


Speaker: Dan Magenheimer



**ORACLE®**

**VM**

ONE COMPLETE SOFTWARE STACK.  
ONE SOURCE FOR SERVER VIRTUALIZATION AND LINUX.  
ONE CALL FOR SUPPORT.

**Topic : The Xen of Virtualization on Linux**

**13<sup>th</sup> May 2008 @**

**CLUE**

**Colorado Linux Users  
and Enthusiasts**



**What is virtualization?**

**A brief history of computing...**

# Three immutable(?) laws of computing

1. Wirth's law: Software gets slower faster than hardware gets faster
2. ...
3. ..

# Three immutable(?) laws of computing

1. Wirth's law: Software gets slower faster than hardware gets faster: "***What Grove giveth, Gates taketh away***"
2. ...
3. ...

# Three immutable(?) laws of computing


1. ...
2. Everything in computing can be solved by ***adding another layer of indirection*** (/ abstraction)
3. ...

# Three immutable(?) laws of computing



1. ...
2. ...
3. There's nothing new in computing that IBM didn't already invent 30 (40, 50) years ago

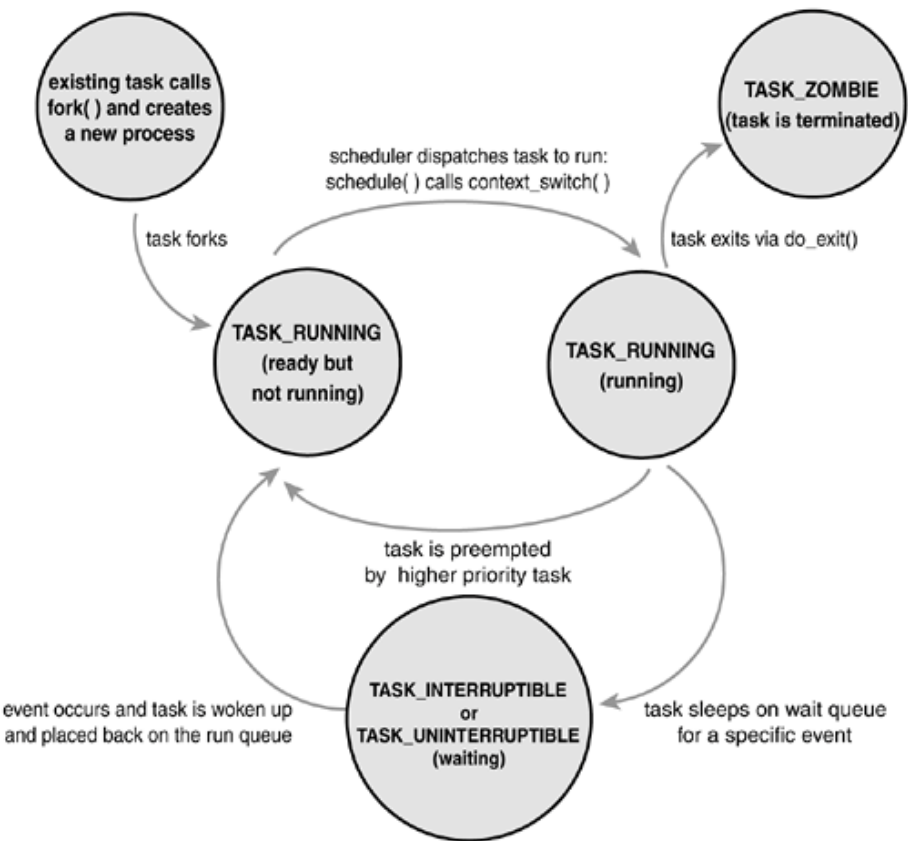

# Three immutable laws of computing -- or are they (immutable?)

1. Software gets slower faster than hardware gets faster, aka “*What Grove giveth, Gates taketh away*”
2. Everything in computing can be solved by adding *another layer of indirection* (/ abstraction)
3. There’s nothing new in computing that IBM didn’t already invent 30 (40, 50) years ago









```

QF d:

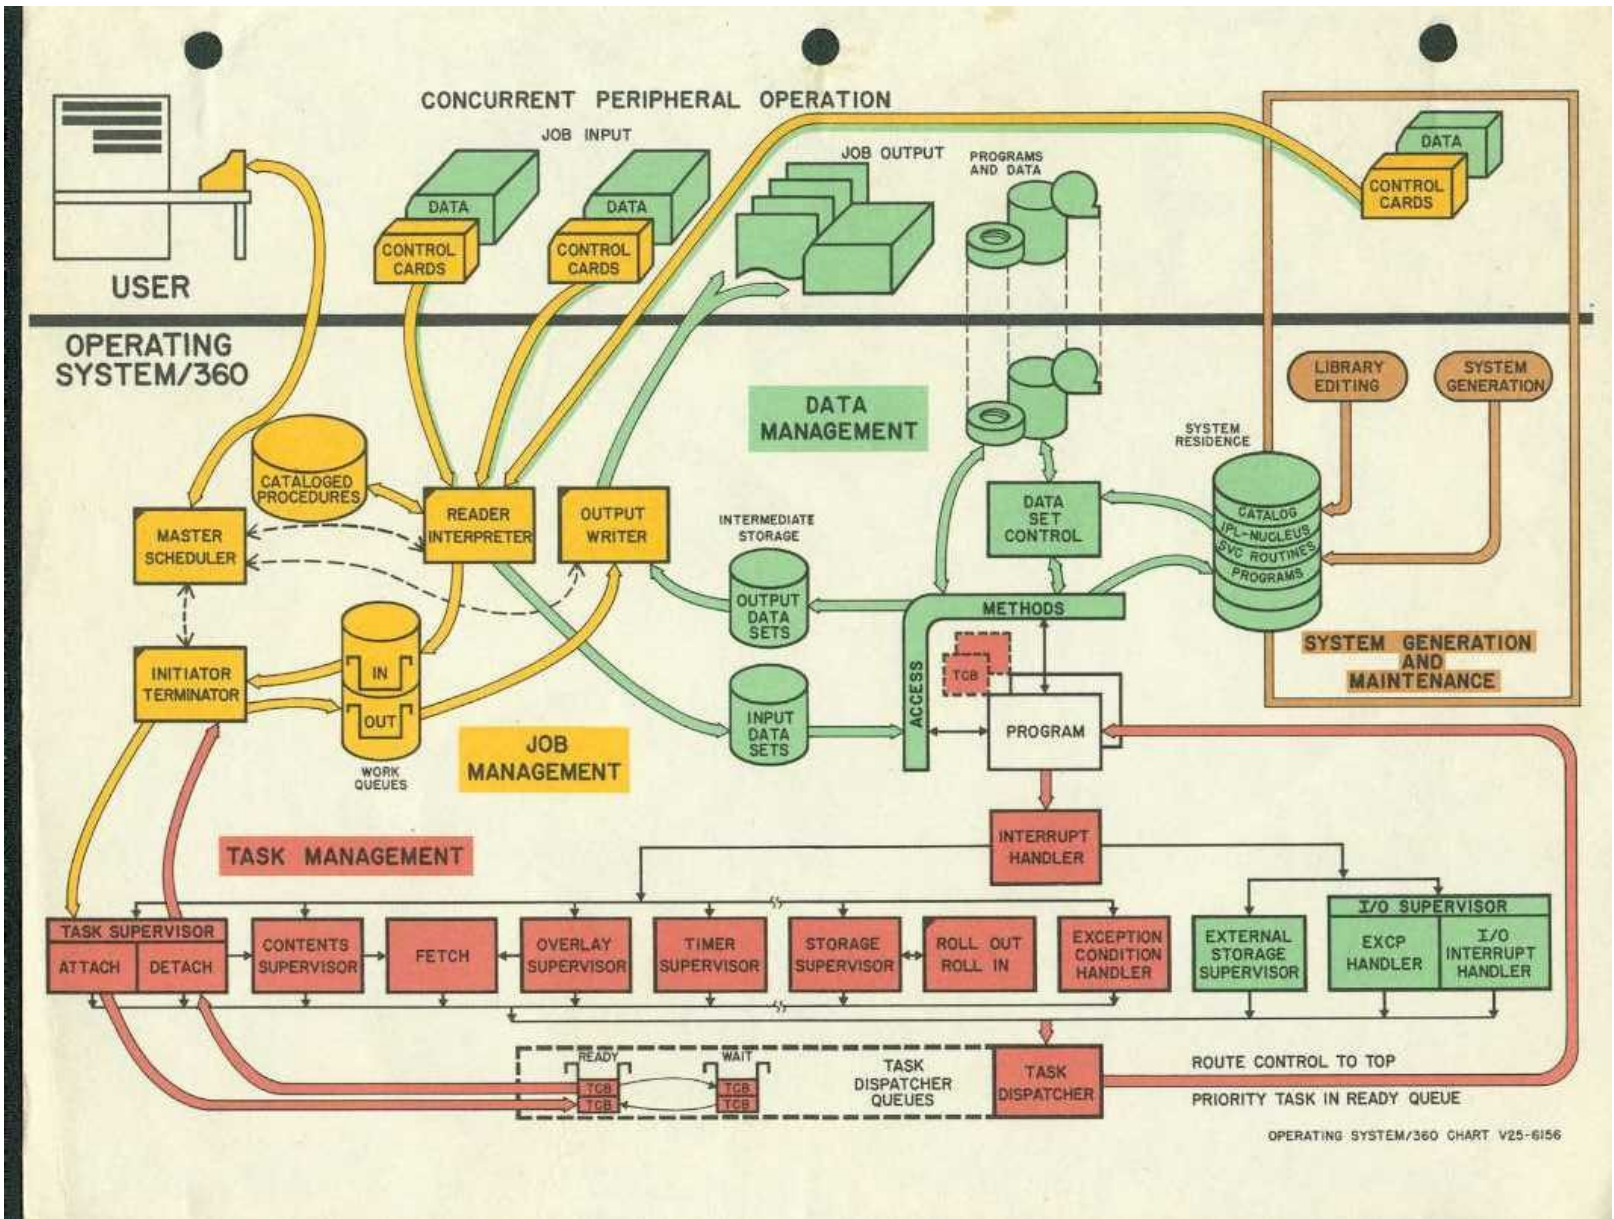
COMMUNICATIONS DIVISION DISPATCH QUERY SCREEN

>QR NAM/
>QTOM NAM/
>QWAR NAM/
>QCHIS NAM/
>QE NAM/
>QPT NAM/
>QPT ENO/

>QA ENO/
>DQ.NAM/
>DQ.SSN/
>DQ.OLN/
>QV.LIC/
>RQ.LIC/
>QVM VIN/
>QVM LIC/

.DOB/
.OLS/
.LIS/
.LIS/
  
