

# Oracle VM Tips and Best Practices

To become a successful Oracle VM professional

---

Presented by: Francisco Munoz Alvarez



June 1, 2015

Pythian  
love your data®  
Pythian  
love your data®

# Francisco Munoz Alvarez

Oracle ACE Director

8/9/10g/11g OCP, RAC OCE, AS OCA, E-Business OCP,  
SQL/PLSQL OCA, Oracle 7 OCM

Oracle 7, 11GR2, 12cR1 and OVM 3.1 and 3.2 and 3.3 Beta Tester

IOUC LA Spokesperson, President of LAOUC, APACOUC, CLOUG and NZOUG

ITIL Certified

2010 Oracle ACE Director of the year by Oracle Magazine

Blog: [oraclenz.com](http://oraclenz.com) - Email: [alvarez@pythian.com](mailto:alvarez@pythian.com)

Twitter : [fcomunoz](https://twitter.com/fcomunoz)

Pythian

Vice President Asia Pacific

[www.pythian.com](http://www.pythian.com)





# Francisco Munoz Alvarez

ORACLE ACE DIRECTOR OF THE YEAR

Oracle ACE director hones skills while helping others.

When it comes to learning something new, Francisco Munoz Alvarez doesn't look for a teacher; he looks for someone who has a problem.

"One of the best ways to learn is by helping others," says Alvarez. "When someone has a problem on an OTN [Oracle Technology Network] forum, I enjoy trying to assist them with it. I've learned a lot by trying to help other Oracle users to solve their problems."

Over the years, he has learned enough to become an Oracle ACE, an Oracle ACE director, and now *Oracle Magazine's* Oracle ACE Director of the Year. Alvarez is also the founder and CEO of Database Integrated Solutions.

For Alvarez, learning—and helping others—takes legwork. This year, as an Oracle ACE director, he has appeared at 22 conferences in 18 countries, where he not only gives talks but takes the time to answer questions and assist users with complex Oracle challenges. In addition, he frequently helps his blog visitors (about 25,000 per month) solve difficult Oracle-related problems.

Given that he's also president of the New Zealand and Chilean Oracle user groups and the Latin American Oracle User Group, Alvarez' energy seems to have no limits when it comes to the Oracle community. He himself puts it best: "I never get tired of helping people learn how to share their knowledge and experience."

© 2010 Oracle Corporation. All rights reserved.



**WINNER STATS**

- Name:** Francisco Munoz Alvarez
- Job title:** Founder and CEO
- Company:** Database Integrated Solutions
- Location:** Auckland, New Zealand
- Award:** Oracle ACE Director of the Year 2010



