



# System X Virtualization with Oracle VM



## Oracle VM Concepts and Foundations

**Oracle/IBM** Joint Solutions Center  
([oraclibm@fr.ibm.com](mailto:oraclibm@fr.ibm.com))

12/21/2008

# Agenda

- Oracle VM Introduction
  - overview
  - guest OS support
  - concepts
- Oracle VM High Availability (auto-restart/failover)
- Secure Live Migration
- Resource Management
- Oracle VM Templates

# Oracle VM

## Server Virtualization Software for both Oracle and non-Oracle Applications :

- Oracle VM Server – Open Source
- Oracle VM Manager – Private Source
- **The *only* Server Virtualization Software supported and certified with Oracle Products**

- Free to Download
- Free to Distribute
- Free License
- Enterprise-quality Support
- Real World deployment Testing
- Advanced High-Availability features
- Secure Live Migration



# Oracle VM

## Server Virtualization & Management

- Oracle VM contains...
  - **Oracle VM Server**
  - **Oracle VM Manager**
- Oracle VM Server
  - **Open source server software tailored by Oracle**
  - **Installs on “bare-metal” servers from a single CD in about a minute**
  - **x86 and x86\_64 based Intel and AMD Systems**
- Oracle VM Manager
  - **Web browser-based management console**
  - **Java-based management server**
  - **Database repository: XE (incl.), or SE, EE, or RAC**



ORACLE®

---

VM

# Oracle VM

## Guest OS Support

- Support for paravirtualized (“virtualization aware”) OS kernels
  - **Highest performance: Great for I/O scalability**
  - **Enterprise Linux 4 and 5 (from Oracle or RedHat)**
- Support for hardware virtualized (“unmodified”) OS versions (hardware support required)
  - **Red Hat Enterprise Linux 3**
  - **Windows 2003 Server (64bit and 32bit), Windows 2003, and Windows XP**
- 64-bit and 32-bit guests
- Up to 64-way SMP hardware
- Up to 32 virtual processors per guest

# Hardware Support for Virtualization (HVM)

## Where does it fit-in?

- **Goal: efficiently insure VM Oses are not “hurting” each other**
  - E.g. Trying to unsafely modify common resources/state, etc.

<u>Techniques:</u>	<u>Purpose:</u>	<u>Benefits:</u>	<u>Comments:</u>
Paravirtualization (PV) e.g. <b>Oracle VM &amp; Xen</b>	Modify the OS and/or drivers so they know how to behave in a virtual environment	<ul style="list-style-type: none"> <li>• <b>Good-to-excellent performance vs. bare-metal</b></li> </ul>	<ul style="list-style-type: none"> <li>• <b>PV OS kernel req'd (but rapidly not an issue)</b></li> </ul>
Emulation / translation, e.g. <b>VMware ESX</b>	Design the virtualization server to intercept or “trap” harmful requests and/or translate requests into appropriate forms	<ul style="list-style-type: none"> <li>• <b>Use unmodified OS</b></li> <li>• <b>HVM hardware not req'd (but rapidly not an issue)</b></li> </ul>	<ul style="list-style-type: none"> <li>• <b>Poor scalability esp. under I/O load</b></li> </ul>
Hardware virtualization (HVM); e.g. <b>Oracle VM &amp; Xen</b>	Design the hardware so it knows how to handle inappropriate requests itself. (Note: PV drivers can be used with an otherwise unmodified OS, e.g. Windows, to improve performance)	<ul style="list-style-type: none"> <li>• <b>Use unmodified OS</b></li> </ul>	<ul style="list-style-type: none"> <li>• <b>Poor perf. today</b></li> <li>• <b>HVM hardware req'd (but rapidly not an issue)</b></li> </ul>