```





User Program 1
Operating System 1
Virtual 370

User Program n
Operating System n
Virtual 370

VM370
Hardware

## VIRTUAL MACHINE ARCHITECTURE (VM370)











```
Microsoft(R) Windows DOS  
(C)Copyright Microsoft Corp 1990-2001.
```

```
C:\>mem
```

```
655360 bytes total conventional memory  
655360 bytes available to MS-DOS  
578352 largest executable program size
```

```
4194304 bytes total EMS memory  
4194304 bytes free EMS memory
```

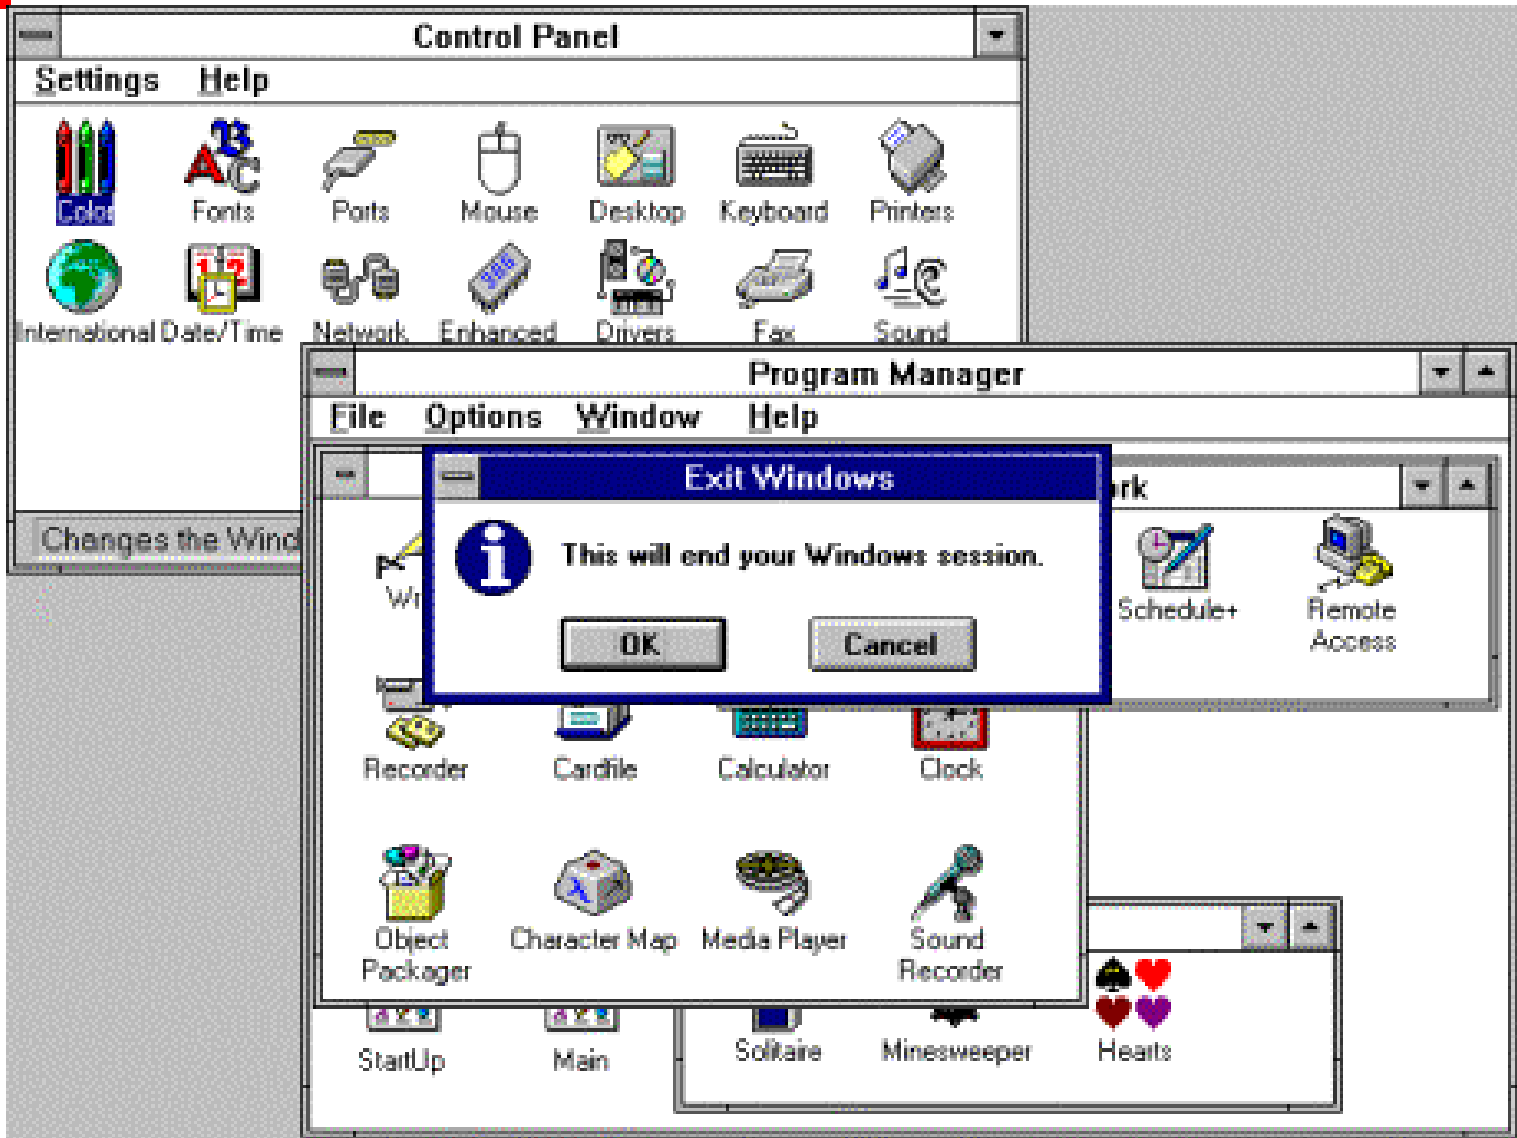
```
19922944 bytes total contiguous extended memory  
0 bytes available contiguous extended memory  
15580160 bytes available XMS memory  
MS-DOS resident in High Memory Area
```

```
C:\>
```

