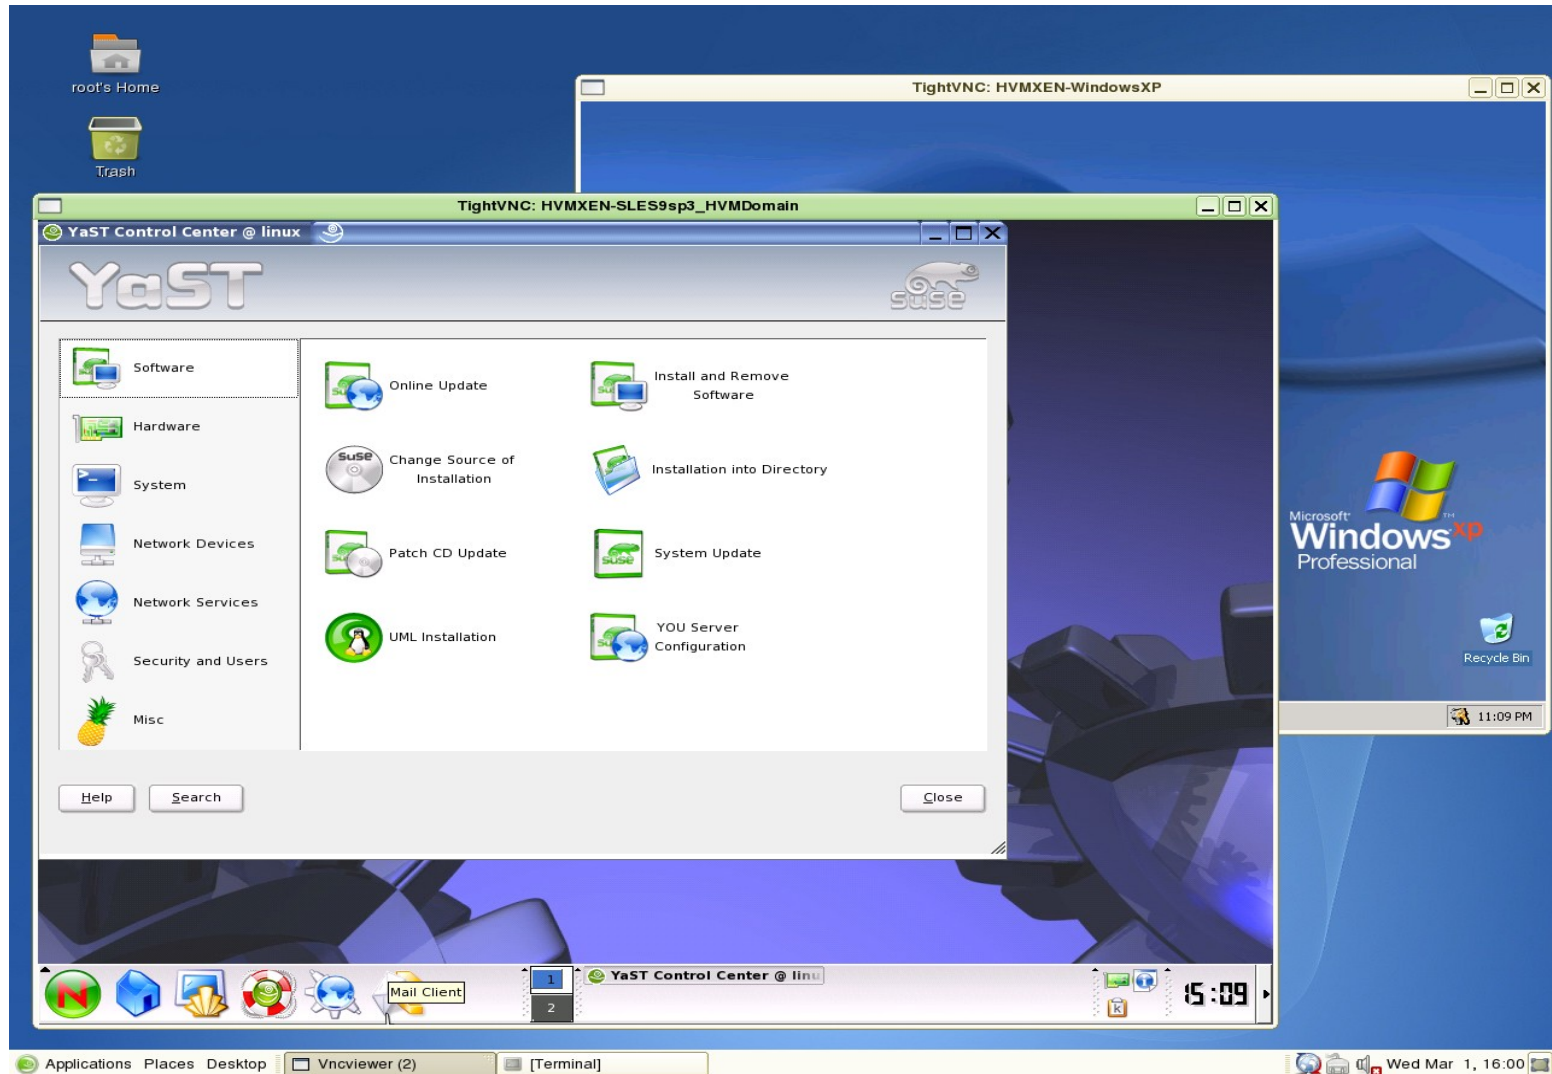




# Terminology and Architecture

# SLES 10 VM Server

## SLES 9 and Windows XP - Fully Virtualized VMs



# Xen Technology Background

Originally a research project from University of Cambridge

Open source

Xen 2.0 released November 2004

Xen 3.0.0 released December 2005

Xen 3.0.2 release May 2006 (SLES 10 Target)

Xen 3.0.4 SLES 10 SP1 Target

# Xen Community Terminology

**Domain:** A container for a running virtual machine. Colloquially, the VM itself.

**Domain 0:** The first domain. Privileged to manage other domains. a.k.a. “dom0”.

**Unprivileged domain:** Any domain other than domain 0. Cannot manage other domains. a.k.a. “domU”.

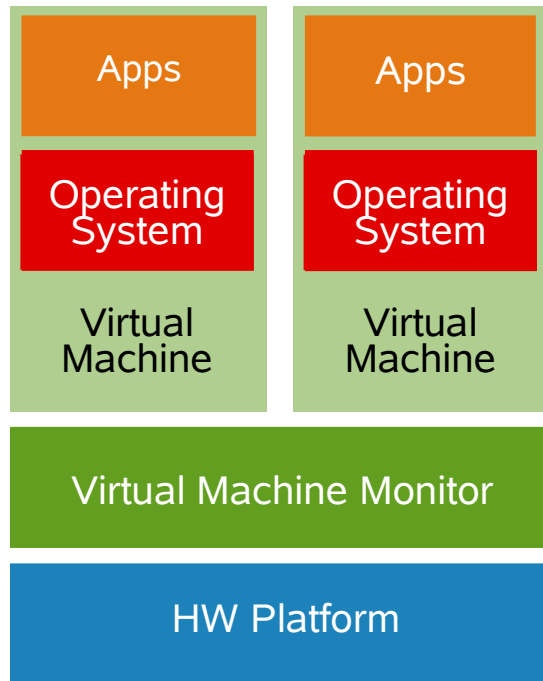
**Driver domain:** A domain that contains physical drivers. Usually this is just domain 0.

**Physical driver:** A device driver (usually in the driver domain) that talks to the hardware.

**Virtual driver:** A device driver (usually in a domU) that fulfills requests by going to the physical driver.


# Full & Paravirtualization Overview

## Full Virtualization



Runtime modification of Guest OS:  
VMM manages the conflict, then  
returns to OS

## Paravirtualization



Static modification of Guest OS prior to  
runtime: Privileged instruction calls are  
exchanged with API functions provided  
by the VMM

- Almost no performance degradation
- Significant scalability

# Novell Terminology

**Fully Virtual:** A VM mode that can run a native, unmodified operating system by emulating all hardware devices.

**Paravirtual:** A VM mode that can run a modified operating system, which cooperates with the VMM.

**VT Computer:** Computer supporting HVM Intel VT, AMD

**Standard Computer:** A computer that does not support virtualization technology and therefore can run Xen VMs only in paravirtual mode.

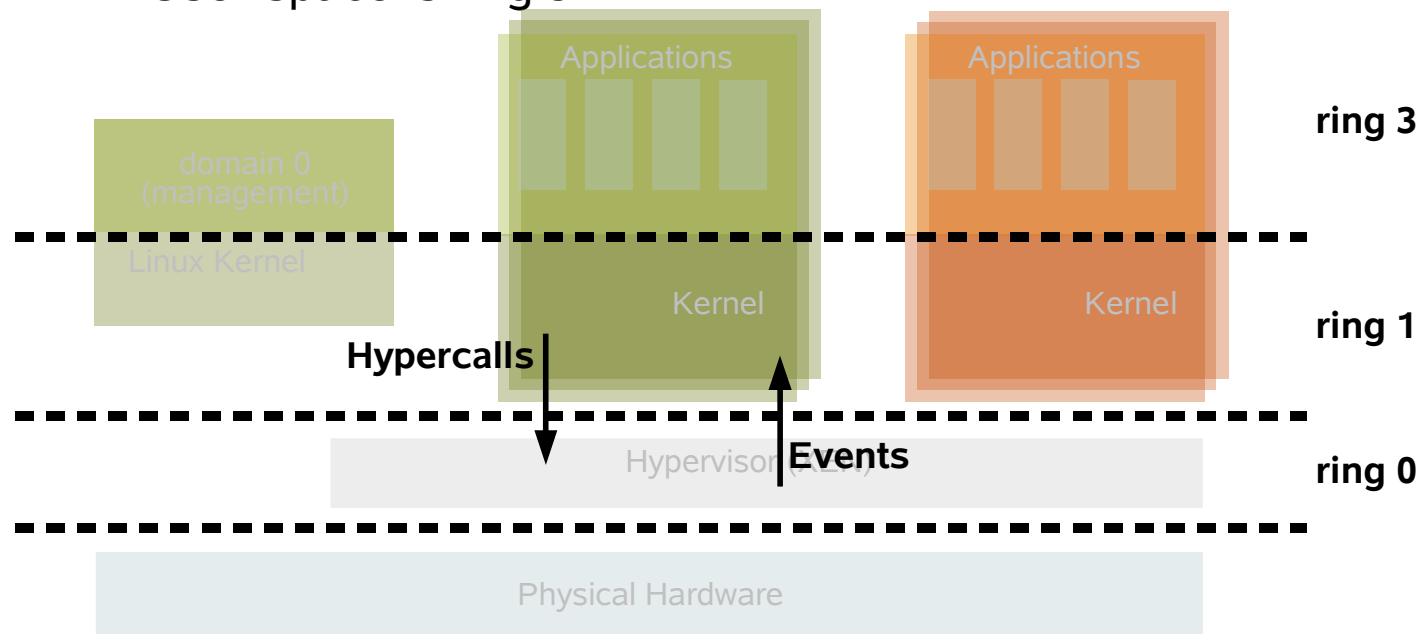
**Native Operating System:** A typical operating system that is not optimized for the VM environment and must run in fully virtual mode.

# Privilege Rings

Xen runs at ring 0 (highest privilege)

All domains run at rings 1 - 3.