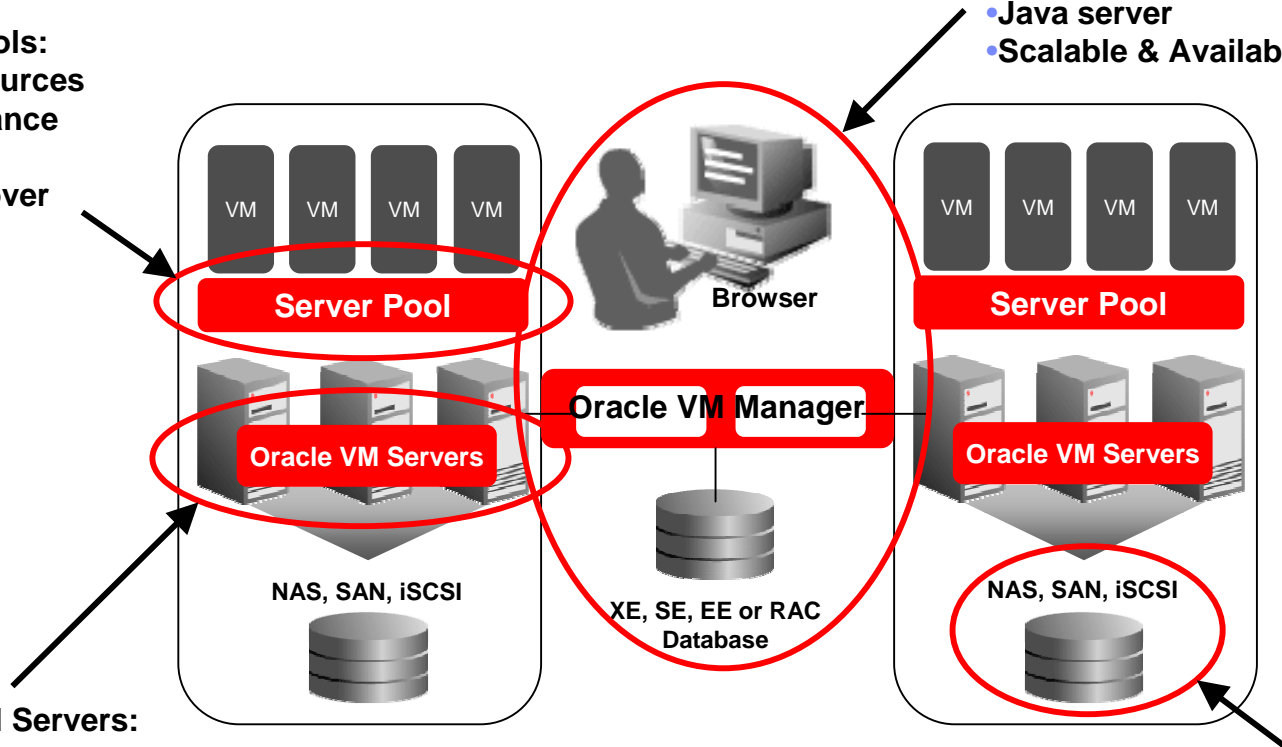
# Oracle VM Architecture

- Server Pools:**
- Pool resources
  - Load balance
  - Migrate
  - Auto failover

- Oracle VM Manager:**
- Browser-based
  - Java server
  - Scalable & Available

- Oracle VM Servers:**
- Host guest VMs
  - Enterprise Linux
  - Windows

- Shared Storage Pool Options:**
- NAS/NFS
  - SAN
  - iSCSI



# Oracle VM Manager

- Browser-based management solution
- Included with Oracle VM
- Full VM lifecycle management:
  - Create
  - Configure
  - Clone
  - Share
  - Boot
  - Migrate