Would you have invested?



Microsoft Corporation, 1978




## Windows

A fatal exception 0E has occurred at 0028:C0011E36 in UXD UMM(01) + 00010E36. The current application will be terminated.

- \* Press any key to terminate the current application.
- \* Press CTRL+ALT+DEL again to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue \_



## Windows

A fatal exception 0E has occurred at 0028:C0011E36 in UXD UMM(01) + 00010E36. The current application will be terminated.

- \* Press any key to terminate the current application.
- \* Press CTRL+ALT+DEL again to restart your computer. You will lose any unsaved information in all applications.

Press any key to continue \_



Linux

BSD

AIX



SCO

1975

1980


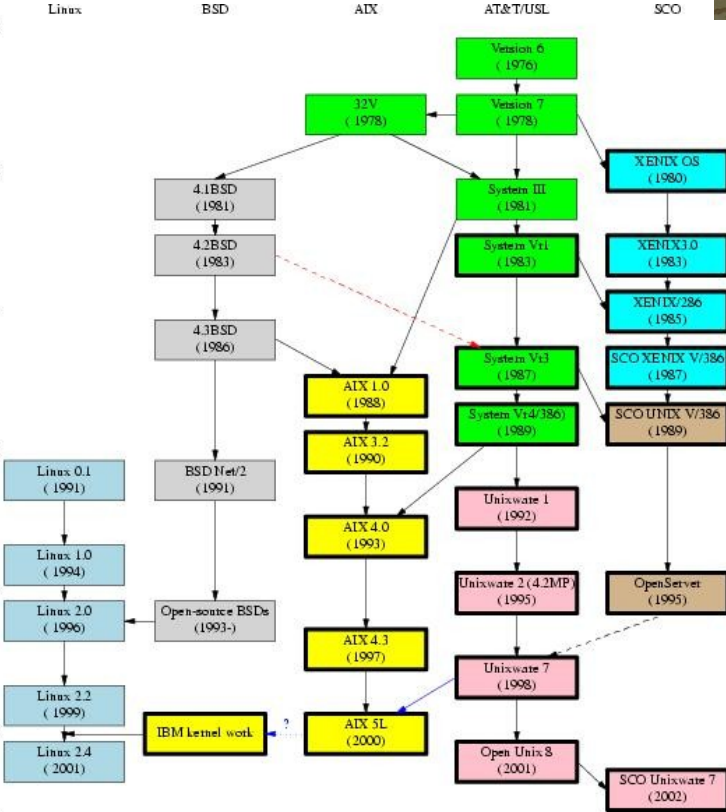
1985


1990

1995


2000

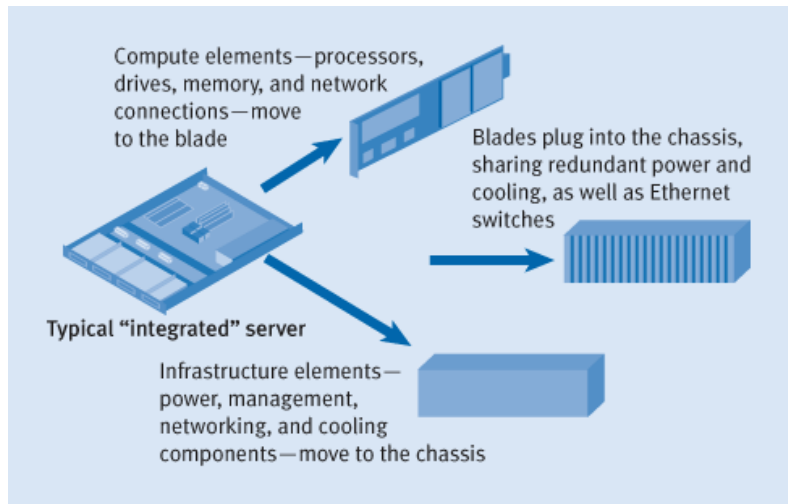
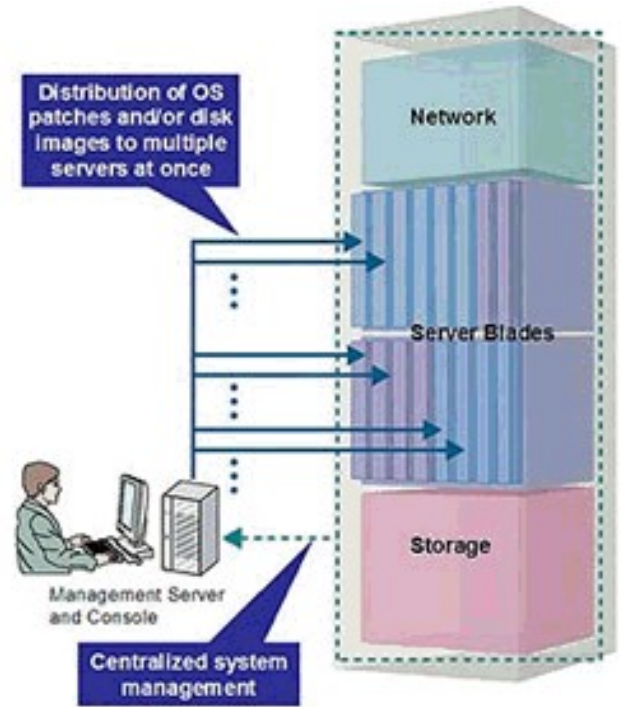
2003

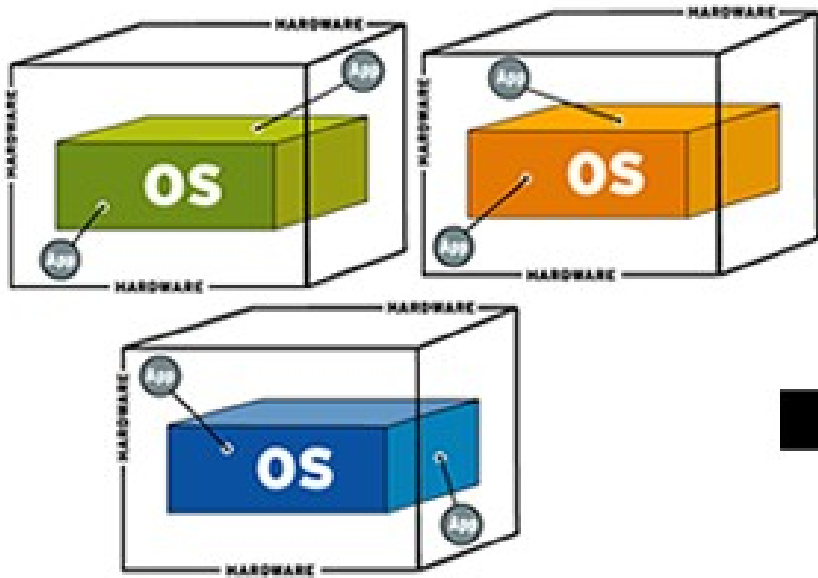














Before



After



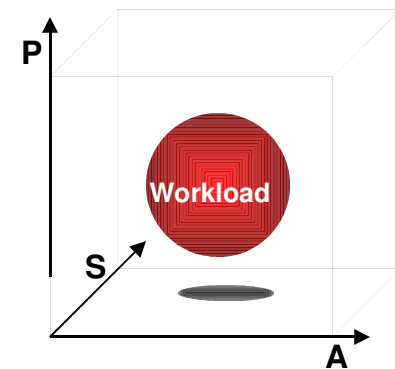


# Overview: Virtualization Technology



# Defining Virtualization


- Abstraction of underlying technology
  - Decouple workloads from resource boundaries
- Benefit:
  - Increased agility in matching resources to the changing needs of your workload:
    - Performance
    - Availability
    - Security
    - Stability