- Kernel is ring 1
- User-space is ring 3



# Hardware assisted virtualization

- **VT Computer**

- run multiple OS concurrently
- protected execution environments
- privilege ring expansion
- simplify hypervisor
- Intel VT for directed I/O (VT-d) - direct assign I/O
  - no emulated drivers necessary



# Extending Intel Virtualization Technology

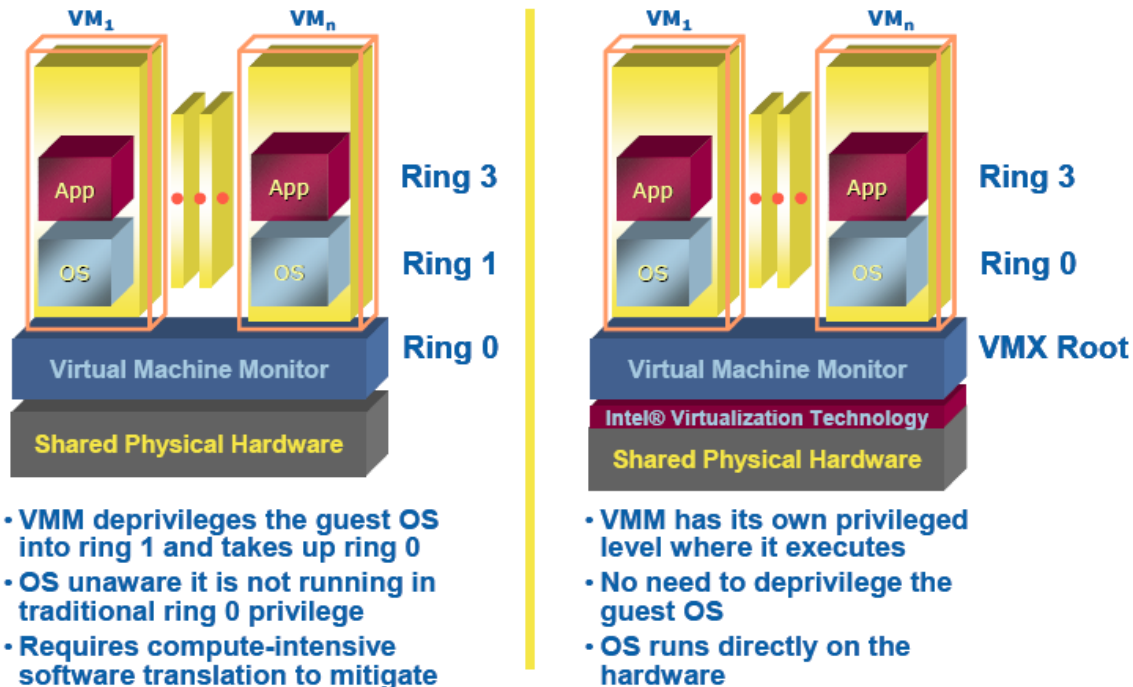
- support for I/O device virtualization
  - direct I/O virtualization to the chipset(“VT-d”).  
Currently, I/O devices aren’t aware of virtualization and must go through the VMM before being assigned to a virtual machine.
- software emulation slow
  - **Performance**
    - > I/O requests must traverse two I/O stacks (guest and host)
  - **Functionality**
    - > Guest OSes “see” only restricted sets of legacy devices
  - **Reliability**
    - > Drivers are potentially undependable if they run as part of privileged software

# Full Virtualization Mode on VT using qemu-dm

- using “device model”
- hypervisor intercepts mmio regions
- forwards request to qemu
- i.e.: read request to harddisk
  
- VM emulates the following devices
- requires the VM's operating system to install, load, and run its native device drivers
  - Network card: AMD PCnet, NE2000
  - Disk drive: IDE
  - Graphics card: Cirrus Logic\* GD5446, VESA-compliant VGA
  - Input: PS/2 mouse and keyboard
  - Sound: Creative\* Sound Blaster 16, Ensoniq\* ES1370

# Intel Pre- and Post-VT


Virtualization Technology: Pre- and Post-VT



Source: Intel, 2006

# AMD IOMMU

- in DomU OS not loaded at address 0
- Xen: direct access to memory difficult-->corruption
- hypervisor intervenes in I/O, apply translation-->overhead
- solutions:
  - rewrite graphics driver ?
  - HW to support IOMMU
- AMD IOMMU -provides isolation and memory protection
- IOMMU: device remap address accessed by HW,




# Ring Buffers

Network, disk I/O is done via shared memory, asynchronous ring buffers.


One ring per VM reduces cross-talk.

Events replace hardware interrupts for notifications.

Page-for-page swap between VM and Xen.



# Xen Architecture – Simple View




Pd = Physical Device Driver  
Vd = Virtual Device Driver



# Xen Architecture – Simple View

Dom0 is the management domain for Xen guests

- controls compute resources dynamically (e.g. memory, CPU, I/O)
- provides interfaces to the physical server.
- provides administration tools and interfaces




Pd = Physical Device Driver  
Vd = Virtual Device Driver



# Xen Architecture – Simple View

Dom0 is the management domain for Xen guests

- controls compute resources dynamically (e.g. memory, CPU, I/O)
- provides interfaces to the physical server.
- provides administration tools and interfaces



Hypervisor

- is the virtual interface to the hardware – virtualizes the hardware
- manages conflicts caused by OS access to privileged machine instructions

Pd = Physical Device Driver  
Vd = Virtual Device Driver






# Xen Architecture – Simple View

Dom0 is the management domain for Xen guests

- controls compute resources dynamically (e.g. memory, CPU, I/O)
- provides interfaces to the physical server.
- provides administration tools and interfaces



DomU is the guest OS

- hosts the application workloads
- typically uses virtual device drivers to connect to the physical drivers in Dom0 by the hypervisor.
- can also use physical device drivers directly
- can be stored in a file-image

Hypervisor

- is the virtual interface to the hardware – virtualizes the hardware
- manages conflicts caused by OS access to privileged machine instructions

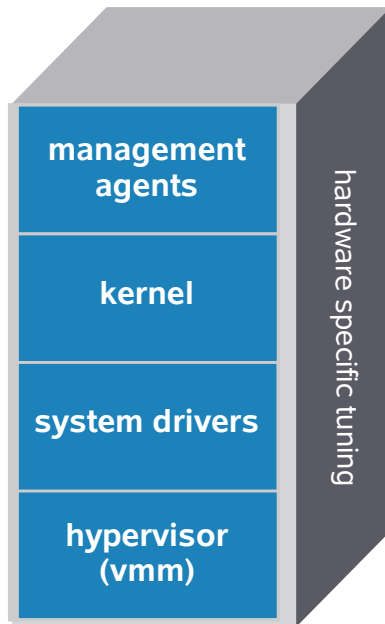
Pd = Physical Device Driver  
Vd = Virtual Device Driver



# Distribution


- pDistro

- Tuned thin, platform distribution specialized for specific physical hardware

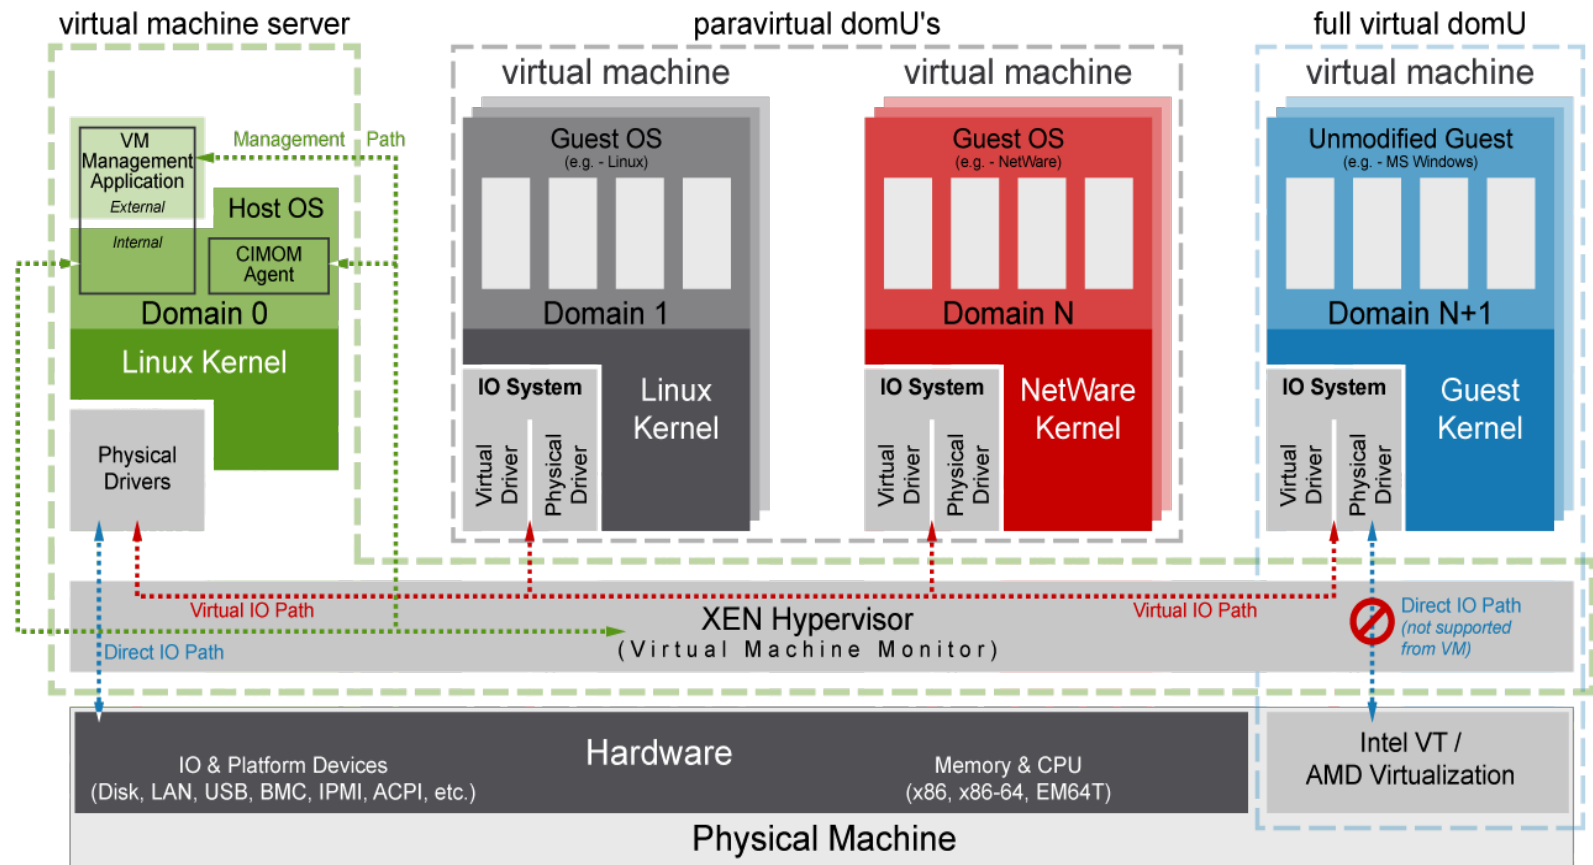


- vDistro

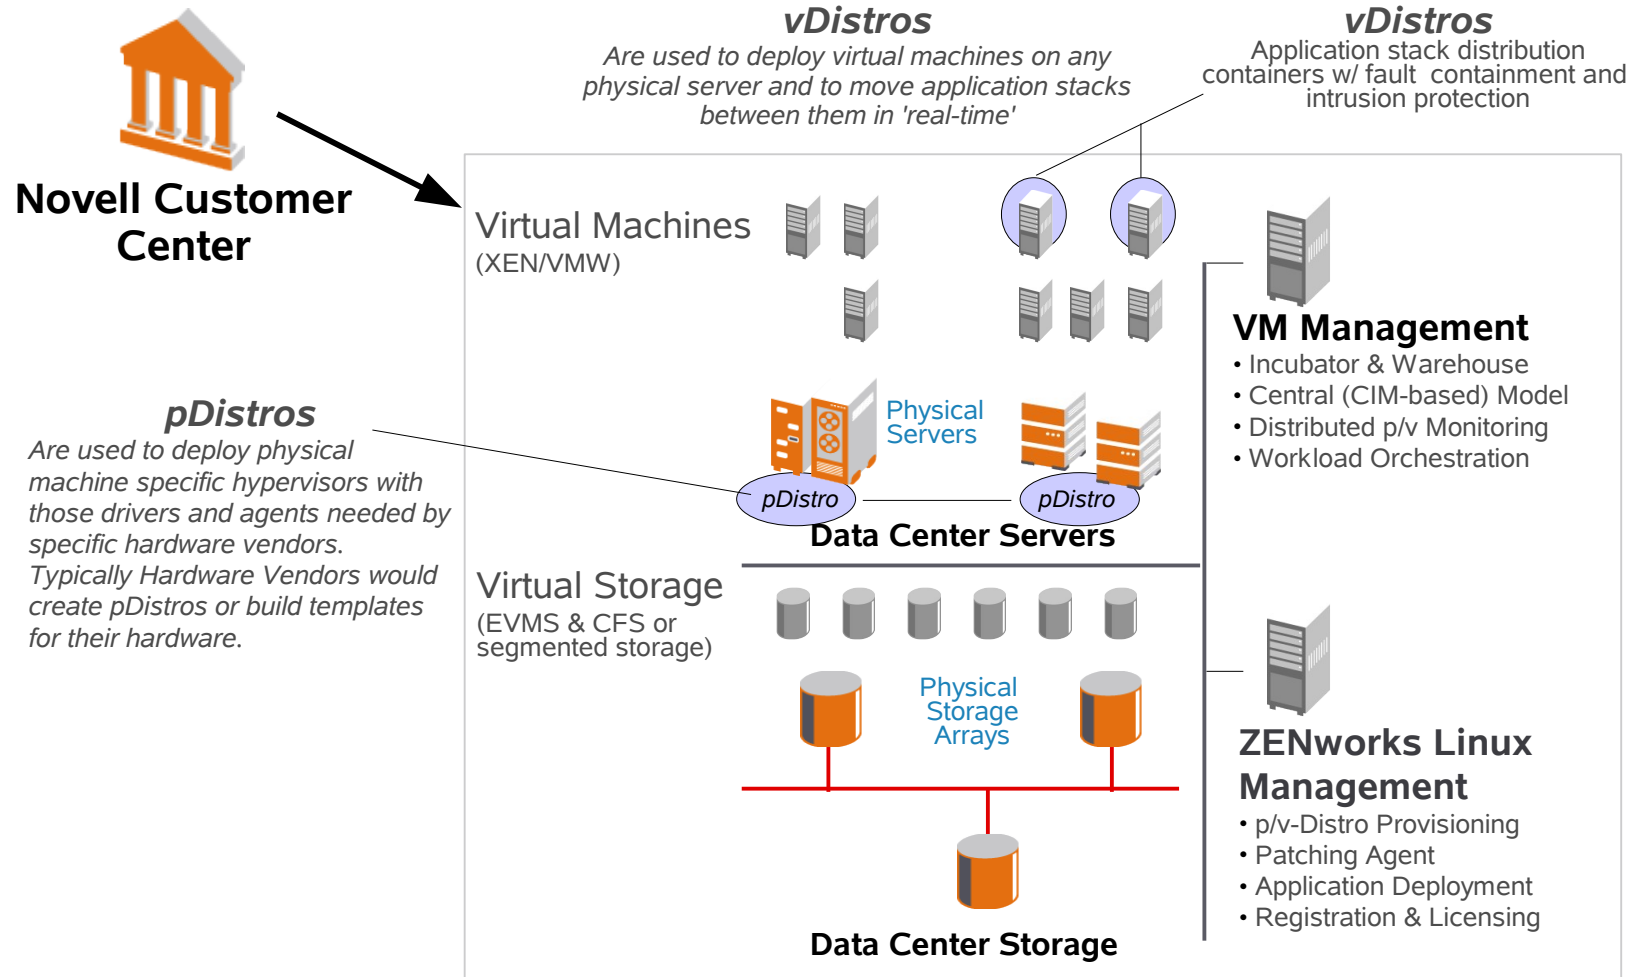
- Tailored for application stacks, able to run on any pDistro