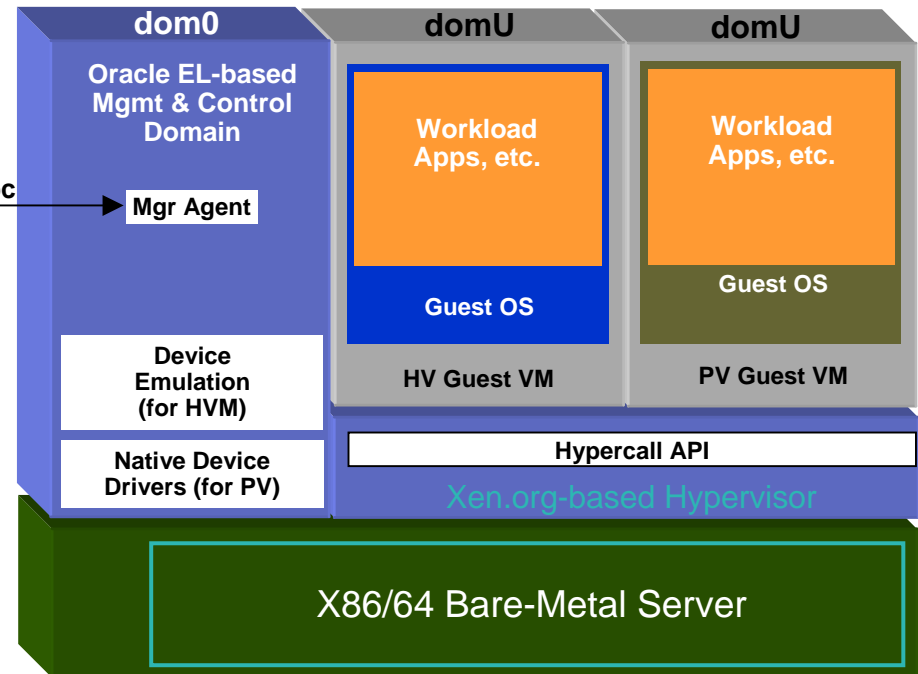
The screenshot displays the Oracle VM Manager web interface. At the top, there is a navigation bar with the Oracle VM Manager logo and links for Home, Profile, Logout, and Help. Below this is a secondary navigation menu with tabs for Virtual Machines, Resources, Servers, Server Pools, and Administration. The current page is titled 'Servers' and shows a search area with fields for Server Pool Name, Server Name, and Server Host/IP, along with a Status dropdown menu. A 'Search' button is located below these fields. A tip indicates that search criteria are case insensitive and that '%' can be used as a wildcard. Below the search area is a table of servers with columns for Select, Server Host/IP, Server Name, Server Type, Status, Server Location, and Server Pool Name. The table contains two rows of server data. At the bottom of the interface, there is a footer with navigation links and a copyright notice for Oracle VM Manager 2.1.

Select	Server Host/IP	Server Name	Server Type	Status	Server Location	Server Pool Name
<input checked="" type="radio"/>	<a href="#">ovs-z.cn.oracle.com</a>	ovs-z.cn.oracle.com	Virtual Machine Server	Active	OARDC	<a href="#">OARDC Server Group</a>
<input type="radio"/>	<a href="#">ovs-y.cn.oracle.com</a>	ovs-y.cn.oracle.com	Server Pool Master, Utility Server, Virtual Machine Server	Active	OARDC	<a href="#">OARDC Server Group</a>

# Anatomy of an Oracle VM Server

- “Dom0” is a privileged management and control domain typically containing...
  - A thin control kernel based on Oracle Enterprise Linux
  - Open / native Linux device drivers
  - Oracle VM Manager agent
  - Device emulation code to support non-PV guests (e.g. Windows)
- **Dom0 should not contain ‘user’ applications as a best practice**
  - Minimize performance & security risks
  - Minimize code size
- “DomU” is an unprivileged/user domain that is a guest VM on the server
  - Run any normal server workload
  - One domU is not aware of another

## Key Concepts: Dom0 & DomU



# VM storage using local disks

- Direct attached disks
  - **SATA, IDE, SAS, SCSI**
- Oracle VM default install uses OCFS2
  - **Locally mounted**
  - **virtual disk images in files**
  - **Robust, simple, high performance**
- Raw block devices
  - **/dev/hdb1, /dev/sdb1, /dev/cciss/c0d1p1**
- LVM is not recommended
  - **Does not support I/O barriers**