Professional Expertise Distilled

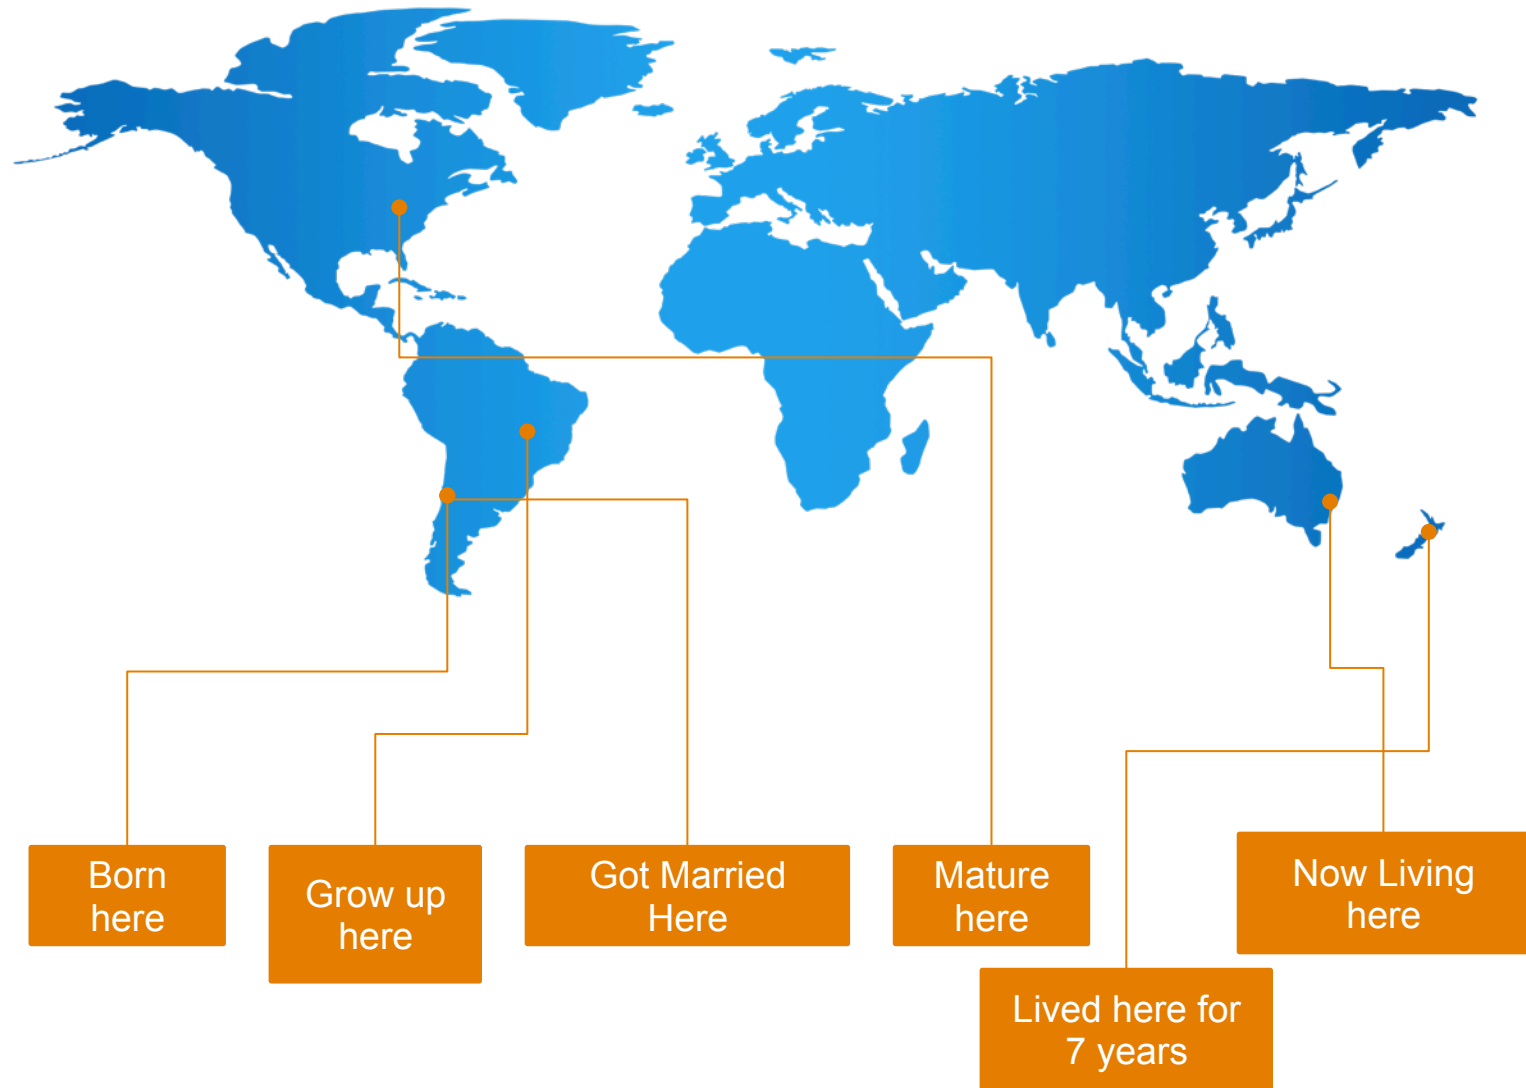
## Oracle Database 12c Backup and Recovery Survival Guide

A comprehensive guide for every DBA to learn recovery and backup solutions

Francisco Munoz Alvarez  
Aman Sharma

[PACKT] enterprise  
PUBLISHING

Available at [AMAZON.COM](https://www.amazon.com)



# Now is time to talk about Oracle VM...

## OVM Tips and Best Practices

1. Be careful with the domain type of the VM you are deploying. Each type will perform different.

The screenshot shows the Oracle VM Manager interface. The left sidebar displays a tree view of Server Pools, including 'WPDBA01C2 B Bottom Server Pool Test' and 'WPDBA01C2B11.OVM.VDC.NZ'. The main panel shows the configuration for a virtual machine named '12c\_Linux\_64bit'. The configuration details are as follows:

Property	Value
Name	12c_Linux_64bit
Status	Running
Operating System	Oracle Linux 6
Keymap	en-us
Max. Processors	4
Processors	4
Max. Memory (MB)	4096
ID	0004fb00000600007e572293f547bfa3
Domain ID	18
Origin	
Description	
Memory (MB)	4096
Processor Cap	100
Priority	50
Mouse Type	Default
Domain Type	Xen PVM
Start Policy	Start on current server
High Availability	No
Repository for Configuration File	OVMM_OVMS_ova
Boot Order	Disk
Network Boot Path	