# Server Virtualization Technologies

## Host OS-based, e.g.

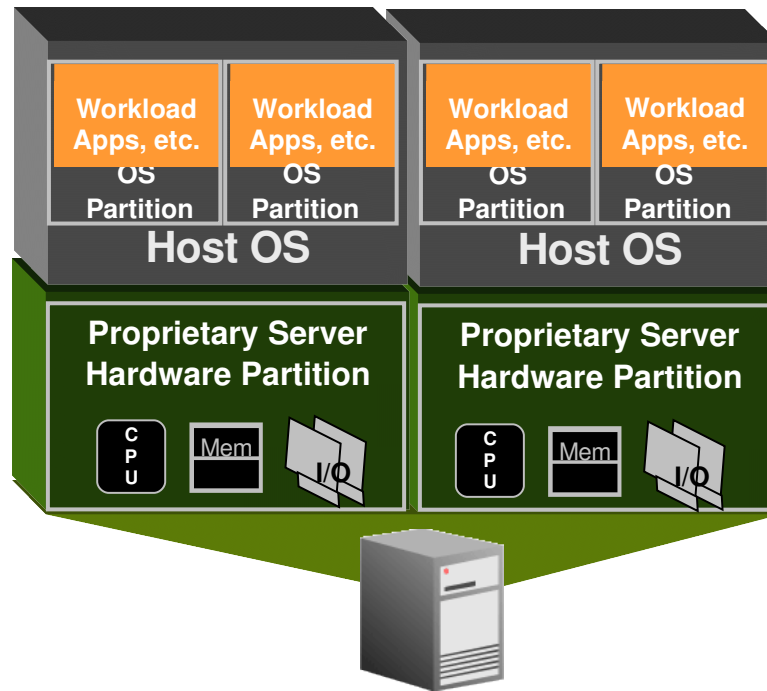
- VMware Workstation
- Microsoft Virtual Server



- Primarily desktop
- Easy to use
- Very slow (2 OSeS)

## Hardware Partitioning, e.g.

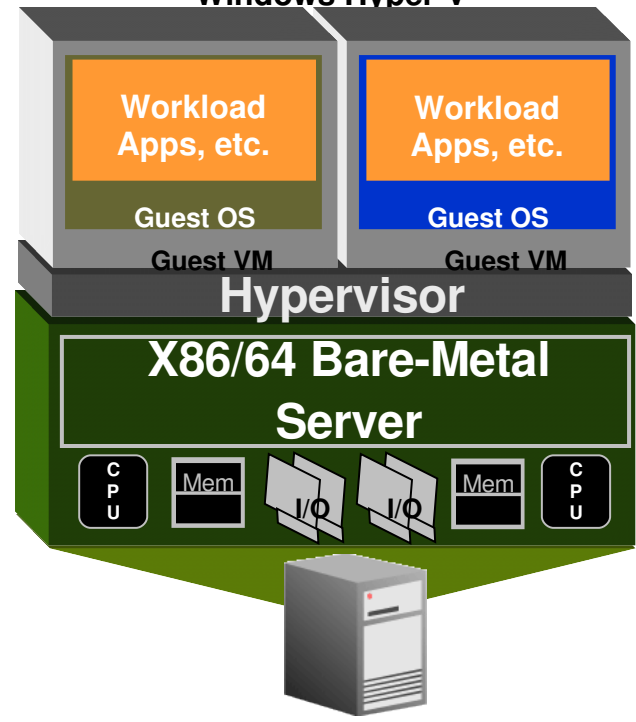
- Sun SPARC Domains
- IBM PPARs
- HP nPARs



- Excellent isolation
- Expensive, proprietary hw
- Coarse grain resources
- Mix OSeS / versions

## OS Partitioning, e.g.

- Solaris Containers
- AIX dLPARs
- HP vPARs



- Only moderate isolation
- Potentially good scalability
- Fine-grained resources
- Cannot mix OS/patch levels

## Hypervisor-based, e.g.

- **Oracle VM**
- VMware ESX Server
- Citrix XenServer
- Windows Hyper-V

- Excellent isolation
- Affordable, multi-source hw
- Fine-grained resources
- Mix OSeS / versions



# “Hypervisor” Vs. “Virtualization Server”?

- Terms commonly used interchangeably but actually technically distinct:
  - A **hypervisor** is a very low-level, very efficient platform code layer that allows multiple operating systems to run on a host computer at the same time
    - Fundamental abstraction of physical compute resources
      - I/O resources may- (VMware) or may not (Xen) be handled by the hypervisor itself
    - Allows resource (re)allocation between VMs without disruption
  - A **virtualization server** includes a hypervisor function but also includes additional functionality, e.g.
    - I/O resource/traffic management (Xen)
    - Server configuration, access, and security management, etc.

# Hardware Support for Virtualization (HVM)

## Where does it fit-in?

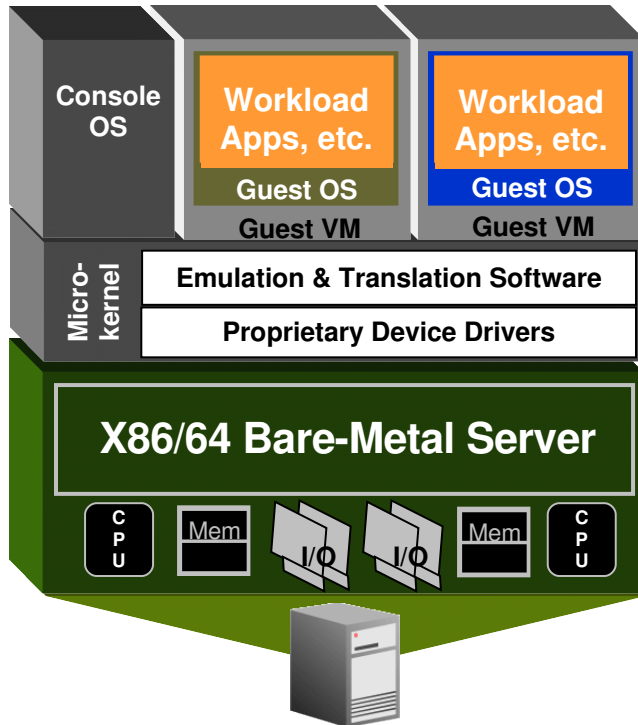
- Goal: efficiently insure VM OSES are not “hurting” each other
  - E.g. Trying to unsafely modify common resources/state, etc.