# Xen Technology – Architecture (cont)



# Emerging Virtual Architectures



**Monitor / Manage / Analyze / Respond**

# VM Installation

# VM Installation using YaST

Novell makes VM installation simple with YaST

## VM Installation using YaST on SLES 10

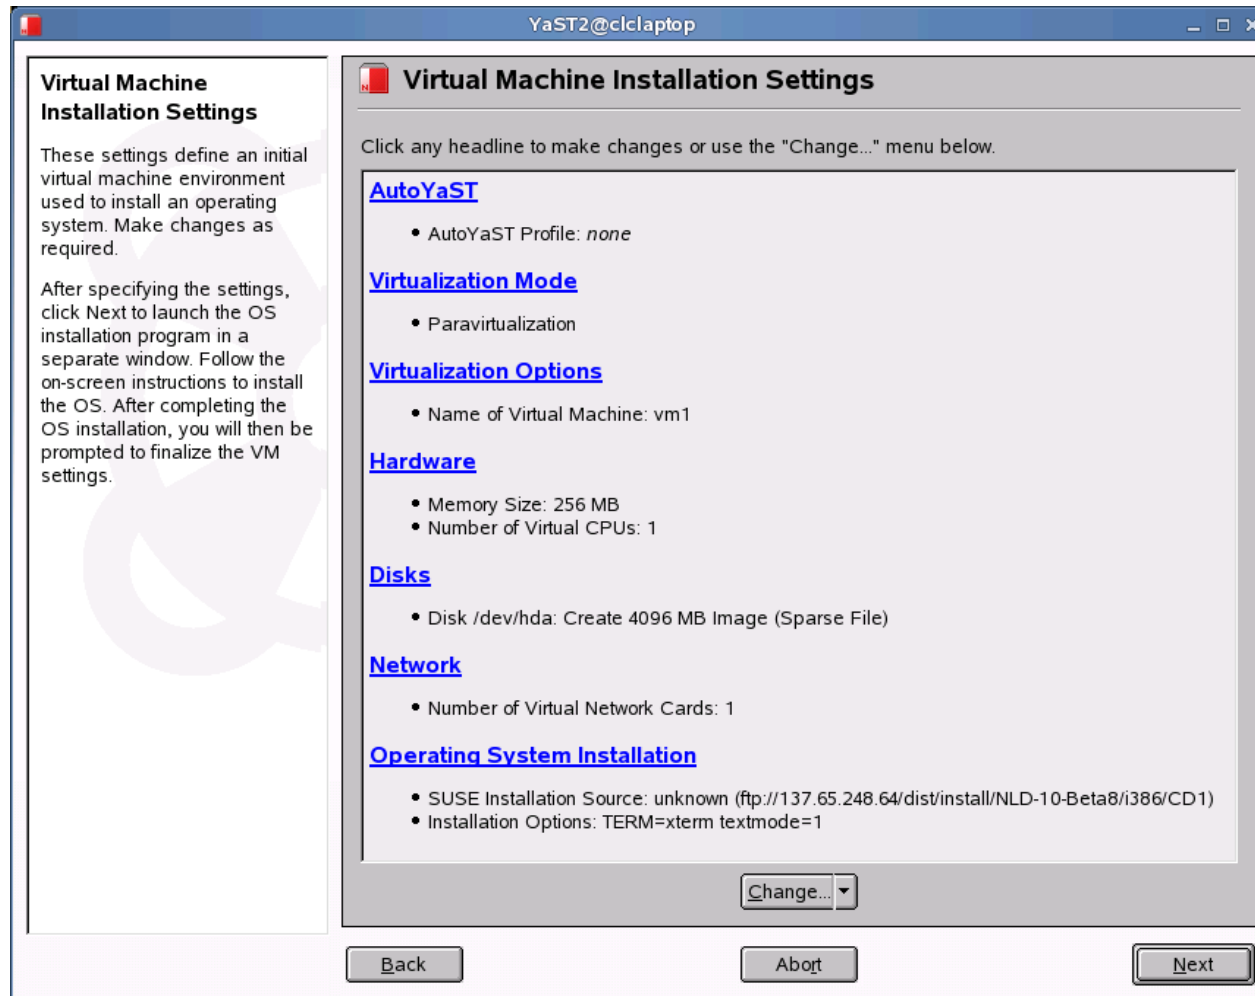
- Support for para-virtualized guests
- Support for fully virtualized guests
- Support for new installation or use of existing disk
- Install over network (SUSE) or from CD/DVD (SUSE & others)

Support for simple lifecycle management:

- Start
- Stop
- View
- Shutdown

# SLES 10

## VM Installation Using YaST



# VM Disks – Supported Storage

- **mapping “virtual” disk (within the VM) and a “physical” disk (within a driver domain).**
- Generic Block Devices (SCSI, SATA, IDE, LVM, ...)
- iSCSI Target (Server), iSCSI Initiator (Client)
- FibreChannel Protocol, SAN
- Native Multipath IO (Kernel Devicemapper - DM-MPIO)
- Multidisk (MD)
- Oracle Cluster Filesystem (OCFS2), Heartbeat2
- Enterprise Volume Manager (EVMS2)
  - Novell® iSCSI, OCFS2 and Heartbeat2 Plugins



# How Does Xen Compare?

- Performance: 1-20% actual overhead (older claim: 1-5%)
- Scalability: 10s of Vms
- Memory overhead: Negligible for para. ~12MB per VM for full.
- Hardware-enforced fault isolation between VMs: Yes
- Ability to dynamically grow/shrink hardware resources: Yes
- Open source
- Commodity hardware
- Ease of deployment: Use YaST
- Support: From Novell

The background of the slide is a solid blue color with a pattern of diagonal stripes in varying shades of blue, creating a sense of motion or depth.

Using Xen

# Booting VM Server (Xen)

If you selected the Xen pattern during installation, Grub should be correct

Installing kernel-xen or kernel-xenpae later should update Grub, also

Select “XEN” boot option from Grub

Grub loads Xen, Linux kernel, and initrd. Xen initializes, then hands control to Linux kernel

# Troubleshooting Boot Loader

Copy/modify the normal entry in /boot/grub/menu.lst:

```
title Xen
    root (hd0,0)
    kernel /boot/xen.gz dom0_mem=65536
    module /boot/vmlinuz-xen root=/dev/hda1 vga=0x31a selinux=0
splash=silent resume=/dev/hda2 elevator=as showopts
    module /boot/initrd-xen
```

Note:

- dom0\_mem is in KB. (Optional. Defaults to taking nearly all memory.)
- Copy kernel args verbatim
- Specify “module” not “initrd”

# Daemons

xend:

- Required
- Runs in VM Server
- Must be running to access management commands
- Manages VM configuration

xendomains:

- Optional
- Starts (or restores) VMs when VM Server starts
- Stops (or saves or migrates) VMs when VM Server stops

# VM Configuration Files

Use an example file as a template:

```
cp /etc/xen/examples/xmexample1 /etc/xen/vm/MyVm
```

Fields to change:

if manual kernel:

**kernel:** The location of the kernel, as visible from dom0.

**initrd:** (\*nix-centric) Extra module for kernel, as per Multiboot specification.

if domU loader:

**bootloader:** /usr/lib/xen/boot/domUloader.py

**bootentry:** paths of kernel and initrd to extract from VM's disk.

**memory:** Memory to give the guest, in megabytes.

**name:** The name must be unique among running VMs.

**vif:** List of virtual network interfaces.

**disk:** This maps the disk device (visible in domain 0) to the device in the guest.

**root:** (\*nix-centric) Device containing root filesystem.

**extra:** (\*nix-centric) Extra kernel arguments, e.g., runlevel.

**i.e.: /etc/xen/vm/sles10pv**

```
disk = [ 'file:/var/lib/xen/images/sles10pv/hda,hda,w', 'phy:/dev/hdc,hdb,r' ]
memory = 256
vcpus = 1
builder = 'linux'
name = 'sles10pv'
vif = [ 'mac=00:16:3e:41:09:60' ]
localtime = 0
on_poweroff = 'destroy'
on_reboot = 'restart'
on_crash = 'restart'
extra = ' TERM=xterm'
bootloader = '/usr/lib/xen/boot/domUloader.py'
bootentry = 'hda2:/boot/vmlinuz-xen,/boot/initrd-xen'
```

# /etc/xen/examples/xmexample.iscsi

```
kernel = "/boot/vmlinuz-xen"  
ramdisk = "/boot/initrd-xen"  
memory = 128  
name = "nbd"  
# Please change MAC  
vif = [ 'mac=00:16:e3:10:00:a0, bridge=xenbr0' ]  
# Replaced the ':' in the iname with '@'  
disk = [ 'iscsi:iqn.2006-09.de.suse@0ac47ee2-216e-452a-a341-  
a12624cd0225,hda,w']  
dhcp = "dhcp"  
hostname= "nbd"  
root = "/dev/hda1"  
extra = "3"
```



# VM Migration

## Migration

- Moving a VM from one physical machine to another, without interrupting the VM's state.
- IP address follows the VM, so networking is not interrupted

## Live Migration

- Migration with minimal down-time (10's of ms)
- Streams VM's pages over network while VM is still running

Disk must still be visible on destination! (e.g., use iSCSI)

# Managing VMs with the xm Command <sup>N</sup>

Must be root.

Use “xm” command (Xen Management)

Querying:

`xm list`

`xm list -l`

`xm console id` (disconnect from VM's console with Ctrl-])

Managing:

`xm create -c configfile`

`xm shutdown id`

`xm destroy id`

`xm migrate id destination-host`

`xm migrate -l id destination-host`

# More xm Commands

Hot-plugging:

xm mem-set

xm block-attach

xm block-detach

Debugging:

xm info

xm top

xm dmesg

For more information:

man xm

xm help --long

# Files & Directories

<code>/etc/xen/xend-config.sxp</code>	Machine-wide configuration
<code>/etc/xen/vm/</code>	VM definition files
<code>/etc/xen/auto/</code>	VMs to auto-start
<code>/etc/xen/examples/</code>	Example VM definition files
<code>/etc/xen/scripts/</code>	Scripts (see <code>xend-config.sxp</code> )
<code>/var/lib/xen/images/</code>	VM disk images
<code>/var/log/</code>	Log files, when things go wrong
<code>/usr/share/doc/packages/xen/README.SuSE</code>	Tips & late-breaking information

# Convert VMware image to Xen image


```
dd if=/dev/zero of=disk.img bs=32k count=32k
#Virtual Hard Disk to Network Block Device mapper
vmware-loop -p SLES9-0.vmdk
modprobe nbd
vmware-loop SLES9-0.vmdk 2 1
#Extract VMware partition from the device to the image file SLES9.img:
dd if=/dev/nbd1 of=/tmp/SLES9.img bs=32k
#inject the root FS into the Qemu image.
file disk.img
disk.img: x86 boot sector, FREE-DOS Beta 0.9 MBR
fdisk disk.img
..
Units = cylinders of 16065 * 512 = 8225280 bytes
..
  Device Boot      Start         End      Blocks   Id  System
disk.img1          1          15       120456   82  Linux swap / Solaris
disk.img2         16         130      923737+   83  Linux
fdisk disk.img
#need to skip swap space, swap space ends at 16065*(16-1) bytes
16065*15=240975
dd of=disk.img seek=240975 if=/tmp/SLES9.img
#fdisk: make partition active, to boot from it
```

The background of the slide is a solid blue color with a pattern of diagonal stripes in varying shades of blue, creating a sense of motion or depth.