Below the configuration details, a table lists other virtual machines:

Name	Status	Event Severity	Server	Max. Memory (MB)	Memory (MB)	Max. Processors	Processors	Keymap	Operating System
Linux01	Running	Normal	WPDBA01C2B11.i.1024	1024		4	4	en-us	Oracle Linux 5
RVRWPDPATE01	Running	Normal	WPDBA01C2B11.i.16384	16384		2	2	en-us	Microsoft Windows Server 2008

## OVM Domain Type

Edit Virtual Machine:12c\_Linux\_64bit

Configuration Networks Disks Boot Order Tags

ID: 0004fb00000600007e572293f547bfa3

\* Name: 12c\_Linux\_64bit

Enable High Availability

Repository: OVMM\_OVMS\_ova

Description:

Operating System: Oracle Linux 6

Mouse Device Type: Default

Keymap: en-us (English, United States)

\* Domain Type: Xen PVM

- Xen HVM
- Xen HVM, PV Drivers
- Xen PVM
- OVMM/SPARC
- Unknown

Start Policy:

Max. Memory (MB):

Memory (MB): 4096

Max. Processors: 4

Processors: 4

Priority: 50

Processor Cap %: 100



## OVM Domain Type, Why?

Normal Performance using different domain types:

	OVM (PVM)	OVM (HVM+PV)	OVM (HVM)
TPS MAX	25	23	21
TPS AVG	14	14	13
TPM MAX	921	891	866
TPM AVG	816	742	719
RT MAX	837	532	1526
RT AVG	44	55	111
Read (mb/s)	423.14	312.37	57.06
Write (mb/s)	380.20	297.32	151.58
TRT	35:02.3	1:08:12	1:30:47
RIPS	81,953	67,198	51,441
DGPS (MB)	6.10	5.07	3.80

Do you see the difference?

If using HVMPV think about using HUGE pages.

## OVM Tips and Best Practices

2. Be careful when doing a V2V or P2V.

	OVM (PVM)	OVM (HVM+PV)	OVM (HVM)
TPS MAX	25	23	21
TPS AVG	14	14	13
TPM MAX	921	891	866
TPM AVG	816	742	719
RT MAX	837	532	1526
RT AVG	44	55	111
Read (mb/s)	423.14	312.37	57.06
Write (mb/s)	380.20	297.32	151.58
TRT	35:02.3	1:08:12	1:30:47
RIPS	81,953	67,198	51,441
DGPS (MB)	6.10	5.07	3.80

When a VM is created using V2V or P2V, it will be by default using the HVM domain type.

## OVM Tips and Best Practices

3. Avoid as much you can using OCFS2.

	OVM (PVM-RD)	OVM (PVM)	OVM (HVM+PV-RD)	OVM (HVM+PV)	OVM (HVM)
TPS MAX	25	25	25	23	21
TPS AVG	18	14	14	14	13
TPM MAX	933	921	912	891	866
TPM AVG	838	816	826	742	719
RT MAX	164	837	215	532	1526
RT AVG	18	44	29	55	111
Read (mb/s)	404.52	423.14	400.32	312.37	57.06
Write (mb/s)	380.20	380.20	388.30	297.32	151.58
TRT	31:52.0	35:02.3	59:03.3	1:08:12	1:30:47
RIPS	93,542	81,953	87,406	67,198	51,441
DGPS (MB)	7.00	6.10	6.50	5.70	3.80

Avoid as much you can using OCFS2

Avoid using virtual disks on OCFS2 and present a LUN (Raw device) directly to the guest VM. You will possible achieve a better performance than a bare metal doing this.

	Bare Metal	OVM (PVM-RD)	OVM (HVM+PV-RD)
TPS MAX	27	25	25
TPS AVG	14	18	14
TPM MAX	924	933	912
TPM AVG	810	838	826
RT MAX	277	164	215
RT AVG	35	18	29
Read (mb/s)	399.15	404.52	400.32
Write (mb/s)	469.87	380.20	388.30
TRT	53:32.0	31:52.0	59:03.3
RIPS	45,161	93,542	87,406
DGPS (MB)	3.40	7.00	6.50