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

# Anatomy of a Virtualization Server

## Emulation-based, e.g.

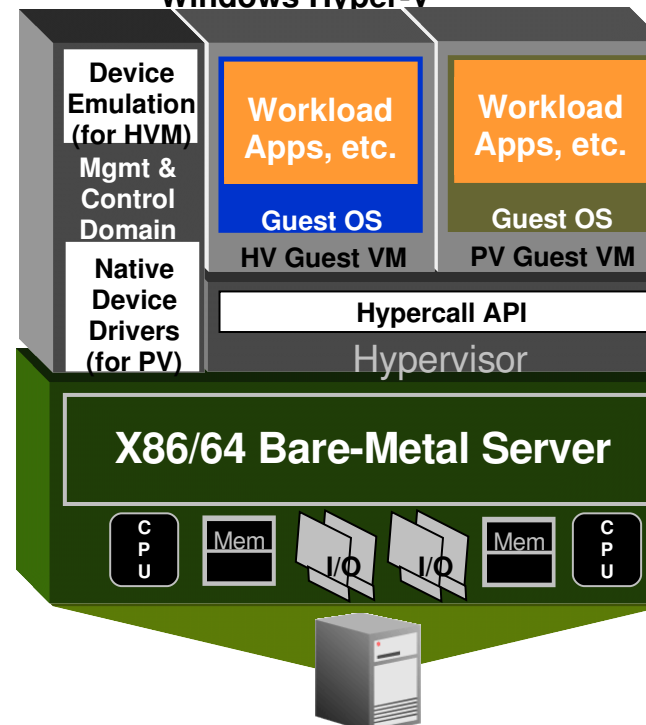
- VMware ESX Server



- Machine images run unmodified: broad compatibility
- Does not leverage or require HV hardware
- Poor I/O scalability due to emulation architecture
- Not open: dependent on virt. vendor for drivers

## Paravirtualization (PV)-based, e.g.

- Oracle VM
- Citrix XenServer
- Windows Hyper-V

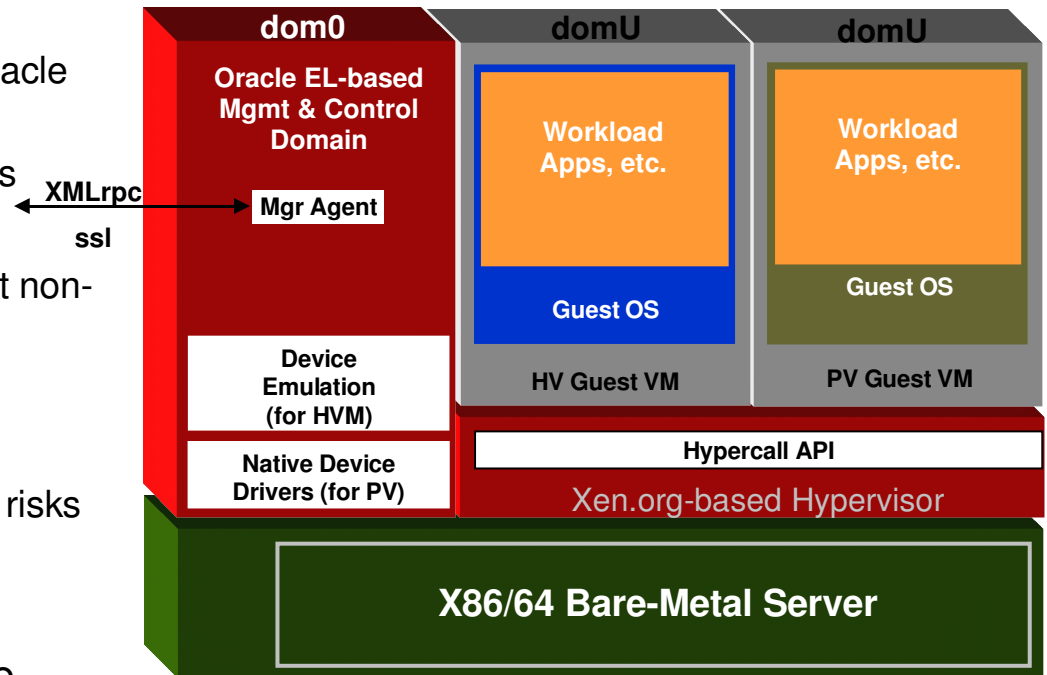


- Excellent scalability, esp. I/O with PV guests
- Requires PV OS kernel for best perf.
- Requires HV hardware for unmodified images
- Open: wide device support: uses native (e.g. Linux) device drivers

# Anatomy of an Oracle VM Server

## Key Concepts: Dom0 & DomU

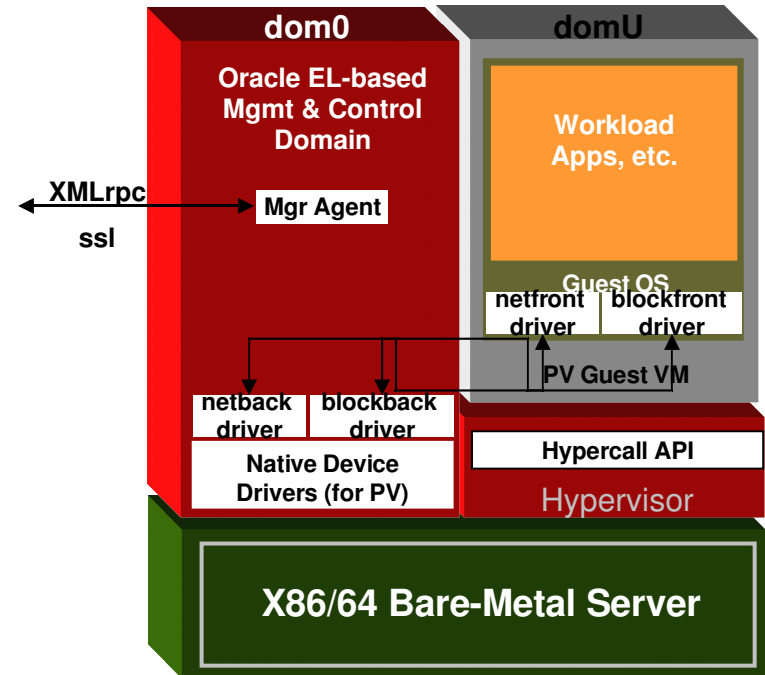
- “**Dom0**” is a privileged management and control domain typically containing...
  - A thin control kernel based on Oracle Enterprise Linux
  - Open / native Linux device drivers
  - Oracle VM Manager agent
  - Device emulation code to support non-PV guests (e.g. Windows)
- Dom0 should not contain ‘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



# Anatomy of an Oracle VM Server

## Key Concepts: Driver Architecture: PV

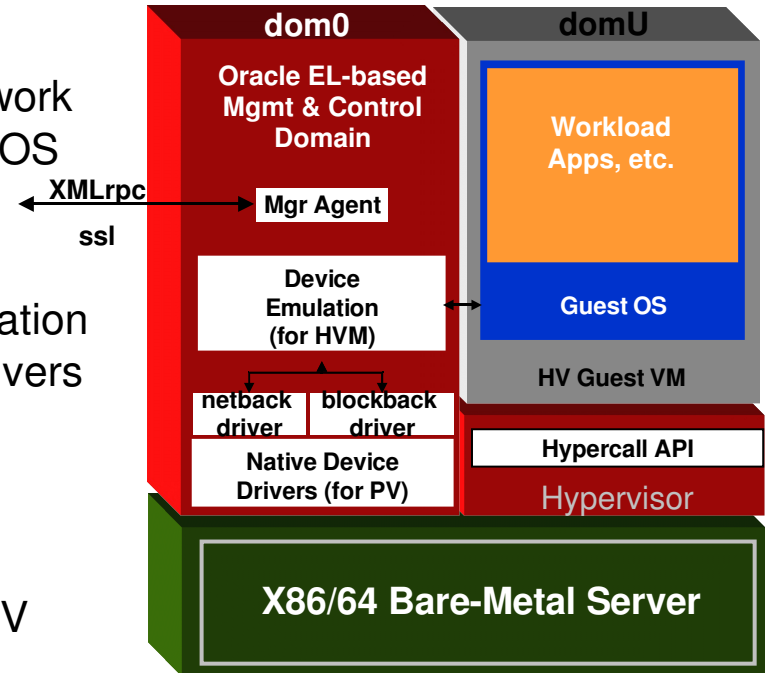
- For PV guests, driver abstraction:
  - Replaces hardware-specific drivers
    - One network driver
    - One block driver
  - Very stable / rarely changes
    - Excellent guest stability
  - **Front-end drivers (net & block)**
    - Inside the VM / domU OS
  - **Back-end drivers (net & block)**
    - In dom0 / shared
  - **Open, native hardware vendor drivers**
    - Uses open Linux drivers
- Virtualization server manages high-performance communication front <-> back
  - Can leverage dom0 kernel security features, e.g. packet sniffing, firewalling, & rate control, etc.