# Case Studies

# Case Study: HA Web Server


- Small office solution
  - High available web server for non-disruptive online business
    - > Redundant setup using 2 machines with SLES 10, LAMP stack and heartbeat2: Xen based apaches, remote storage access via iSCSI and OCFS2, web shop application
    - > Automated load balancing and high availability, easy extension of compute and storage resources



# Case Study: HA mail, file and print

- Small office solution

- Small scale consolidation with increased throughput and improved fault tolerance
  - > Past: 4 dedicated servers for mail, file, print and an Internet gateway
  - > Today: 2 state-of-the-art servers running mail, file, print and Internet gateway in dedicated and thus isolated VMs, DMZ realized with firewall VM, all VMs instrumented with Heartbeat v2 to fail-over services on standby virtual machines.
  - > More reliability, fault tolerance and thus high level of availability for business critical IT services. Improved performance by using latest, but still cost effective hardware, flexibility to do 'rolling' upgrade if performance requirements increase over time (services stay online)





# Novell Xen Roadmap

# Novell Virtual Server Architecture

## Overview

- SUSE Linux Enterprise Server 10 includes Xen 3.0.2
- Supported hardware architectures: x86 (32bit) / x86 PAE (32bit) / x86-64 (64bit) including features to run unmodified OS
- Unmodified OS operation (full virtualization) requires Intel VT and AMD Virtualization (AMD-V) hardware
- One IO and system management VM, many application VMs
- Tight cooperation with virtualization technology providers to maintain stable interfaces between VM Server and VMs
- Virtual Machine Server and Virtual Machine are verified layers by Novell YES certification
- Certified systems, supported solution stacks and their combinations will be listed in YES certification bulletins

# SUSE Linux Enterprise Server 10

## Xen 3.0.2 support status



### Overview

- Runs both SLES 10 VM Server / SLES 10 VM (32/32bit and 64/64bit)
- VM server RAM support: x86 4GB, x86 PAE 16GB, x86-64 32GB
- Multiple vCPUs, NICs and disks for VM Server
- Multiple vCPUs, vNICs and vDisks for VM
- Technical preview allows to run different unmodified OS if hardware support is present (full virtualization)

### Maintenance update Oct 2006

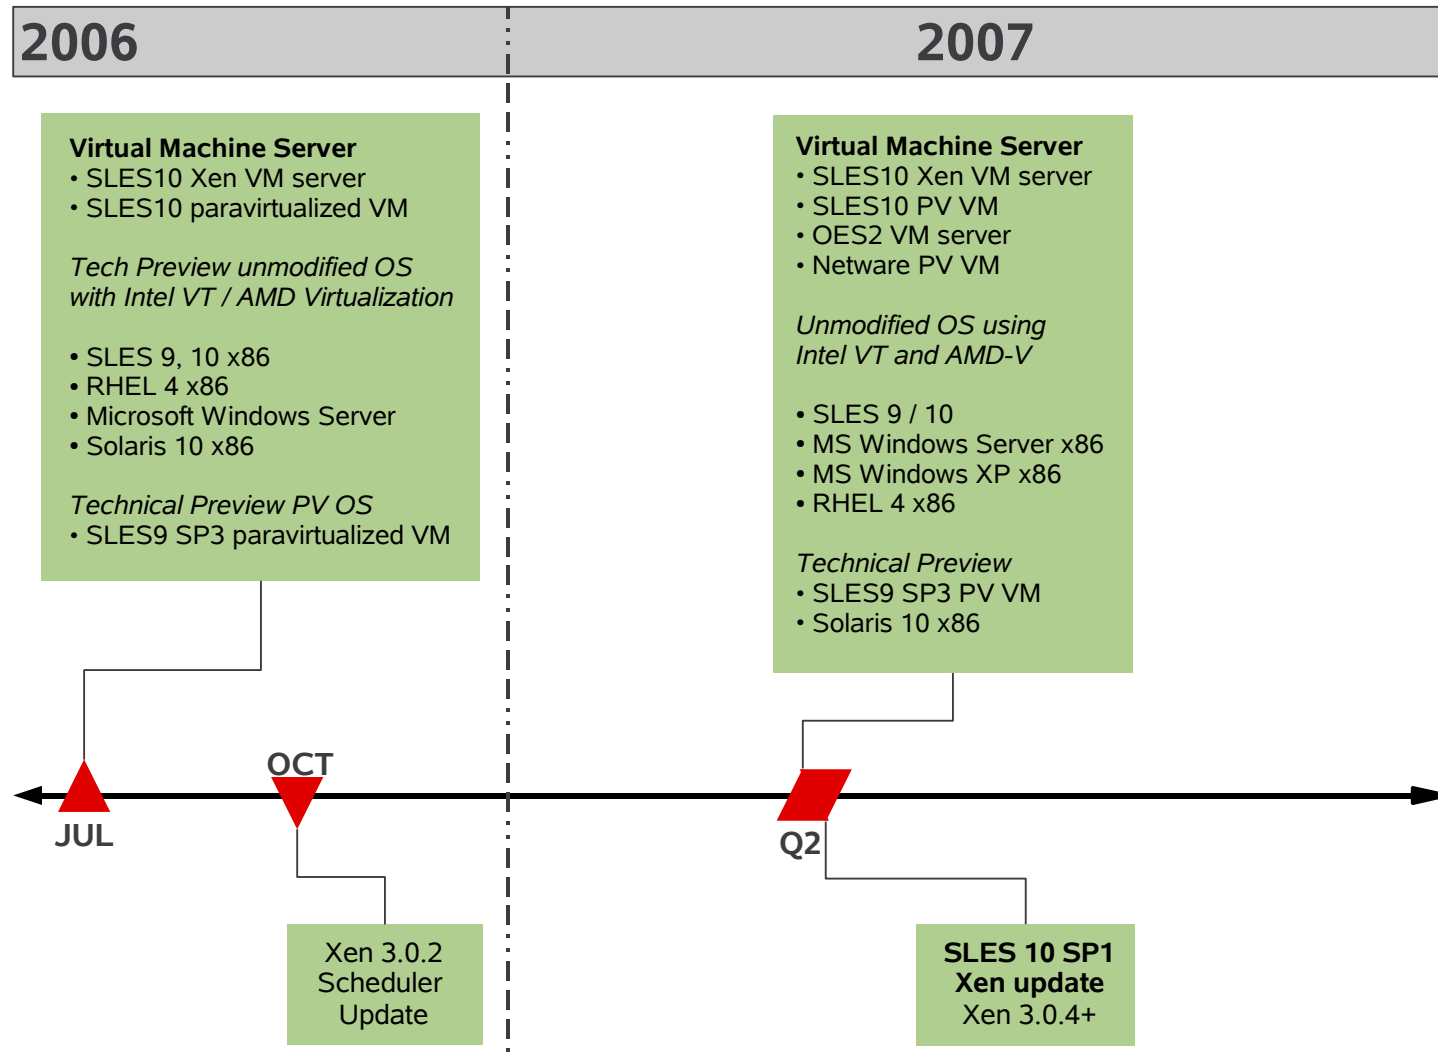
- Updated hypervisor to credit scheduler of Xen 3.0.3 for increased performance and scalability
- Several fixes to improve full virtualized OS installation and support
- more minor fixes (see changelog)

# Xen 3.0.2 support status (2)

## Overview – Xen unsupported features

- F: Save & Restore (plan: get to supported status with a maintenance update)
- F: Live Migration (plan: get to supported status with a maintenance update)
- F: Use of asynchronous IO with sparse file mounted loopback in Virtual Machine Server (dom0) (plan: get to supported status with a maintenance update)
- F: Direct physical device access from Virtual Machines
- Support: technical preview full virtualization support: track bugs based on best effort

# Roadmap - Xen



# **SUSE Linux Enterprise Server 10 SP1 N**

## **Xen 3.0.4+ planned support Q2/2007**

### Overview

- CIM provider for VM lifecycle management, monitoring and resource management
- CIM client for VM lifecycle management, monitoring and resource management
- Update from Xen 3.0.2 to 3.0.4+ (due to system management)
- SLES 9 PV installation support (preview / documented)
- OES2 using Netware 6.5 SP7 PV
- VM server RAM support: x86 4GB, x86 PAE 64GB, x86-64 256GB
- Basic NUMA support

# Xen 3.0.4+ support status (1)

## Overview – planned Xen supported features

- x86 PAE max VM server physical memory: 64GB
- x86 PAE max VM memory: 62GB
- x86-64 max VM server physical memory: 128GB
- x86-64 max VM memory: 126GB
- 32 logical hardware threads per VM Server, 256 VMs per VM Server
- 8 vCPUs per VM, 8 VMs per physical CPU core
- 8 vNIC per VM

# Xen 3.0.4+ support status (2)

## Overview – planned Xen supported features

- Save & Restore of Virtual Machines
- Live Migration of paravirtualized VMs
- use of asynchronous IO with sparse file mounted loopback in Virtual Machine Server (dom0)
- paravirtualized device drivers for IO and network loads for Linux and MS Windows to boost fullvirtualized OS performance



# OES2 virtualization support

- OES2 shares same common code base of SLE 10 SP1
- NetWare runs as paravirtualized OS 32-bit mode on 32bit and 64bit processor architectures that support the x86/x86-64 instruction set.
  - This includes x86-64 (AMD64 / Intel EM64T)
  - In this mode NetWare will be fully functional as a 32bit VM while allowing other virtual machines to fully benefit from the additional address space available on such architectures
- Workgroup servers, NetWare migration to recent hardware, High Availability

# Virtualization Pilot Program

- Intel-VT or AMD-V hardware assisted
- **Full Virtualization** with PV Device Drivers:
  - SLES 9 and RHEL 4
  - MSFT WinXP/2000/2003
- **increase in performance**

cluster aware logical volumes for  
hosting Xen domUs

# cluster aware hosting of Xen domains<sup>N</sup>

- Howto use Heartbeat and EVMS (in dom0) to create cluster aware logical volumes for hosting Xen domUs
  - Heartbeat 2.0.4 or later
  - EVMS 2.5.5 or later
  - hb2-1.0.0.so is the EVMS plugin for Heartbeat2
- 
- Add these lines to your `/etc/ha.d/ha.cf`
  - `respawn root /sbin/evmsd`
  - `apiauth evms uid=hacluster,root`

# cluster aware hosting of Xen domains<sup>N</sup>

## cont

- Start Heartbeat2 and it'll start the EVMS daemon on each node
  - Renewal time frames are not well communicated or defined
- verify the cluster aware behavior of EVMS
  - evmsgui->Settings->Node Administered...
  - all cluster nodes listed
  - EVMS and Heartbeat interact properly

# cluster aware hosting of Xen domains<sup>N</sup>

## Summary

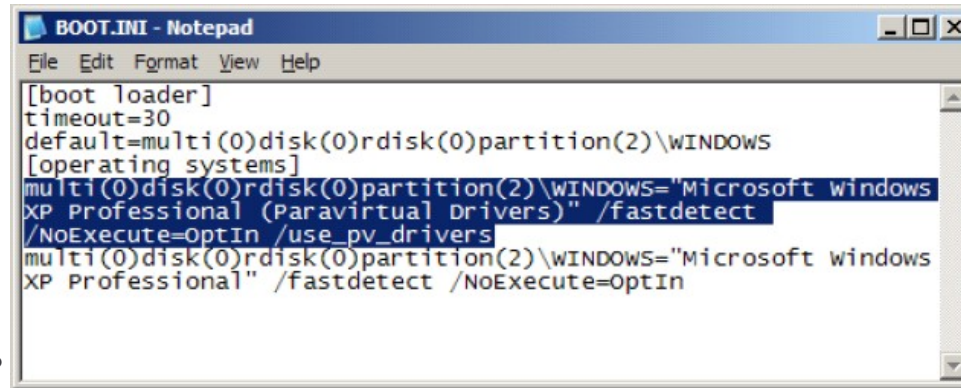
- use EVMS to create a logical volume per each Xen domU OS image
  - e.g. 10G image of hda
  - /dev/evms/MyVirtualMachineHda
  - copy guest OS image into that logical volume
  - logical volume will have same (persistent) device name across all cluster nodes
  - corresponding Xen control file will be location transparent too
  - reference /dev/evms/MyVirtualMachineHda as path for domU's hda on all nodes

> [http://sourceforge.net/mailarchive/message.php?msg\\_id=15485661](http://sourceforge.net/mailarchive/message.php?msg_id=15485661)

The background of the slide is a solid blue color with a pattern of diagonal stripes in varying shades of blue, creating a sense of motion or energy.