## OVM Tips and Best Practices

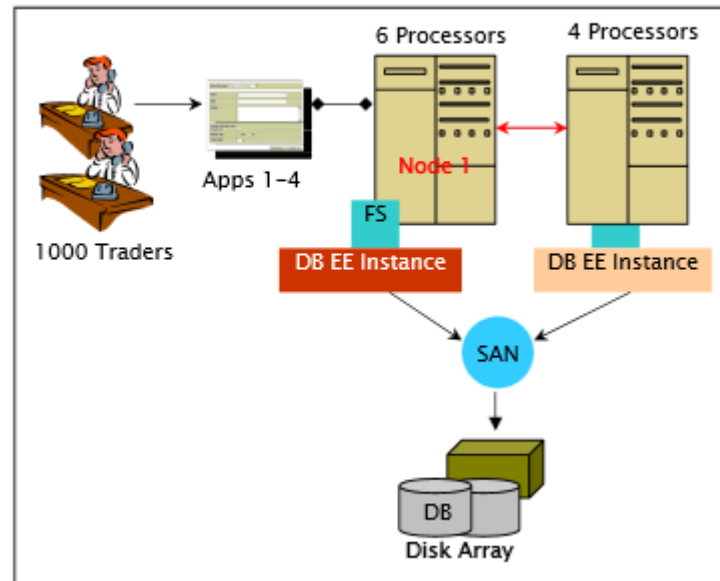
4. Implement HA for your Oracle VM environment.
  - a) Have a DR site
  - b) Always activate HA on the VMs
  - c) When possible always use Live Migration to move your VMs (Even if the CPUs are pinned)

## Software Investment Guide

*Failover* - In this type of recovery, nodes are arranged in a cluster and share one disk array. A Failover cluster is a group of systems, bound together into a common resource pool. In this type of recovery method, the Production node acts as the primary node. When the primary node fails, one of the surviving nodes in the cluster acts as the primary node. Solutions like Oracle Failsafe (included with Oracle Database EE or SE, SE1), or third party vendor solutions (e.g. Veritas, HP Service Guard, HACMP, Linux HA - Heartbeat) are used to manage Failover environments. In this type of environment, Oracle permits licensed Oracle customers to run some Technology Programs on an unlicensed spare computer for up to a total of ten separate days in any given calendar year. Once the primary node is repaired, you must switch back to the primary node. Once the failover period has exceeded ten days, the failover node must be licensed. In addition, only one failover node per clustered environment is at no charge for up to ten separate days even if multiple nodes are configured as failover. Downtime for maintenance purposes counts towards the ten separate days limitation. Any other use requires the environment to be fully licensed. . In a failover environment, the same license metric must be used for the production and failover nodes when licensing a given clustered configuration. Additionally, when licensing options on a failover environment, the options must match the number of licenses of the associated database. See illustration #4.

## Software Investment Guide

Illustration #4: Failover



- Failsafe included with Oracle Database EE is installed on nodes 1 and 2; it is running on node 1 only
- If node 1 fails, node 2 takes over (there is always an idle node with Failsafe or other vendor cluster solutions)

# Migrating your VMs with pinned vCPUs

When live migrating a Guest VM that have vCPUs pinned (due that you need to patch/upgrade the Oracle VM server where the VM is running), always remember to:

1. Manually pin the VM vCPUs to physical CPUs in the destination server just after the live migration of the VM is completed.
2. When the patching/upgrade required on the Source server is completed, live migrate the guest VM back to this server and manually pin the vCPUs again.
3. Be careful to not use Live Migration for more than 10 days per year (as per the example in the previous 2 slides). If you violate this policy you will need to license your whole cluster.



# OVM Tips and Best Practices

## 5. OEM12c is a good friend!

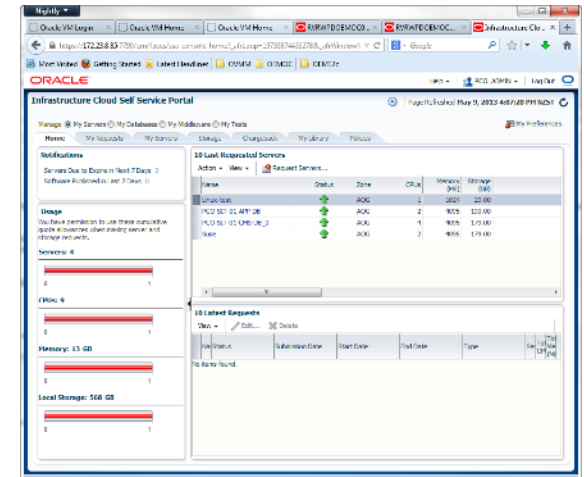
The screenshot displays the Oracle Enterprise Manager Cloud Control 12c interface. The main content area shows the configuration and monitoring for a VM Guest named '12c\_Linux\_64bit'. The interface is divided into several sections:

- General:** OVM Display Name: 12c\_Linux\_64bit, OVM State: Running, Oracle VM Server Pool: WPDBA01C2 B Bottom Server Pool Test, Oracle VM Server: WPDBA01C2B11.OVM.VDC.NZ, Host: 172.29.62.4, Oracle Assembly Instance: Description.
- Job Activity:** Summary of jobs whose start date is within the last 7 days. A table shows no jobs found.
- Configuration:** VCPUs: 4, CPU Priority: 50, CPU Utilization Cap: 100, Allocated Storage (GB): 100, Allocated Memory (MB): 4096, Memory Required (MB): 4193, High Availability: Disabled, Domain Type: Xen PVM.
- Networks:** A table with columns for MAC Address, IP Address, IP Address Type, and Ethernet Network. One entry is shown: MAC Address: 00:21:f6:00:00:02, IP Address Type: RVR Mgmt.
- Monitoring Charts:**
  - CPU Utilization:** A line chart showing Physical CPU Utilization over time, with a peak around 11:08 AM.
  - Network Activity:** A line chart showing Total Throughput (kb/s) over time, with a peak around 11:07 AM.
  - Disk Activity:** A section indicating 'No data available'.

## OEM12c

### Benefits when using Portal:

- User will only see their own VM's
- When the user right click on the VM they will get the following options:
  - **Modify Configuration**
  - **Clone**
  - Delete
  - Start
  - Stop
  - Restart
  - Start and stop
  - Suspend
  - Resume
  - Launch VNC console



They will not be able to Clone or modify the configuration of any VM if you have removed their quotas.

**And it is free for IaaS without charge back!**

## OVM Tips and Best Practices

6. Do not over allocate the CPU of a Oracle VM Server.

Example:

- OVM Server has 2 hexa core = 12 CPUs
- With Multithread activated = 24 vCPUs

*Always reserve 1 physical CPU per Processor. In this case 2 hexa core = 2 CPUs = 4 vCPUs.*

**Maximum available for VMs: 20 vCPUs**

```
# xm info
nr_cpus : 8
nr_nodes : 1
cores_per_socket : 4
threads_per_core : 2
cpu_mhz : 3200
```

## OVM Tips and Best Practices

7. If want to save money in Licenses, always remember to PIN your VMs vCPU.

- PIN using the Oracle VM Utilities (ovm\_vmcontrol) [OVM 3.1 and up]
- PIN it at configuration file Level (vm.cfg) [OVM 2]

*“Hard partitioning means binding a virtual machine CPU to a physical CPU or core, and preventing it from running on other physical cores than the ones specified. “*

# `xm vcpu-list` command shows a summary of which virtual CPUs are running on which physical CPUs.

## Hard Partitioning

The Oracle VM 3 Utilities are a collection of command line scripts that allow you to perform a set of basic management tasks. The Oracle VM Utilities are available for download as a .zip file from My Oracle Support, search for patch ID 13602094.

## OVM Tips and Best Practices

8. Have on mind all OS you can run on Oracle VM.

- Red Hat
- Oracle Linux
- Solaris
- Windows

Is these all?

Think Again.....

## Guest OS for PV, HVM and HVMPV

**Table 1. 64-bit CPU Hardware Virtualized Supported Guest Operating Systems (Linux and Solaris)**

Guest Operating System	Hardware Virtualized 32-bit	Hardware Virtualized 32-bit with PV Drivers	Hardware Virtualized 64-bit	Hardware Virtualized 64-bit with PV Drivers
Oracle Linux 7.x	N/A	N/A	Yes	Yes
Oracle Linux 6.x	Yes	Yes	Yes	Yes
Oracle Linux 5.x	Yes	Yes	Yes	Yes
Oracle Linux 4.x	Yes	Yes	Yes	Yes
Oracle Solaris 11 <sup>[a]</sup>	N/A	N/A	N/A	Yes
Oracle Solaris 10 <sup>[a]</sup>	N/A	Yes	N/A	Yes
Red Hat Enterprise Linux 7.x	N/A	N/A	Yes	Yes
Red Hat Enterprise Linux 6.x	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 5.x	Yes	Yes	Yes	Yes
Red Hat Enterprise Linux 4.x	Yes	Yes	Yes	Yes
CentOS 6.x	Yes	Yes	Yes	Yes
CentOS 5.x	Yes	Yes	Yes	Yes
CentOS 4.x	Yes	Yes	Yes	Yes
SUSE Linux Enterprise Server 11.x <sup>[b]</sup>	No	No	Yes	Yes

## Guest OS for PV, HVM and HVMPV

Oracle Solaris support begins with Solaris 10 10/09. The Solaris 10 or Solaris 11 OS runs as a hardware virtual machine (HVM), which requires HVM support (Intel VT or AMD-V) on the underlying hardware platform. By default, Solaris 10 or Solaris 11 OS already has the required paravirtualized (PV) drivers installed as part of the OS. Oracle Solaris 10 supports x86 32-bit and 64-bit architecture. 32-bit or 64-bit mode is selected at OS boot time by examining the hypervisor and the underlying hardware. Oracle Solaris 11 supports x86 64-bit architecture only.