# Anatomy of an Oracle VM Server

## Key Concepts: Driver Architecture: HVM

- HVM guests driver choices:
  - **Unmodified native driver(s)**
    - OS typically installs basic native network and block drivers that come with the OS (not with the virtualization server)
    - Device support provided via device emulation / translation in the virtualization server on top of net- & block-back drivers in dom0
    - Slower than PV due to overhead
  - **PV drivers** (front- & back)
    - HVM (unmodified) kernels can use PV drivers
    - Leverages PV driver stack same as previously described
    - Excellent performance

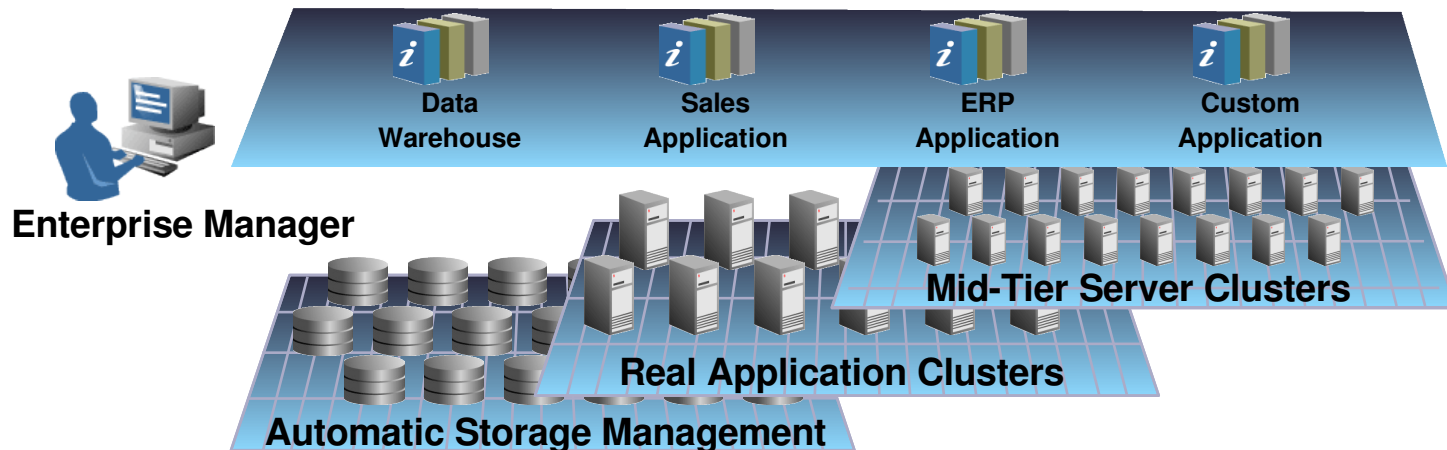


# Oracle's Virtualization Platform: Oracle VM




# Oracle Grid-Based Computing A History in Virtualization

- Grid-Based Storage
  - Oracle Automatic Storage Management
- Grid-Based Server Clusters
  - Oracle Database
  - Oracle Middleware
- Grid-Based Management
  - Oracle Enterprise Manager Grid Control





# Oracle VM



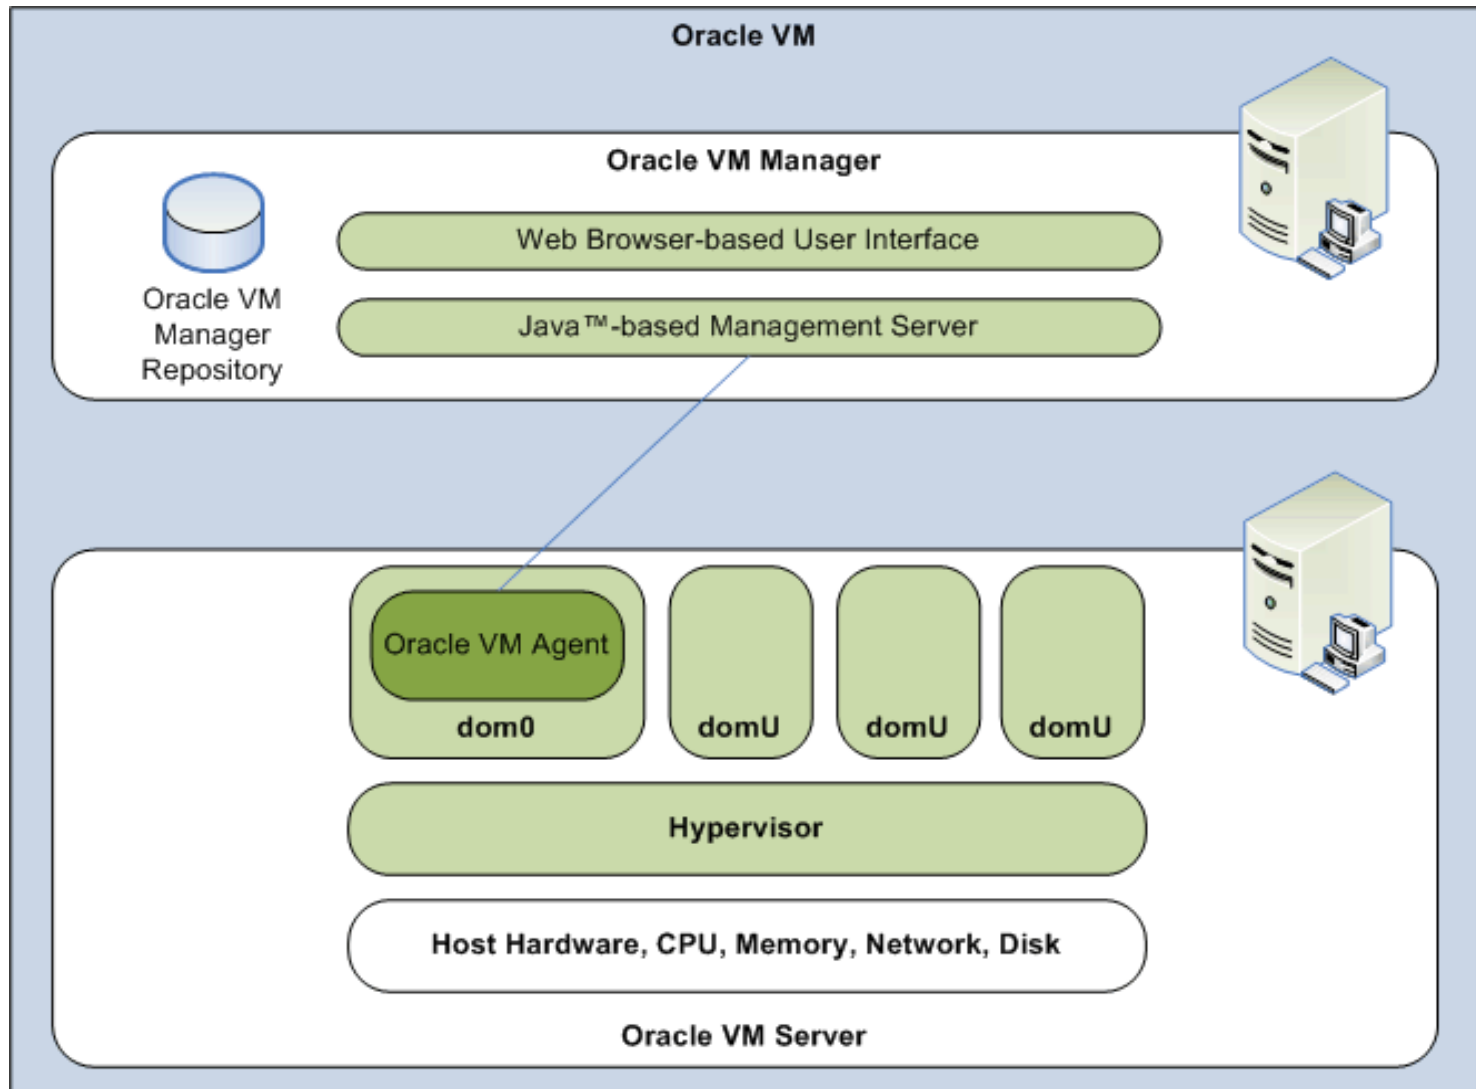
- Oracle tested and supported server virtualization
- Maximizes consolidation of Linux and Windows servers
- Saves on power, cooling and space
- Improves operational agility through the lifecycle