# Shared VM storage using SAN

- Storage array attached via SAN fabric
  - **accessible to multiple VM Servers**
- OCFS2 in shared LUN
  - **virtual disk image stored in file**
- Raw SAN LUN
  - **Virtual disk image stored in /dev/sdb1, /dev/mapper/mpath1**
- LVM is not supported
  - **Not cluster aware**

# Oracle VM Server

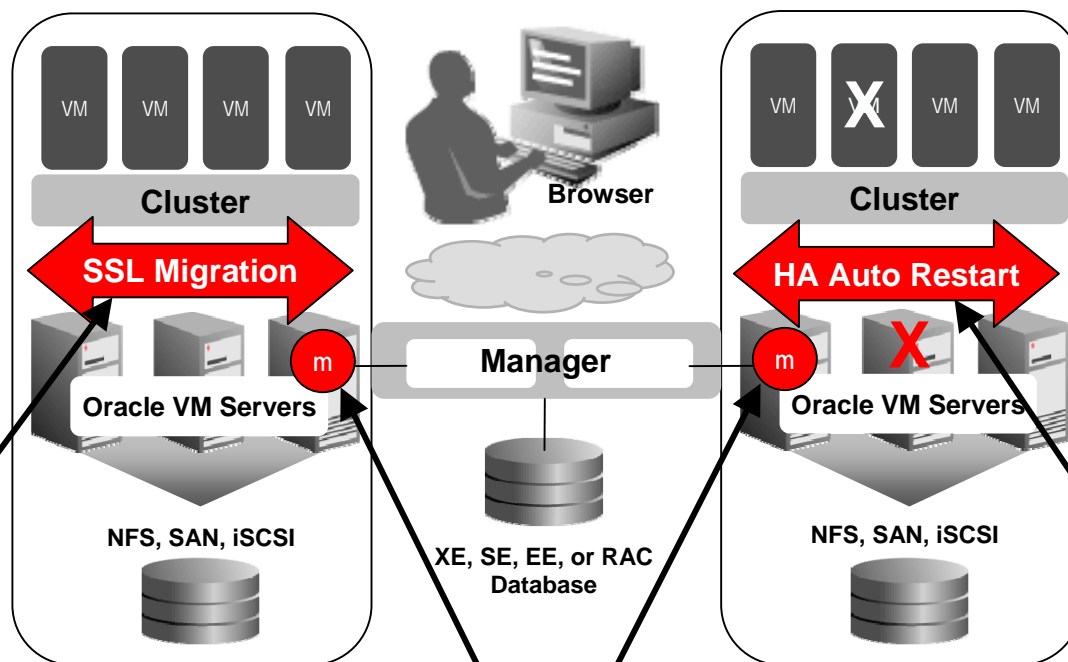
## Network/NAS: Redundant NICs and Switches

- Install redundant NICs for each network on each server
- Ideally, configure redundant switches
- Configure NIC Bonding
- Active/Passive for switch redundancy
- See: Support Note 738269.1
- If configured in Domain-0, all VMs take advantage of the redundant network configuration, or configure in each VM if required

# Guest VM High-Availability

- Automatic restart of failed VMs across the pool
  - **Server failure (all VMs restarted)**
  - **Individual VM failure**
- Reliable restart based on proven Oracle Clusterware technology
  - **Sophisticated heartbeat and lock management**
  - **Reliable failure detection and corruption prevention**
- Maximize up-time without complexity of traditional HA clustering
  - **Cost-effective solution**
  - **No guest VM agents or modifications required**

# Oracle VM Guest High Availability



## Planned Events:

- E.g. maintenance or upgrades
- Secure Live Migration
- Zero interruption

Pool Masters assure  
Secure Migration or  
HA restarts complete  
in the event of a  
Manager outage