PV drivers

# PV Driver for Windows



```
[boot loader]
timeout=30
default=multi(0)disk(0)rdisk(0)partition(2)\WINDOWS
[operating systems]
multi(0)disk(0)rdisk(0)partition(2)\WINDOWS="Microsoft windows
XP Professional (Paravirtual Drivers)" /fastdetect
/NoExecute=OptIn /use_pv_drivers
multi(0)disk(0)rdisk(0)partition(2)\WINDOWS="Microsoft windows
XP Professional" /fastdetect /NoExecute=OptIn
```

- Novell PV Xenbus Driver for Windows
- PCI Device --> Update Driver
  - > Xen Virtual Block Device
- Add a New Hardware Device
- Network Drivers
- Xen Virtual Nic --> Update Driver



# setting up PV drivers for SLES9

- create new SLES9 VM
  - edit /boot/grub/menu.lst
    - > kernel vmlinuz-kernel\_ver append hda=noprobe
  - shutdown VM
  - xm delete vm\_name (reference in xen store)
  - edit /etc/xen/vm/vm\_cfgfile
    - > remove parameter from vif line EXCEPT mac=
  - xm new xm\_cfgfile, restart VM --> new HW found
  - /sys/class/net/ethx/drivers/vif-x
  - verify PV drivers in use:
    - storage device drivers: /sys/block/hdx/device/nodename
    - network card driver: /sys/class/net/ethx/driver/vif-x

# Novell Xen based offerings

# Infos on Pricing & Policies

Easy pricing and policy is in place

- Today we are charging for first instance (physical or virtual)
  - first SLES 10 on physical machine or VMware etc
  - <http://www.novell.com/products/server/virtualization.html>
- Outlook
  - SLES 10 / SLED 10 or vice versa -> pay first server (like current policy)
  - OES2 / Netware PV is planned to pay per user
  - SLES 10 / OES2 PV o. FV / Netware PV – pay first server plus per user
  - OES2 / Netware / SLES 10 – pay first server plus per user

# Microsoft and Novell Collaborating on Virtualization and Interoperability



Collaboration on software to enable SUSE Linux Enterprise Server 10 to run as a fully virtualized guest on SP1 of Virtual Server 2005 R2

- SUSE Linux Enterprise Server 10 add-ins will be available in 2007

Collaboration on software that translates between virtualization technologies

- Windows Server “Longhorn” supporting paravirtualized SUSE Linux Enterprise Server 10 guest
- SUSE Linux Enterprise Server 10 supporting enlightened Windows Server “Longhorn” guest

The background of the slide is a solid blue color with a pattern of thin, light blue diagonal lines radiating from the right side towards the left, creating a sense of motion or depth.

Where Do I find Information ?

# External Resources

## Where do customers get more information

- Novell Virtualization home page
  - <http://www.novell.com/linux/virtualization/>
  - Virtualization Technical Library, Whitepapers
- Novell's pricing for virtualization and available support offerings
  - <http://www.novell.com/products/server/virtualization.html>
  - <http://support.novell.com/linux/>
- Novell online documentation on Xen
  - <http://www.novell.com/documentation/vmserver/>
- Supported hardware and technical limits
  - <http://developer.novell.com/yessearch/Search.jsp>  
search in category “Novell Product: SUSE Linux Enterprise 10 x86 with Xen”
  - [http://support.novell.com/products/server/supported\\_packages/](http://support.novell.com/products/server/supported_packages/) search for SLES10 and Virtualization (Xen Support Matrix)

# Other Resources

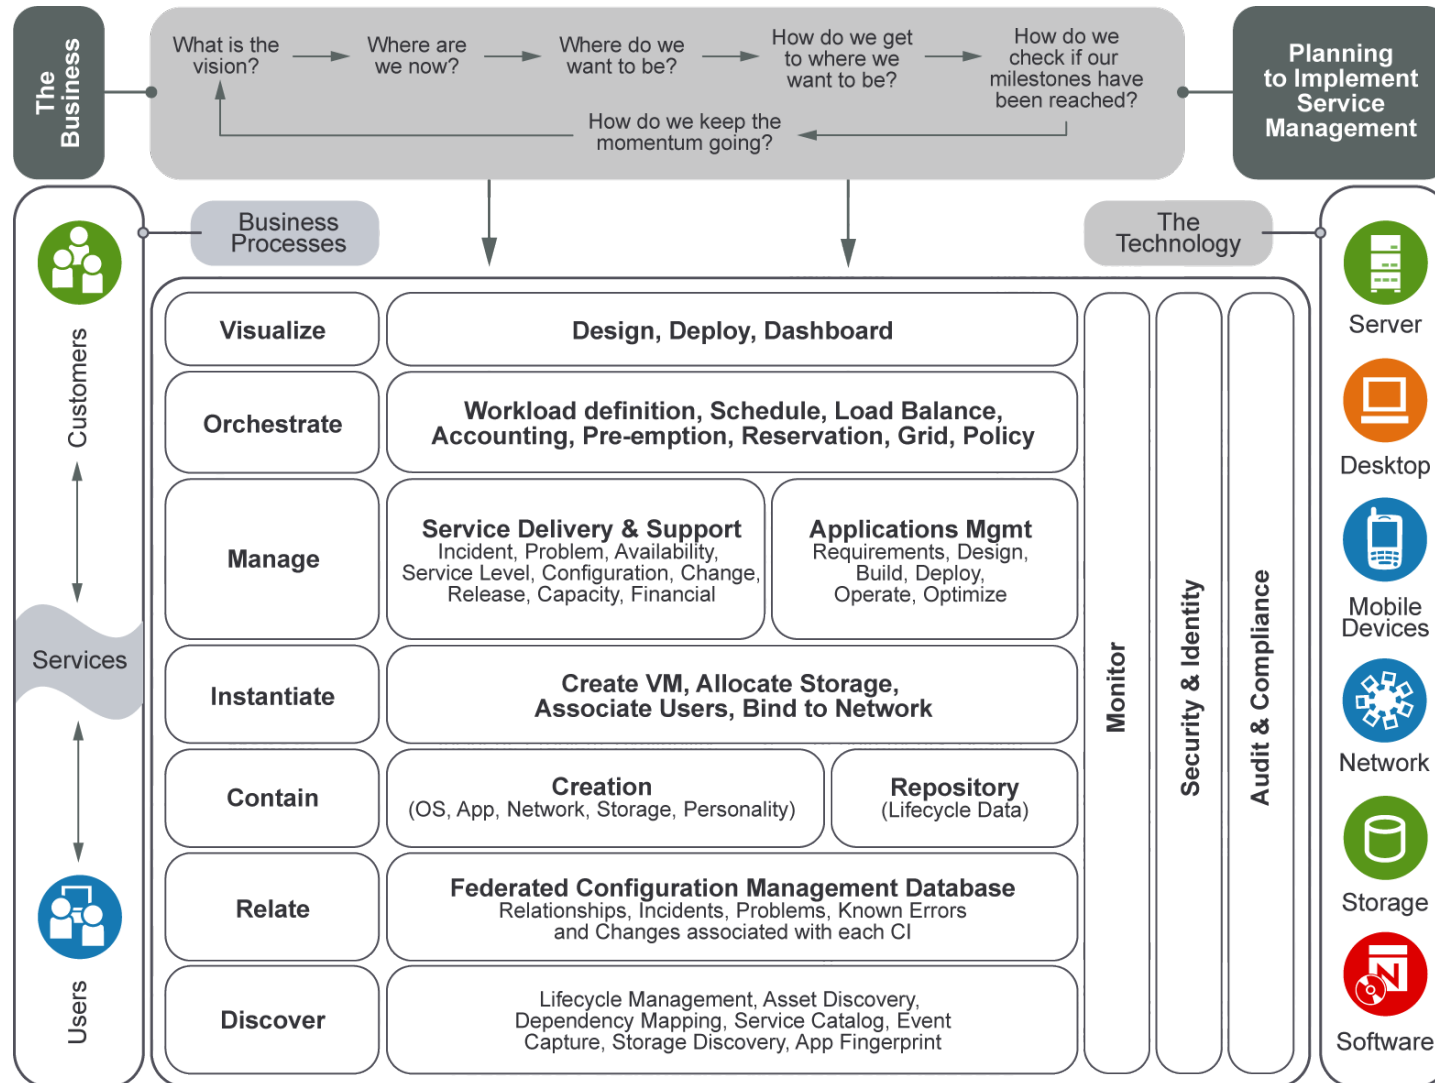
## Recommended links

- Xen project home <http://www.xensource.com/xen/> wiki
  - Xen summit presentations
  - Xen architecture docs, developer discussions, future roadmap
  - Recent changes, indexes, doc, FAQs
  - Deployment examples, third party projects related to Xen
- Latest Novell's Xen externally available technical preview <http://forge.novell.com/modules/xfmod/project/?xenpreview>
  - Latest code for test (SLES 10, SLES9, ...), not production