# Oracle VM

- Server virtualization software and support
  - Enterprise-class hypervisor
  - Free product download; low-cost support
- Runs both Linux and Windows Guests
  - Paravirtualized and hardware virtualized guests
    - HV on x86 hardware with HVM support
  - 64-bit and 32-bit guests
  - Up to 64-way SMP hardware
  - Up to 32 virtual processors per guest
  - Includes Live Migration at no additional cost
  - Integrated, browser-based management console
  - Free downloadable VM images
- Enterprise-quality support worldwide

**ORACLE**<sup>®</sup>  
VM

# Oracle VM



# Oracle VM Features

- Simplified installation
  - Single CD server installation
  - Installs and configures in a minute
- Faster stack deployment
  - Pre-configured Virtual Machine images of Oracle Database and Oracle Enterprise Linux
- Live VM migration included
  - Better uptime, agility, and HA
- Linux and Windows guest support
  - Oracle Enterprise Linux 4 and 5;
  - RHEL3, RHEL4 and RHEL5
  - On HV capable hardware: Windows 2000, Windows 2003, Windows Server 2003 (32- and 64bit) and Windows XP

# 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, the title "ORACLE VM Manager" is visible, along with navigation links for "Home", "Profile", "Logout", and "Help". Below the title, there are tabs for "Virtual Machines", "Resources", "Servers", "Server Pools", and "Administration". The "Servers" tab is currently selected. The page shows a search area with fields for "Server Pool Name", "Server Host/IP", "Server Name", and "Status" (set to "All"). A "Search" button is located below these fields. A tip below the search area states: "TIP Search criteria are case insensitive. Use '%' as a wildcard, for example prod%". Below the search area, there is a table of servers. The table has columns for "Select", "Server Host/IP", "Server Name", "Server Type", "Status", "Server Location", and "Server Pool Name". Two servers are listed: "ovs-z.cn.oracle.com" (Virtual Machine Server) and "ovs-y.cn.oracle.com" (Server Pool Master, Utility Server, Virtual Machine Server). Both are in an "Active" status and located in "OARDC". Below the table, there are "Refresh" and "Add Server" buttons. At the bottom of the page, there is a footer with the text "Copyright © 2007, Oracle. All rights reserved. Oracle VM Manager 2.1".

ORACLE VM Manager

Home Profile Logout Help

Virtual Machines Resources Servers Server Pools Administration

Logged in as admin

Servers

Refresh Add Server

Search

Server Pool Name:  Server Host/IP:

Server Name:  Status: All

TIP Search criteria are case insensitive. Use '%' as a wildcard, for example prod%

Servers

Select and

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>

Refresh Add Server

Virtual Machines Resources Servers Server Pools Administration

Copyright © 2007, Oracle. All rights reserved. Oracle VM Manager 2.1

# Oracle Product Certification with Oracle VM

- Oracle Database
- Oracle Application Server
- Oracle Enterprise Manager
- Oracle Berkeley DB
- Oracle TimesTen
- Oracle E-Business Suite
- Oracle PeopleSoft
- Oracle Siebel
- Oracle Hyperion
- More information on Metalink Note 464754.1



## Oracle VM 2.1

### Validated Configuration Details

Publication Date	26-DEC-07
Version	1
Server Platform	HP ServerBlade BL480c
Storage Model	SmartArray P400 StorageBlade 40c,StorageBlade, SB600c
Guest VMs	Oracle Database Single Instance on OEL 4 Update 5 x86_64 , OEL 5 x86_64

**Oracle Validated Configurations  
For  
Oracle VM**

### Guest VM Configurations

Oracle recommends Paravirtualized guests

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 4 Update 5 x86\\_64 ,Hardware Virtualized with Oracle 10.2.0.3](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 5 x86\\_64 ,Hardware Virtualized with Oracle 11.1.0.6](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 4 Update 5 x86\\_64 ,Hardware Virtualized with Oracle 11.1.0.6](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 5 x86\\_64 ,Hardware Virtualized with Oracle 10.2.0.3](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 5 x86\\_64 ,Paravirtualized with Oracle 10.2.0.3](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 5 x86\\_64 ,Paravirtualized with Oracle 11.1.0.6](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 4 Update 5 x86\\_64 ,Paravirtualized with Oracle 11.1.0.6](#)

[Oracle VM Server 2.1 , Guest VM - Oracle Enterprise Linux 4 Update 5 x86\\_64 ,Paravirtualized with Oracle 10.2.0.3](#)

### Server and Storage Platform Details

Server Model	2 X HP ServerBlade BL480c
Processors	2 X Quad Xeon
Hardware Virtualization Support	Yes