The minimum required kernel version for SLES 11 is 3.0.31-0.9. To start a PVHVM guest on SLES 11, the following steps are required:

1. Install the operating system and upgrade to the 3.0.31-0.9 kernel version.
2. Shutdown the virtual machine.
3. In Oracle VM Manager, edit the Domain Type for the virtual machine to set it to "Xen PVM".
4. Restart the virtual machine.



## Guest OS for PV, HVM and HVMPV

**Table 2. 64-bit CPU Hardware Virtualized Supported Guest Operating Systems (Microsoft Windows)**

Guest Operating System	Hardware Virtualized 32-bit	Hardware Virtualized 32-bit with PV Drivers 3.x	Hardware Virtualized 64-bit	Hardware Virtualized 64-bit with PV Drivers 3.x
Microsoft Windows Server 2012 R2	N/A	N/A	Yes	Yes
Microsoft Windows Server 2012	N/A	N/A	Yes	Yes
Microsoft Windows 8.1	Yes	Yes	Yes	Yes
Microsoft Windows 8	Yes	Yes	Yes	Yes
Microsoft Windows™ 7 SP1	Yes <sup>[a]</sup>	Yes	Yes <sup>[a]</sup>	Yes
Microsoft Windows™ Vista SP2	Yes <sup>[a]</sup>	Yes	Yes <sup>[a]</sup>	Yes
Microsoft Windows™ Server 2008 R2 SP1	N/A	N/A	Yes <sup>[a]</sup>	Yes
Microsoft Windows™ Server 2008 SP2	Yes <sup>[a]</sup>	Yes	Yes <sup>[a]</sup>	Yes
Microsoft Windows™ Server 2003 R2 SP2	Yes <sup>[a]</sup>	Yes	Yes <sup>[a]</sup>	Yes
Microsoft Windows™ Server 2003 SP2	Yes <sup>[a]</sup>	Yes	Yes <sup>[a]</sup>	Yes

Note: HVM-only mode is used to facilitate the Windows PV Drivers installation. Windows PV Drivers are required on the Windows OS if available in order to be supported as a guest OS.

## More **Oracle VM Server/Windows Guests**

Windows paravirtual (PV) drivers are high-performance network and disk drivers that significantly reduce the overhead of the traditional implementation of I/O device emulation. These drivers provide improved network and disk throughput to run fully virtualized Windows guests in an Oracle VM Server for x86 environment.

Oracle VM Windows PV Drivers are signed by Microsoft. Oracle is a participant of Microsoft Windows Server Virtualization Validation Program (SVVP). Windows PV Drivers along with Oracle VM Server for x86 have passed the SVVP requirements, [read the details here](#). The Microsoft SVVP program enables the vendor such as Oracle to validate Oracle VM Server for x86 along with Windows PV Drivers so that Microsoft customers running copies of Windows Server they have acquired and licensed from Microsoft directly can receive technical support for Windows Server in virtualized environments,