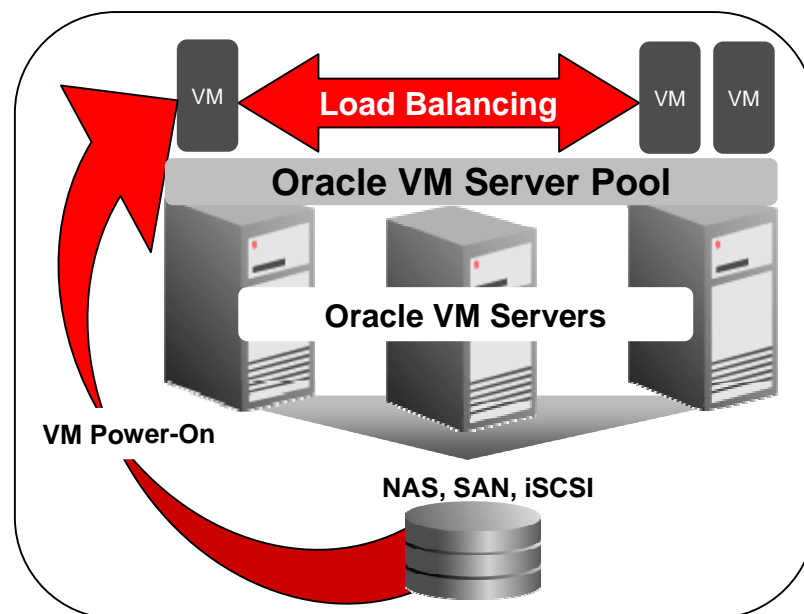
## Unplanned Events

- E.g. Server or VM failure
- HA Auto-restart in pool
- No manual intervention

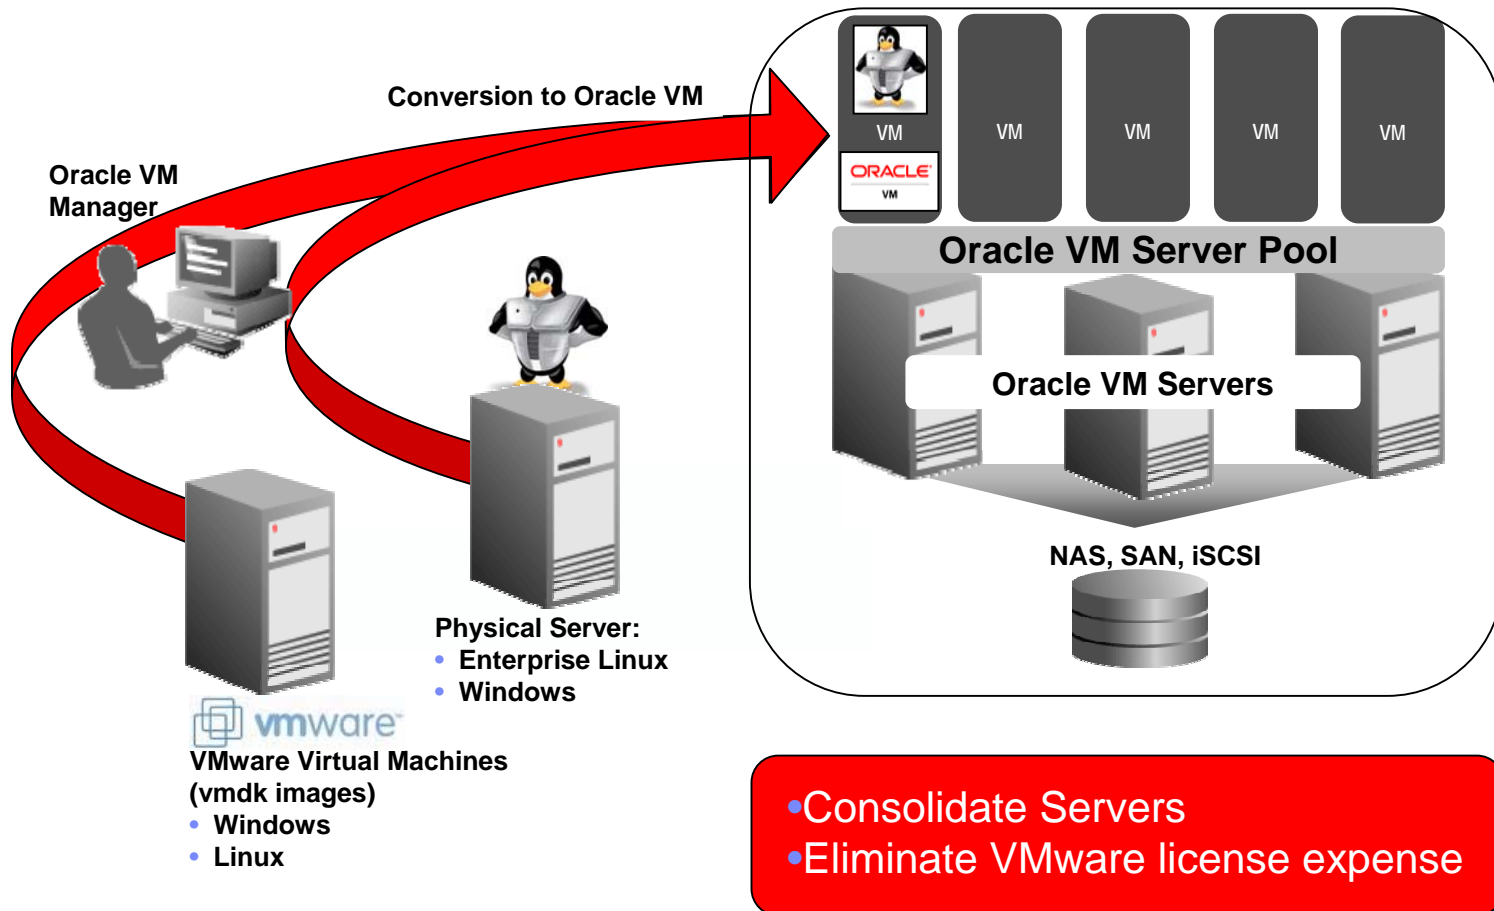
# Pool Load Balancing

## Reliable Start-Up & Resource Efficiency

- Automatic load balancing across all servers in (sub) pool
- User defined Preferred Server sub pools per guest VM
- Algorithm selects server with the most available compute resources (CPU and Mem)
- Efficient resource usage
- Assures that “down” server will not block VM power-on



# Physical-to-Virtual / Virtual-to-Virtual Machine Conversion



# Oracle VM Templates

## Rapid Application Deployment


ORACLE E-Delivery

Download from Oracle

- Pre-built, pre-configured VM
- Complete app, middleware, DB installation
- Complete Siebel CRM, Database 11g, Enterprise Manager...



Import to Oracle VM Manager



Start-Up in Oracle VM Pool

Save days or weeks in installation and configuration time



Customize & Save as Golden Images

Oracle VM Server Pool

Oracle VM Servers

NAS, SAN, iSCSI



# Performance Testing

- **Extensive real-world testing**
  - LMbench: microbenchmarking
    - Context switching, Null calls, TCP Selects, etc
  - Swingbench: DB workload benchmarking
    - Typical OLTP workloads
    - Varying SGA size, # of users, # of vCPUs
  - Paravirtualized domains with Enterprise Linux 4, 64bit
- **Results:**
  - On average **three times less overhead** compared server virtualization products from other vendors.
  - Virtual SMP scalability of Oracle VM is at 90+%
- **Small workloads are at same performance as real hardware**

# What Differentiates Oracle VM?

**Enterprise-quality support**

**24x7, global, large-scale support organization**

---

**Complete. Integrated Solution**

**Full software stack as opposed to a point product**

---

**Real-world Testing and Best Practices**

**Oracle VM Templates; Oracle Validated Configurations available at no additional charge; Faster and More Efficient**

---

**Lower total cost; Affordable**

**No license fees; All options are included; lower support cost; consolidate support contracts**

---

**Ease of doing business**

**One vendor for the complete solution, no finger pointing**