**Oracle & Alliance Partner Collaboration:  
Making Linux & Virtualization Easy to Deploy**

# Partner Support





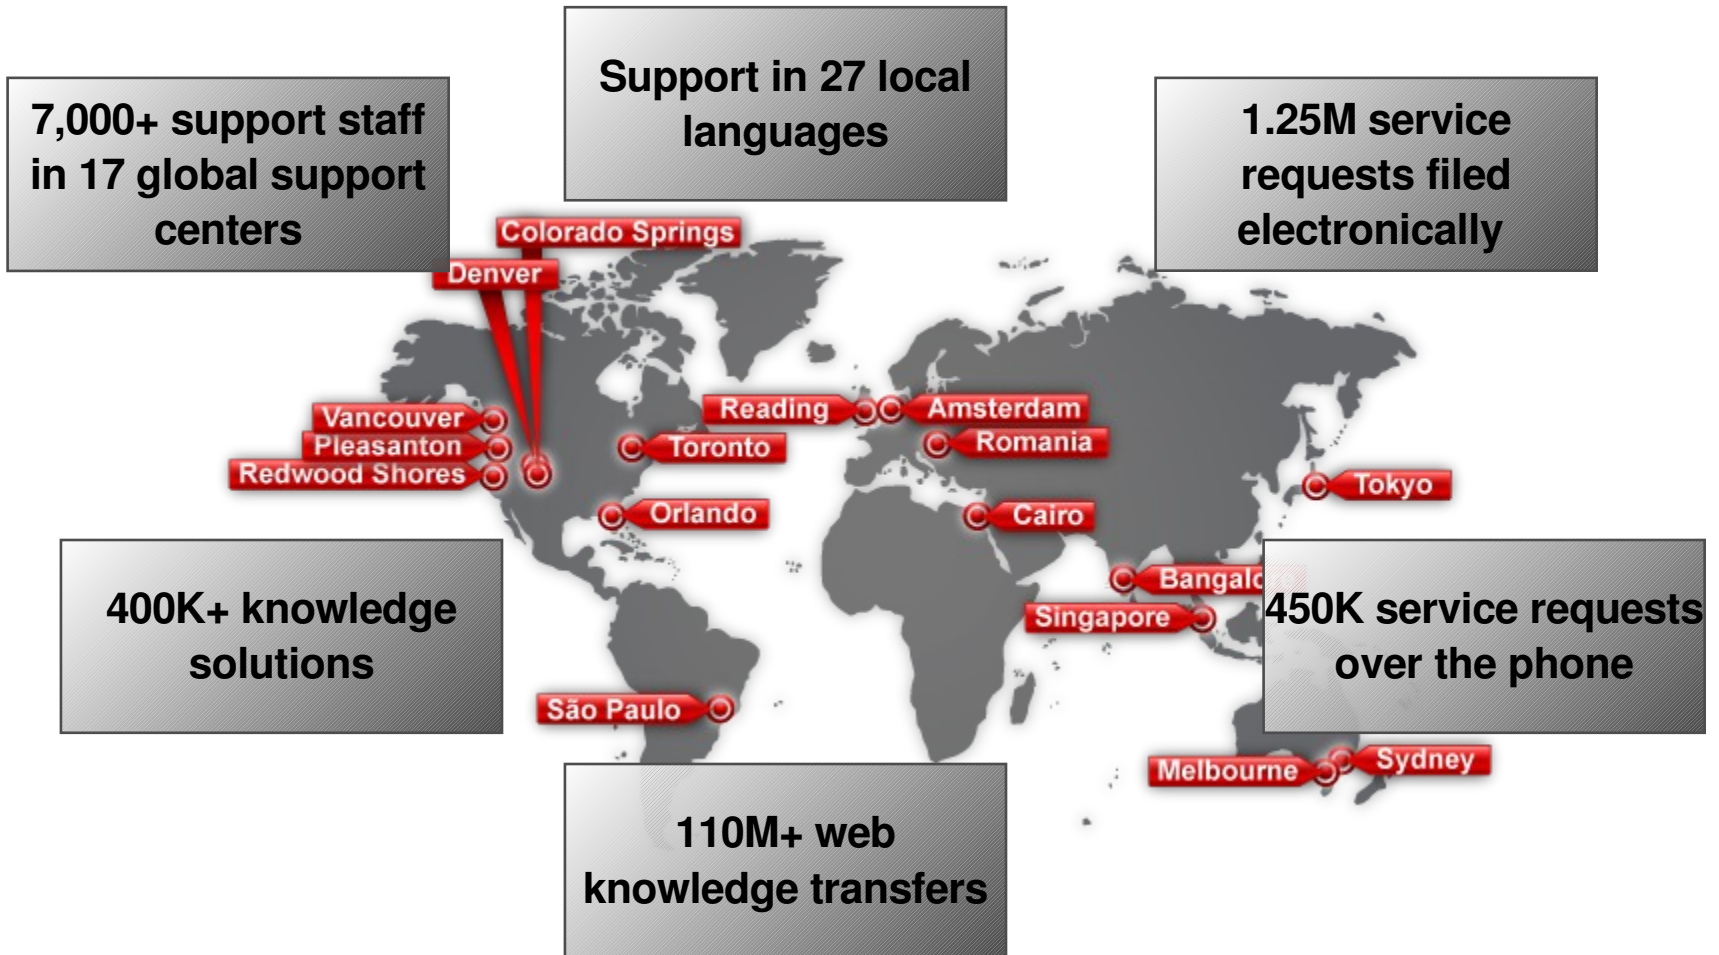
# Oracle's Contributions

- Enhanced and optimized Xen technology
  - I/O overhead
  - Memory overhead
  - Process Scheduling
- Community involvement
  - Dedicated Xen development team
  - Code and bug fix contributions to Xen community
  - Members of Xen community at Oracle
  - Member of XenSummit committee
- Significant testing
  - Real-world testing with Oracle On Demand workloads
  - Testing with Oracle Validated configuration workloads

# Performance Testing

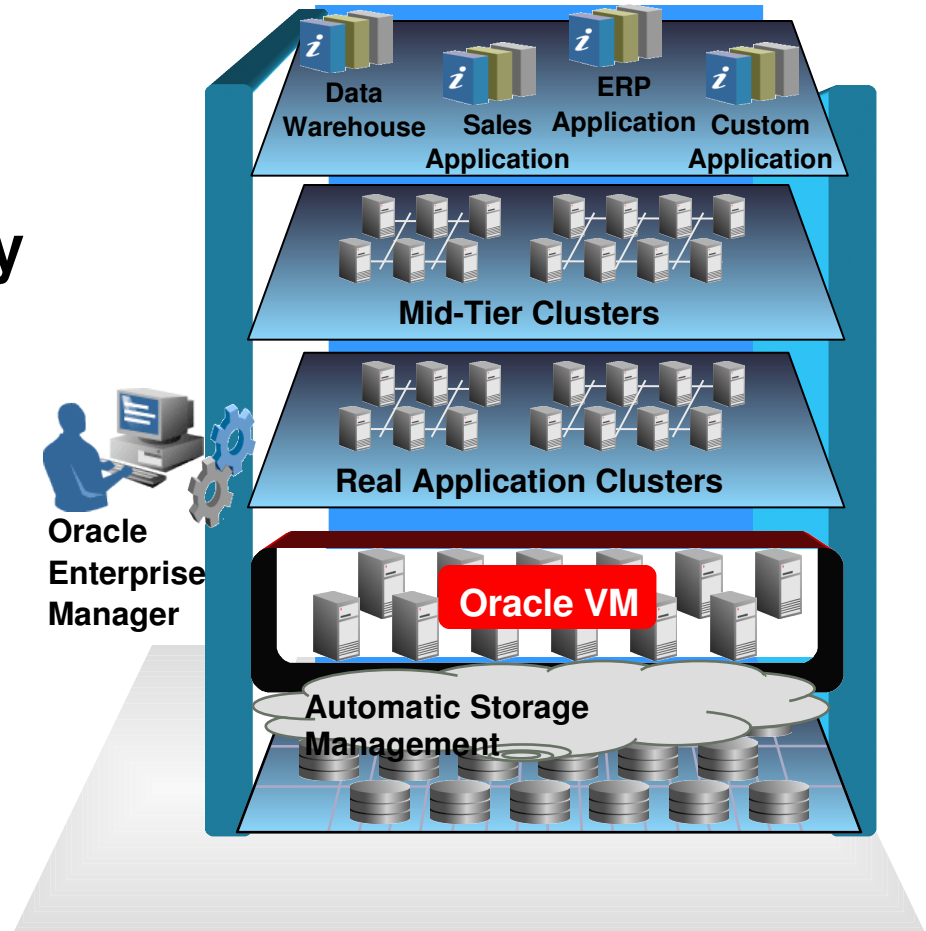
- Extensive real-world testing
  - LMbench: microbenchmarking
    - Context switching, Null calls, TCP Selects, etc
  - Swingbench: DB workload benchmarking
    - Typical OLTP workloads
    - Varying SGA size, # of users, # of vCPUs
  - Paravirtualized domains with Enterprise Linux 4, 64bit
- Results:
  - 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

# One Support Call for the Complete Stack



# Summary

- **Oracle Uniquely Combines**
  - Proven Grid capabilities
  - Server virtualization
- **Full Virtualization Strategy**
  - Applications
  - Data
  - Servers
  - Storage
- **Real-World Benefits**
  - Maximized consolidation
  - Optimized agility
  - High availability
  - Affordable enterprise-class support



# Resources

- **Virtualization Center on OTN**  
[otn.oracle.com/goto/virtualization](http://otn.oracle.com/goto/virtualization)  
- wiki, downloads, discussion forum, howtos
- **Oracle VM Home Page**  
[oracle.com/virtualization](http://oracle.com/virtualization)
- **Free Download: Oracle VM**  
[edelivery.oracle.com/oraclevm](http://edelivery.oracle.com/oraclevm)