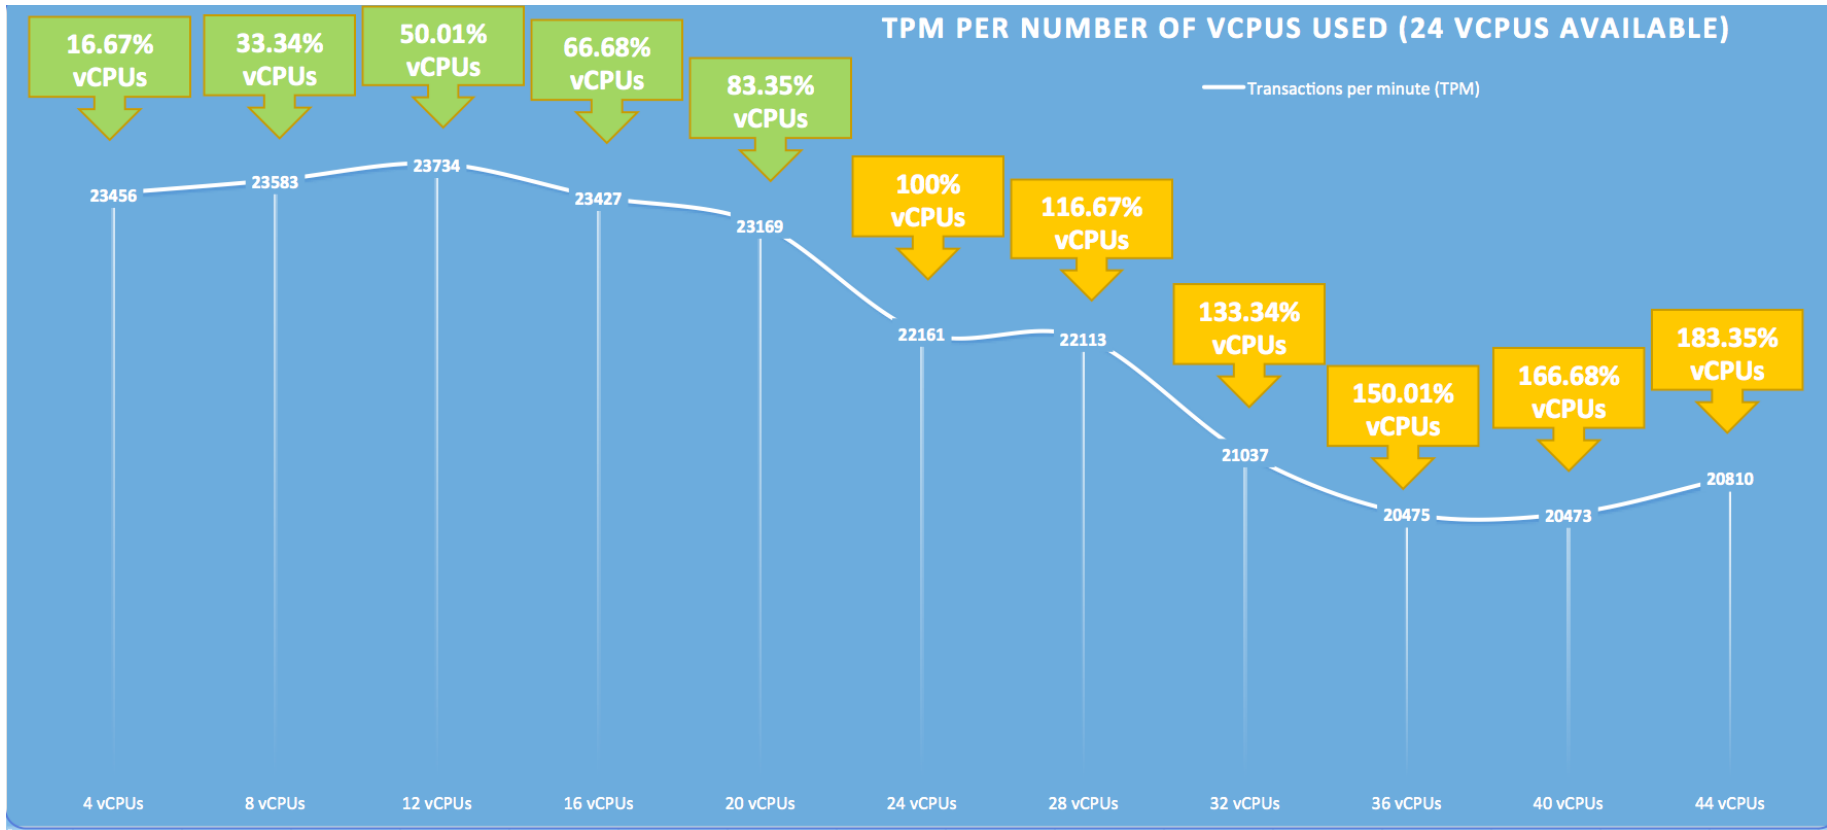
[learn the details of the supported Microsoft server software \(e.g. SQL Server, Exchange Server, etc.\) here](#).