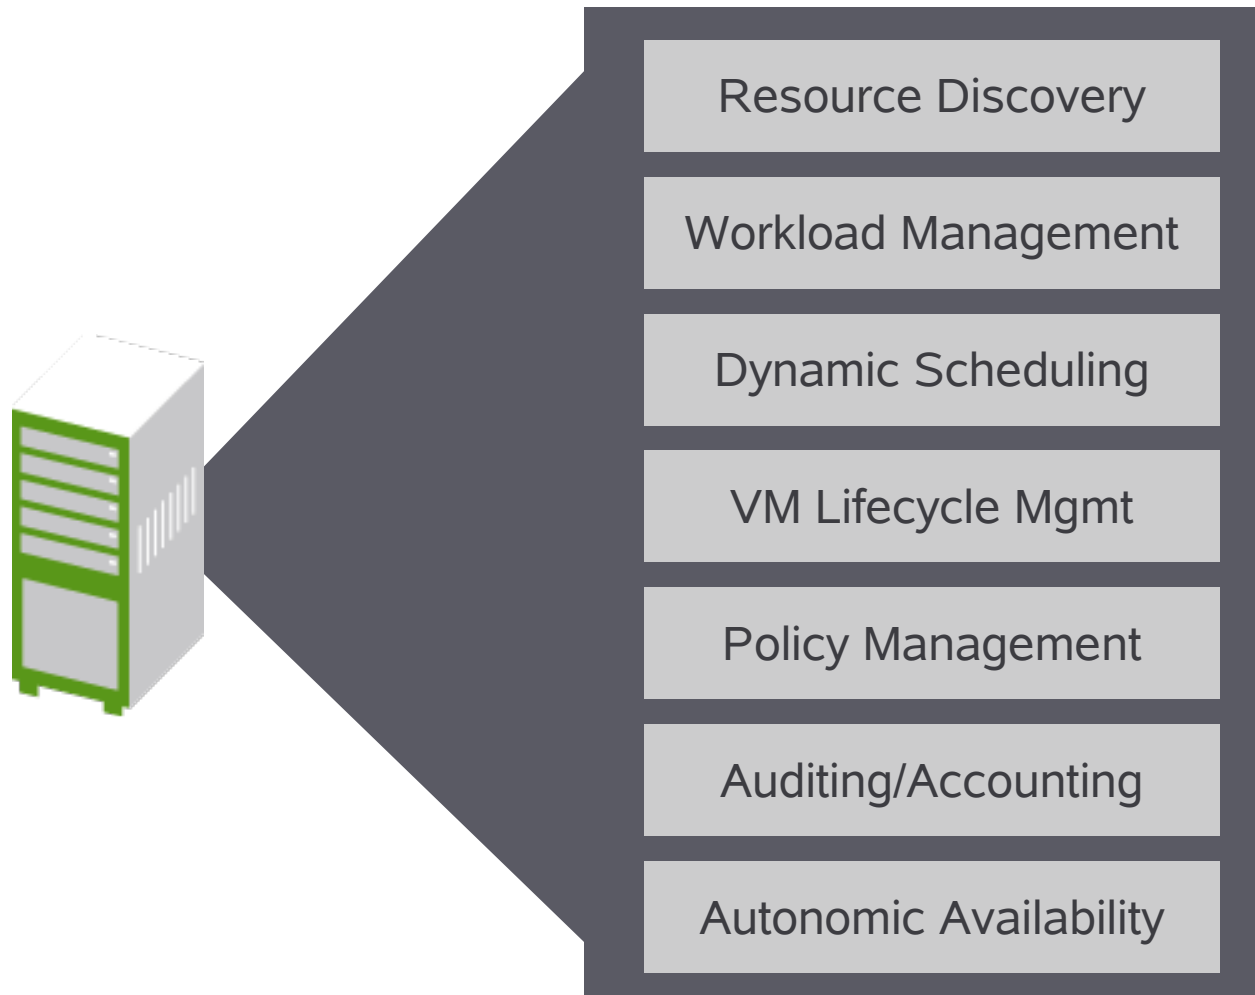
# Outlook: Data Center Automation



# The Solution – Management Blueprint



# ZENworks® Orchestration Server



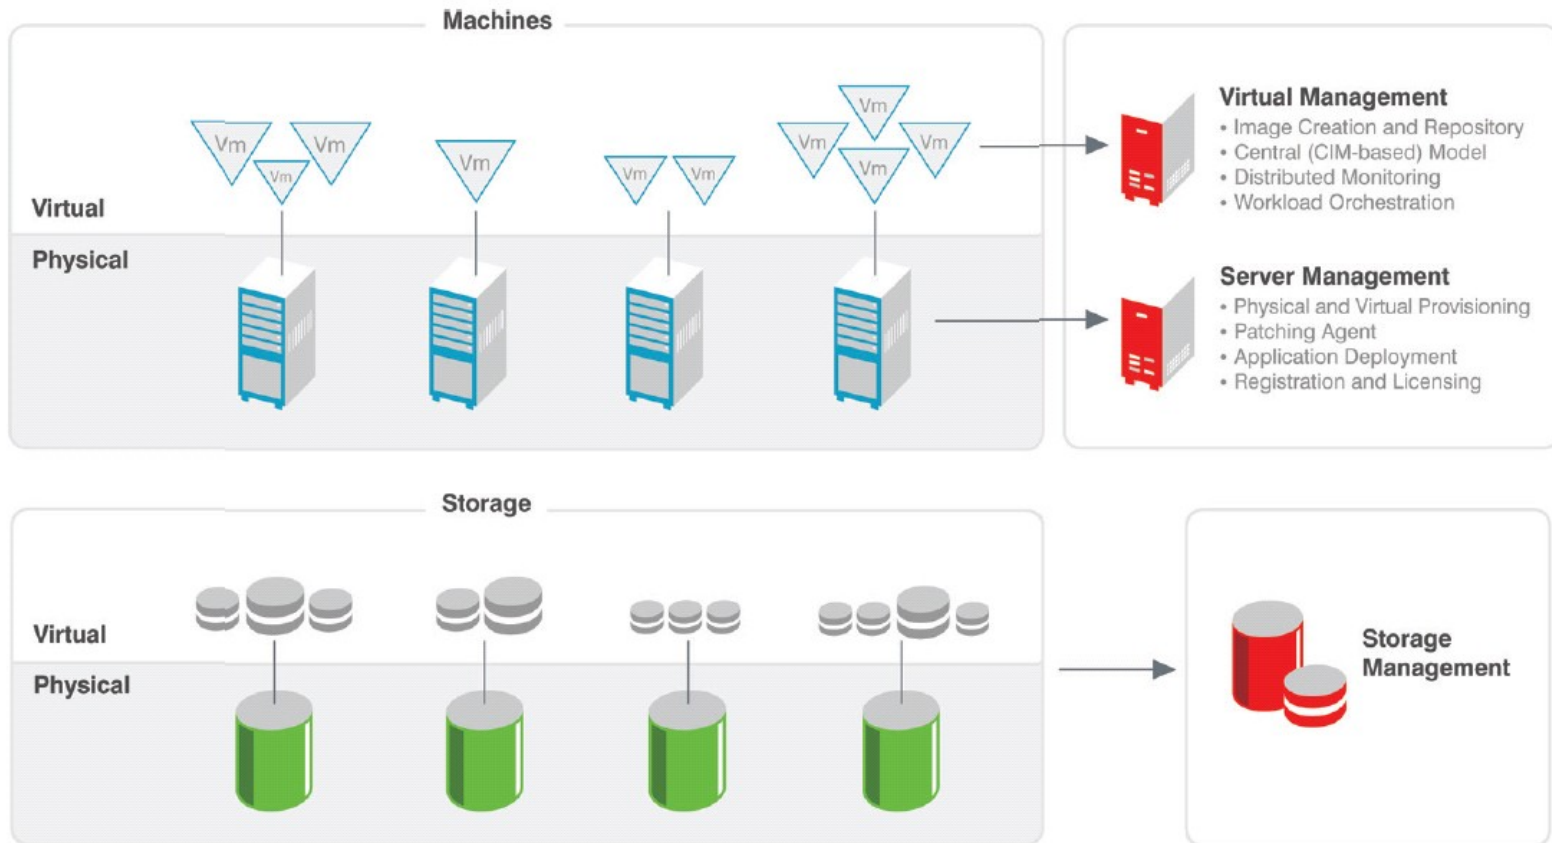
# ZENworks Orchestrator

## Virtual OS

Application stacks with fault containment and intrusion protection

## Physical OS


Comprised of a hypervisor, device drivers and agents needed by specific hardware



# ZENworks Orchestrator

## Features

- Agenten basiertes Entdecken
- Paralleles Ausführen
- Abschätzung der Workloads
- Planen
- Überwachung
  - Nutzdaten, Abrechnung



# ZENworks Orchestrator

## Features - Continued

- Erweiterbare Job Definition
  - Sprache: Python
  - Regeln: XML Format
  - Voreinstellungen historischer Daten
  - Ressourcen Optimierung
- Regelbasiert
  - Reservieren der Ressourcen
  - Vorrang kritischer Aufgaben
  - ereignisorientiert


Orchestration Server



# ZENworks Virtual Machine Management <sup>N</sup>

## Features


- Heterogenes Virtual Machine Management
  - VMware, Xen, Microsoft
  - Entdecken der Server für VM Kommissionierung
  - Off/On-line Vm's & Templates
  - Deployment, re-deployment/rollback
  - Assoziieren physischer, virtueller und Storage Compute Nodes



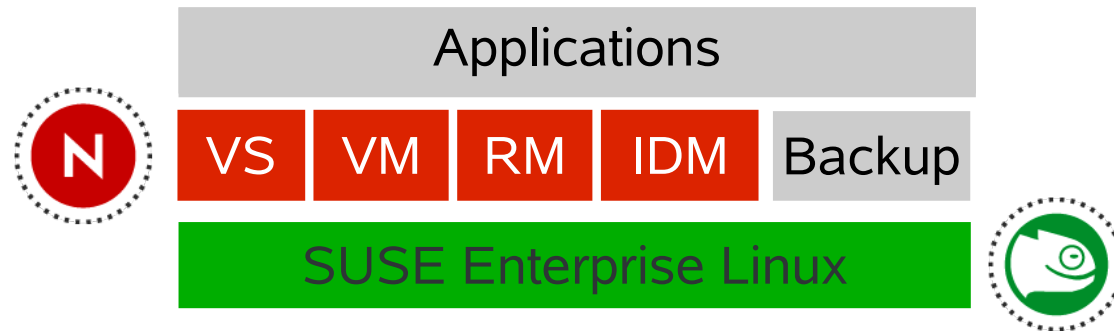
# ZENworks Virtual Machine Management Features



- Cluster-aware Virtualisierung
  - Redundanz, Hochverfügbarkeit, Disaster Recovery
- Regel-basiertes, dynamisches Workload Deployment
  - Job Zuweisung & Regelausführung
  - Verteilen und Bewegen der Workloads dynamisch



# Novell Portfolio Integration



## Virtual Machines

- Consolidation
- Migration
- Effective hardware utilization

## Resource Management

- Scheduling
- Patch management
- Deployment
- Grid

## Virtual Storage

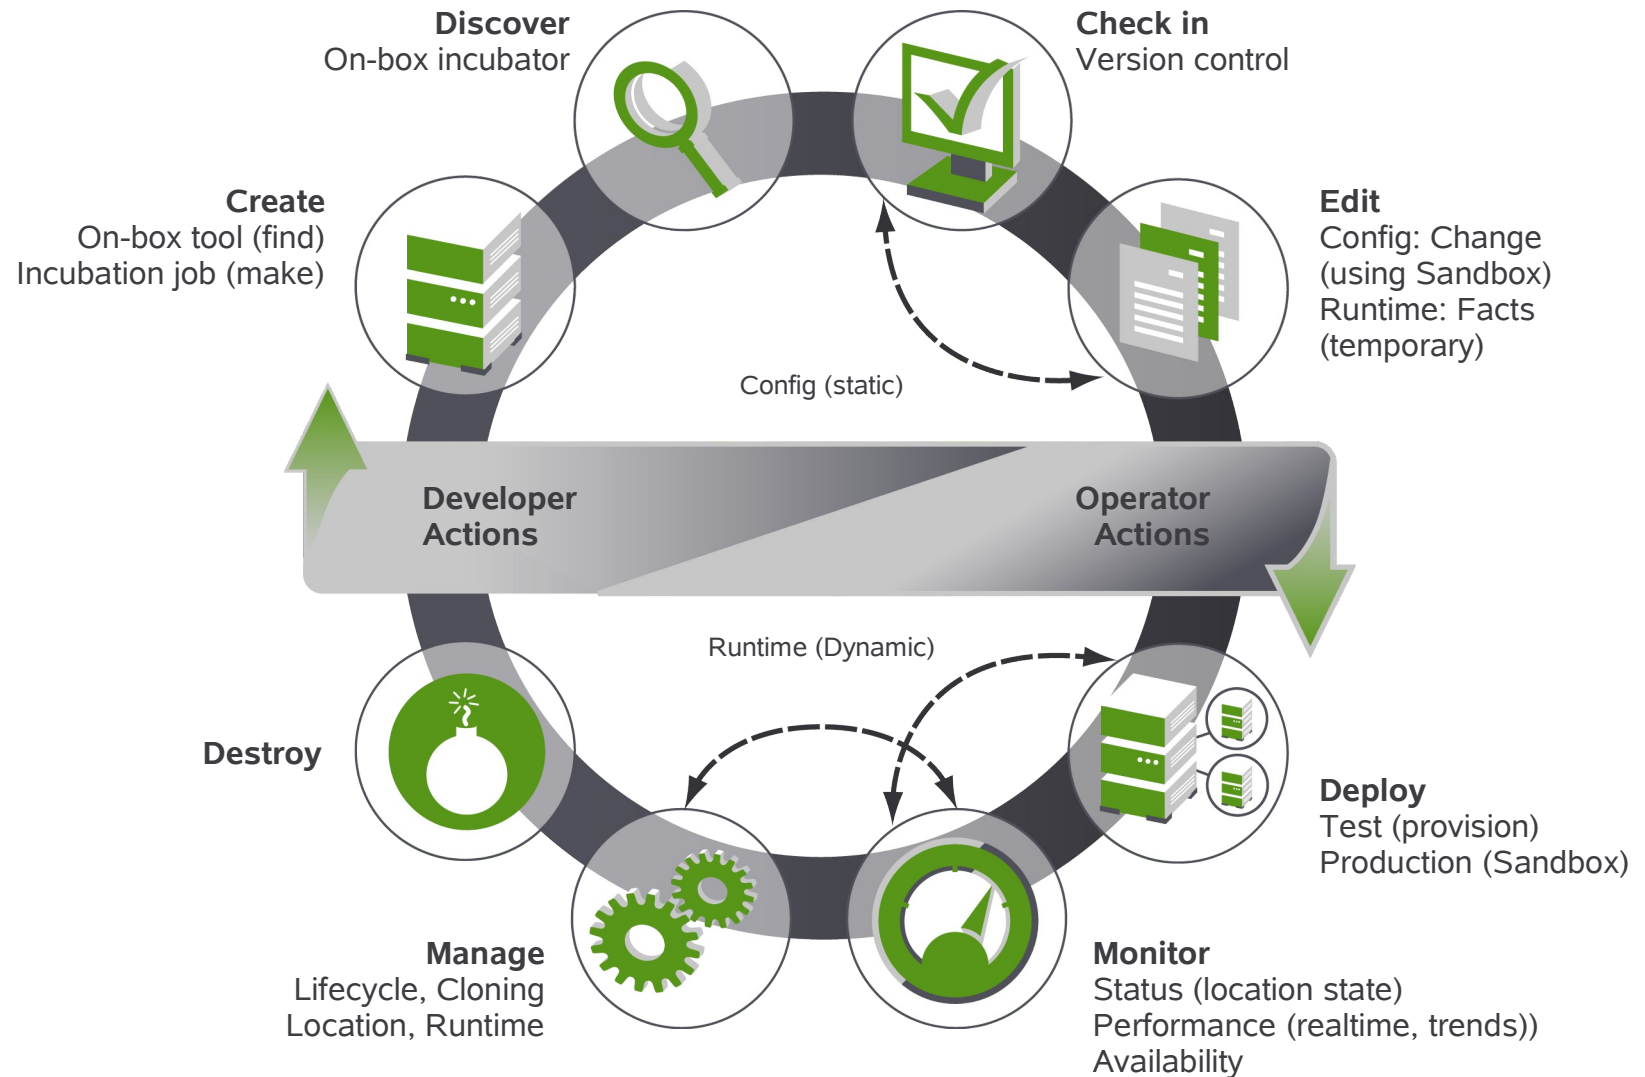
- Volume manager
- High availability software
- Cluster file system
- Business Continuation