## Plus **Oracle VM Server for x86 SPARC Guest Operating Systems**

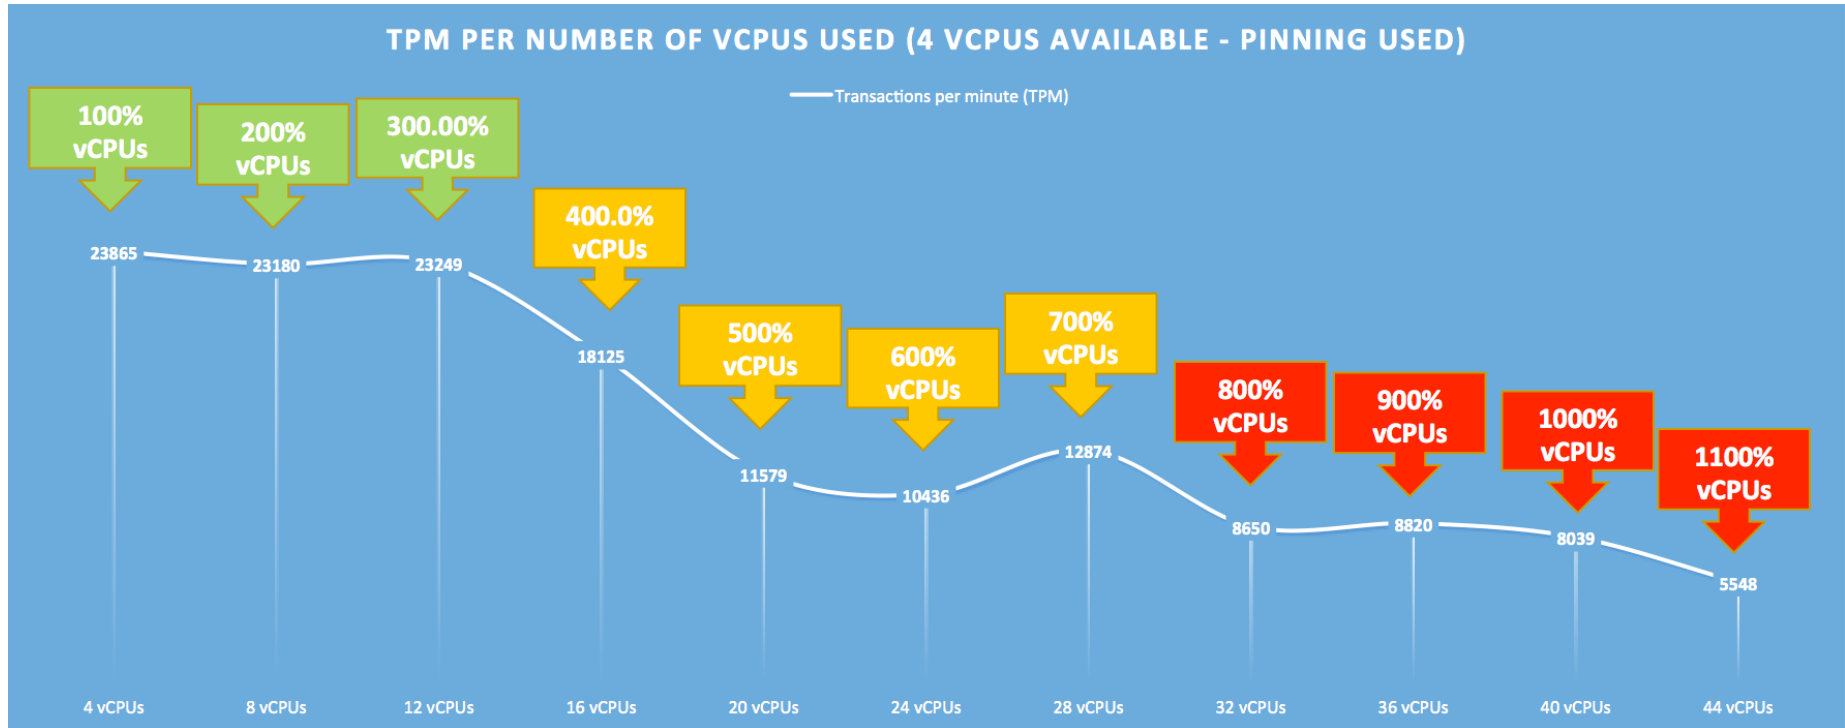
The guest virtual machines you create in an Oracle VM Server for SPARC-based server pool must be one of the following supported configurations.

- Oracle Solaris 11
- Oracle Solaris 10 8/11 or later

# Oracle VM and CPU Over Committing (2 sockets, 12 cores)



# Oracle VM and CPU Over Committing (1 core for Oracle DB License)



## Conclusions

- Choose the domain type correctly due that it will affect the performance of your VM.
- **Be careful when doing V2V or P2V.**
- If possible do not use virtual disk, present a LUN directly to the VM instead.
- **Implement HA in your environment and use live migration even if you have your VM vCPUs pinned.**
- Use OEM12c to manage your Oracle VM environment.
- **Do not over use physical CPU of a Oracle VM Server.**
- Use Hard Partitioning to save in Oracle Licenses.
- **Oracle VM is flexible and can run multiple OS as a Guests.**
- Oracle VM and CPU over committing

References:

My Blog:

<http://www.oraclenz.org>

Certificate Software on Oracle VM:

[My Oracle Support 464754.1](#)

Oracle VM Resource page:

<http://www.oracle.com/us/technologies/virtualization/oraclevm/resources/index.html>

# Questions? Comments?

Francisco Munoz Alvarez

Pythian APAC  
alvarez@pythian.com



# Thank you!!