## Identity Management

- Users, groups, resource association
- Pervasive infrastructure that relates these services together
- Utility Computing



# Virtual Machine Lifecycle



# Virtual Machine Lifecycle Control

- System maintains library of VMs and images, hosts
  - Like physical resources VMs have ‘facts’ describing attributes
  - VMs can be grouped
- Actions can be performed on VMs:
  - Provision, Shutdown, Suspend, Create Template, Create VM, Create Template From Physical, Create VM From Physical, Clone VM, Clone Online VM, Clone Template, Destroy, Restart, Migrate, Check Status, Template to Instance, Instance to Template, Affiliate With Host, Make Standalone, Checkpoint, Restore, Delete, Cancel Action

# Virtual Machine Lifecycle Control (cont.)

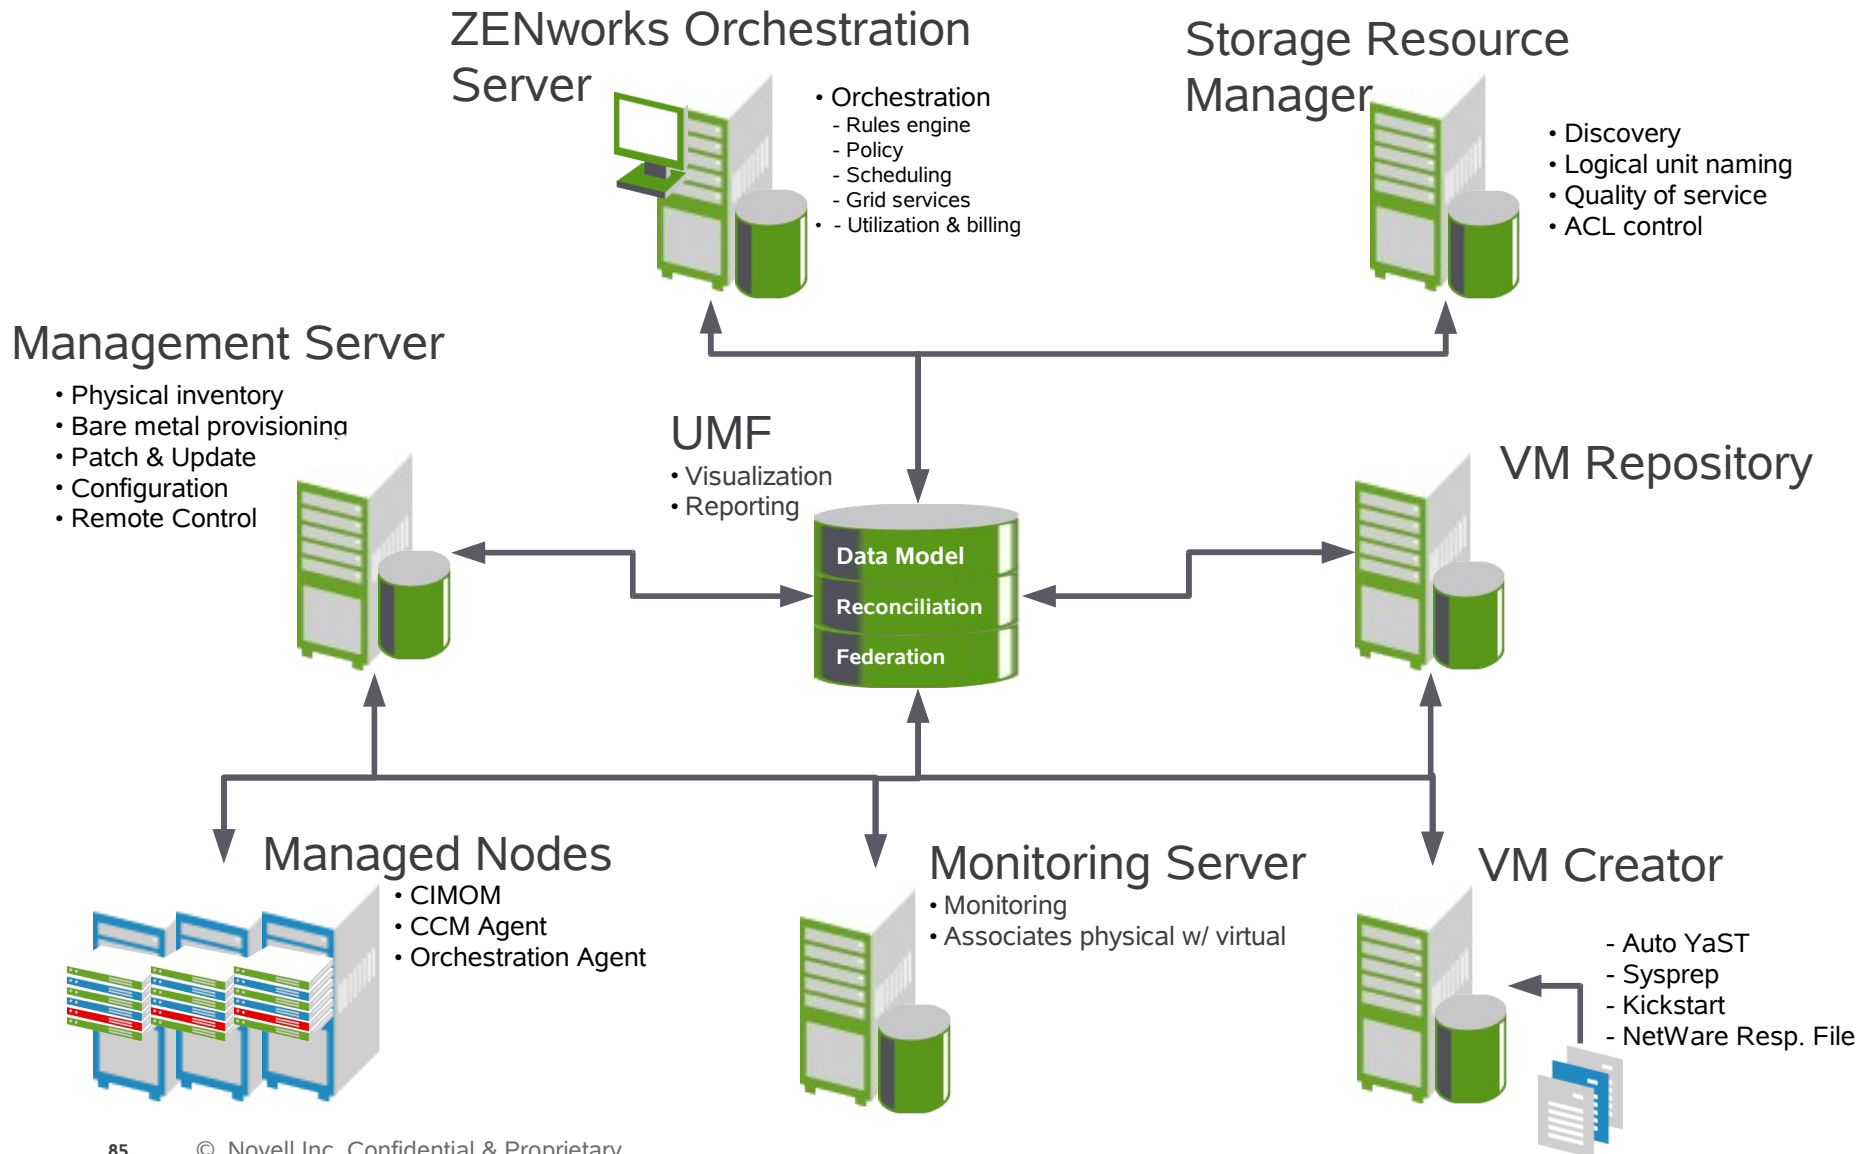
- VM lifecycle controlled by:
  - Programmatically (Job / JDL)
  - Manually (through mgmt. console)
  - Automatically on demand
    - > A job makes request for unavailable resource... Suitable VM image is located, host is located, image is provisioned, instance is initially reserved for calling job, logic is invoked to make use of new resource
- VM host/instance selection/placement is similar to resource selection:
  - Governed by policies, priorities, queues and ranking
- Provisioning Adapters provide VM abstraction:
  - Special ‘provisioning’ jobs perform operations for each integration with different VM technologies
  - **The provisioning adapter is a JOB!**

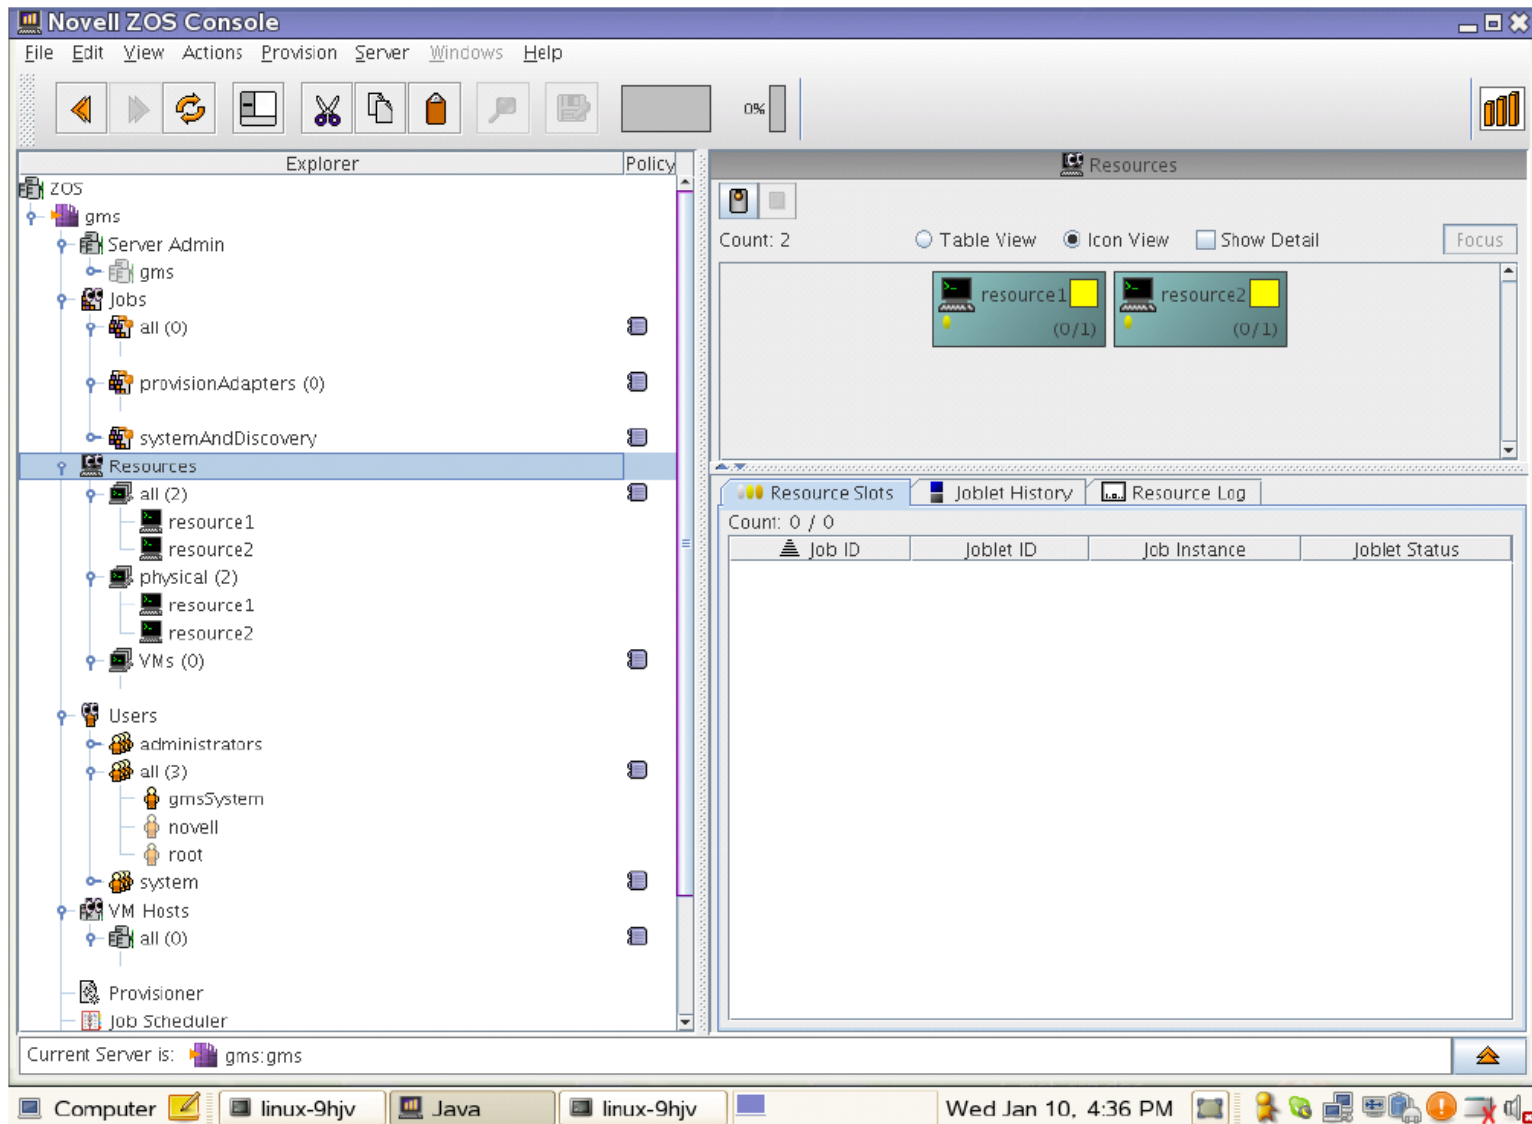
# Virtual Provisioning and Life Cycle Management



- Discovery- provisioning adapter job
- Two types of VM; Instance & Template
- Provisioner requests a VM host
- Reservations, Constraints

# Architektur



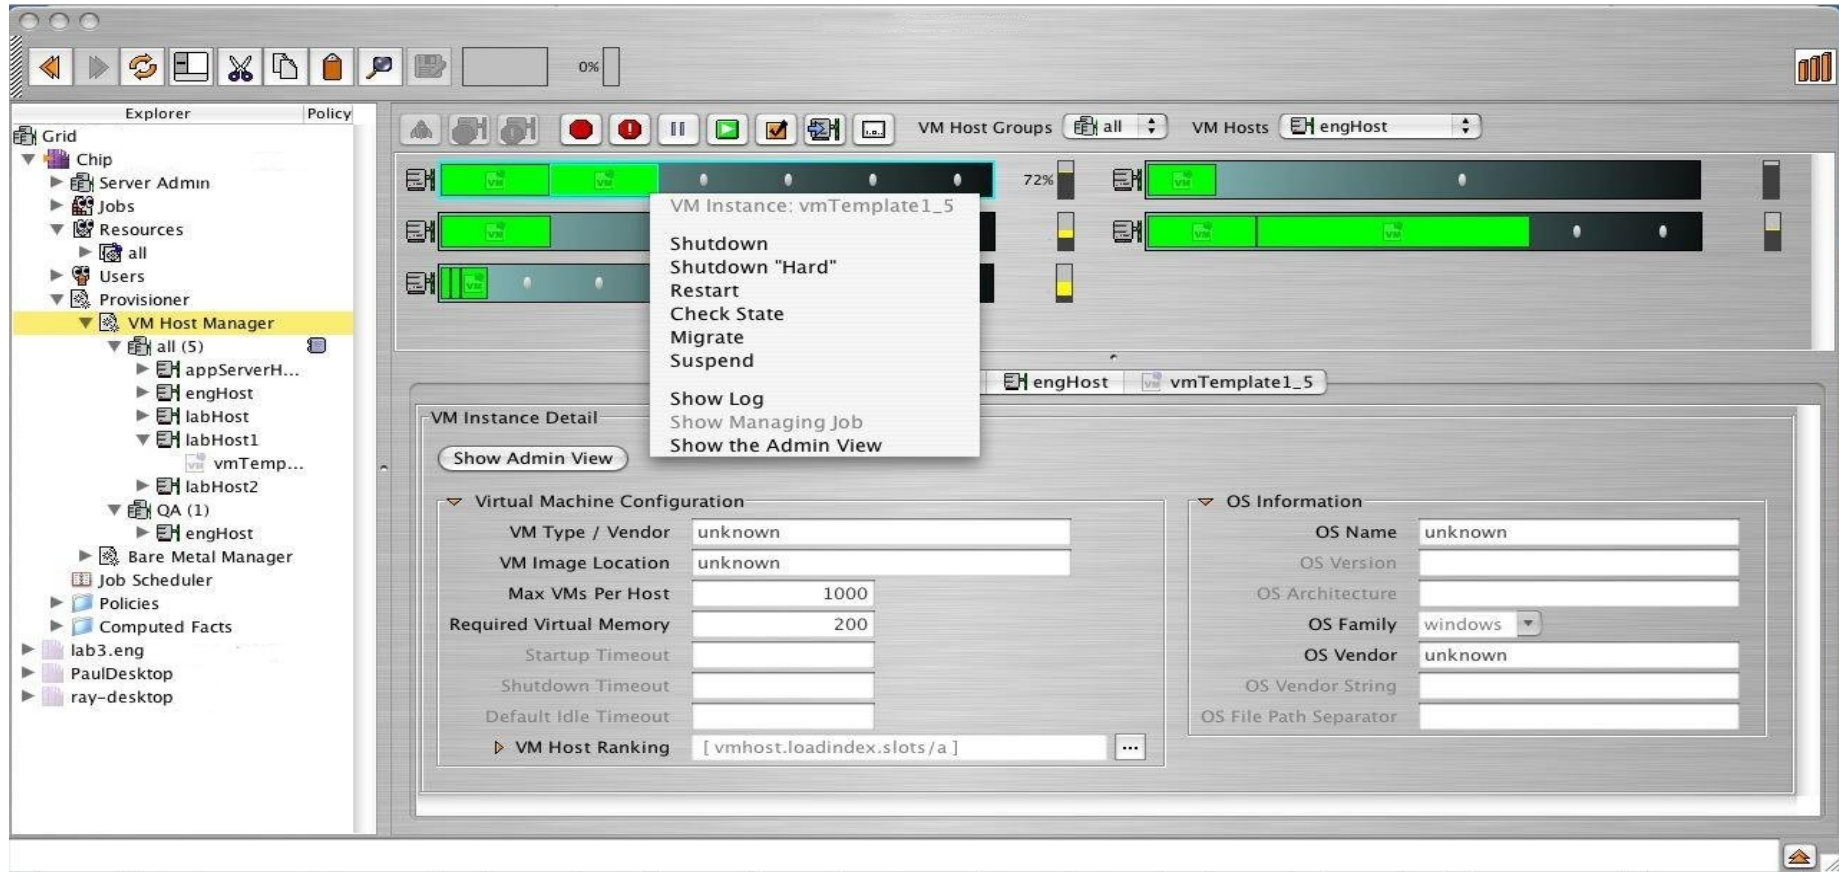


The screenshot displays the Novell Provisioning Server (PVS) management console. The interface is divided into several sections:

- Explorer (Left Panel):** A tree view showing the hierarchy of the environment, including Grid, Chip, Lab, PaulDesktop, Server Admin, Jobs, Resources, Users, and VM Hosts. The VM Hosts section is currently selected.
- VM Hosts (Main Panel):** A table listing various ESX hosts and their resource usage. The hosts are grouped into VM Host Groups (all, esx3) and VM Hosts (esx1, esx2, esx3, esx4, esx5, esx7, vmhost1, vmhost2, vmhost3, vmhost4, vmhost5). Each host has a progress bar indicating its resource usage percentage.
 

Host Name	Resource Usage (%)
esx1	77%
esx2	51%
esx3	66%
esx4	24%
esx5	35%
esx7	71%
vmhost1	13%
vmhost2	95%
vmhost3	62%
vmhost4	79%
vmhost5	-
- VM Instance Detail (Bottom Panel):** A section showing the details of a specific VM instance, currently 'redhat\_4'. It includes a tab for 'Available VMs' and a 'VM Instance Detail' section.
 

Property	Value
Resource Enabled	<input checked="" type="checkbox"/>
Shutting Down	<input type="checkbox"/>
Host Name	
IP Address	
Billing Rate (\$/hr)	1.00
Bill For	walltime
Power Factor	1.000
Load Average	0.000





File Edit View Matrix Provision Server Windows Help

0%

Explorer

- Grid
  - Chip
  - lab
  - PaulDesktop
    - Server Admin
    - Jobs
      - all (3)
        - provision
        - quickie
        - subjobs
      - provisionAdapters (2)
        - provision
        - quickie
    - Resources
      - all
      - production
      - testing
      - VMs (4)
        - redhat
          - websrvr
          - win2k
          - winxp
      - Users
        - all (2)
          - devPortal
          - gmsSystem
        - VM Hosts
          - Provisioner
          - Job Scheduler
          - Policies
          - Computed Facts

Count: 8

Request Time	VM Inst...	Req User ID	Req Job ID	VM Host	State	Current Action	Provisio...	Action Elap...	VM Template/Refere...	Startup Ti...	Reservation	Mode
18:25:50	redha...	gmsSyste...		esx6	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx4	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		vmhost1	VM Active (no agent ...		gmsSy...	n/a	redhat	18:26:02	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx3	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx1	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx5	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx2	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...
18:25:50	redha...	gmsSyste...		esx7	Launching VM	Provision/...	gmsSy...	0:00:20	redhat	n/a	n/a	Automa...

Provision Log Host Assignment Log Autoprep Data

VM has been assigned the vmhost: 'esx1'. Reasoning based on default '[vmhost.loadindex.slots/a]' ranking is summarized below:

Host Name	Log Text
esx1	** Assigned **
esx2	- Possible 2nd choice, but not optimal
esx3	- Possible 3rd choice, but not optimal
esx4	- Possible 4th choice, but not optimal
esx5	- Possible 5th choice, but not optimal
esx6	- Possible 6th choice, but not optimal
esx7	- Possible 7th choice, but not optimal
esx8	VM host is not online
vmhost1	- Possible 8th choice, but not optimal
vmhost2	- Possible 9th choice, but not optimal
vmhost3	- Possible 10th choice, but not optimal
vmhost4	- Possible 11th choice, but not optimal
vmhost5	- Possible 12th choice, but not optimal

## Filter By:

Last 

## Other Views:

[Text Host View](#)

## Datacenter Monitoring Metrics Service

Navigation: [Grid](#) > [Datacenter](#) > [provo.novell.com](#)

Current View: Graphical Host View for Fri, 02 Mar 2007 17:42:41 -0700

## Time and String Metrics



boottime	Tue, 13 Feb 2007 07:33:23 -0700
gexec	OFF
gmond_started	Fri, 02 Mar 2007 17:26:58 -0700
last_reported	0 days, 0:00:18
machine_type	x86
os_name	Linux
os_release	2.6.16.21-0.8-smp
uptime	17 days, 10:09:13
vm_type	

## Constant Metrics

cpu_num	2 CPUs
cpu_speed	2791 MHz
mem_total	2075688 KB
swap_total	2104472 KB

## Node Status

This host is up and running.



**Novell®**

## **Unpublished Work of Novell, Inc. All Rights Reserved.**

This work is an unpublished work and contains confidential, proprietary, and trade secret information of Novell, Inc. Access to this work is restricted to Novell employees who have a need to know to perform tasks within the scope of their assignments. No part of this work may be practiced, performed, copied, distributed, revised, modified, translated, abridged, condensed, expanded, collected, or adapted without the prior written consent of Novell, Inc. Any use or exploitation of this work without authorization could subject the perpetrator to criminal and civil liability.

## **General Disclaimer**

This document is not to be construed as a promise by any participating company to develop, deliver, or market a product. Novell, Inc., makes no representations or warranties with respect to the contents of this document, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this document and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes. All Novell marks referenced in this presentation are trademarks or registered trademarks of Novell, Inc. in the United States and other countries. All third-party trademarks are the property of their respective owners.

