

VMotion between Apples and Oranges

Understanding CPU Compatibility Constraints for VMware® VMotion™

Matthias Hausner

Manager, Software Reliability

VMware



VMWORLD 2006

Common Questions

- I have two XXX servers and plan to buy 4 new YYY systems. Will I be able to use VMotion between them?
- When I attempt to VMotion between hosts X and Y, I get an error message “problem detected at CPUID level 0x1 register ecx.”

What's wrong?

Compatibility:

- w2k3onSanMelody
- ▲ Unable to migrate from giza.vmware.com to eeng183.vmware.com: The CPU of the host is incompatible with the CPU feature requirements of virtual machine; problem detected at CPUID level 0x1 register 'ecx'



Outline

- Terminology
- Why are there CPU compatibility constraints?
- Incompatible CPU features
- Compatibility groups of CPUs
- Relaxation of restrictions
- Questions

Terminology

- Moving a VM from one host to another is called migration.
- **Cold Migration**
Power off VM on host A, power on on host B
- **Hot Migration**
Suspend running VM on host A, resume on host B
- **Migration with VMotion**
Form of Hot Migration
Migrate running VM from host A to host B

Requirements for VMotion

- Source and destination hosts must be:
 - Part of the same Datacenter
 - Connected to the same Gigabit network
 - Connected to the same storage (SAN, NAS, iSCSI)
- Recommended: dedicated Gigabit network for VMotion
- Destination host must have enough resources
- Source and destination hosts must have **compatible CPU models**

Why Are Some CPUs Not Compatible?

- Why can't I VMotion between any two CPUs?
 - Most likely outcome is guest crash if CPUs are not compatible
 - VMware ESX is a stable, reliable platform: crashing is not an option
- CPUs are pass-through devices
 - Guest software sees the physical CPU of the host
- Example 1: 64 bit support
 - Can not run a VM with Windows XP x64 Edition on a host which does not have a 64 bit CPU
 - Can not migrate a VM with a 64 bit guest OS to a host which does not have a 64 bit CPU
- Example 2: Different CPU vendors
 - AMD and Intel CPUs have slightly different instruction sets
 - Different Model-specific Registers (MSRs)
 - Therefore: cannot VMotion from one CPU vendor to another

Which CPU Features matter?

- **CPU Vendor**
- **CPU Family** (Pentium 4, Intel Core, AMD K8, AMD K7)
- **Support for 64 bit guests:**
 - Intel: EM64T and VT
 - AMD: AMD64 on Rev D and newer CPUs
- **NX/XD**
 - Ability to mark memory pages a non-executable
- **SSE3**
 - New instructions for optimized signal processing, 3D graphics, and hyperthreading.

Which CPU Features matter?

- **SSSE3** (aka SSE4, Merom New Instructions)
 - Intel Core 2 family only
 - Supplemental SSE3 instructions
- **FFXSR**
 - AMD only (Rev D)
 - Optimization for FXSAVE/FXRSTOR instructions
- **RDTSCP**
 - AMD only (Rev F)
 - Read serialized TSC Pair
- **CMPXCHG16B**
 - AMD only (Rev F), 64bit only
 - All VMware supported Intel EM64T CPUs have this feature

CPUID Instruction

- CPU reports vendor, family, model and feature information with **CPUID** instruction
- Input value: CPUID level
- Output value: 4 32-bit values in registers

Feature	CPUID function
NX/XD	cpuid(80000001h).edx[20]
SSE3	cpuid(1).ecx[0]
SSSE3	cpuid(1).ecx[9]
FXSR	cpuid(80000001h).edx[25]
RDTSCP	cpuid(80000001h).edx[27]
CMPXCHG16B	cpuid(1).ecx[13]

Compatibility Groups: Pentium 4 Family

Pentium 4 CPU Model	CPU Feature		VI3 Groups	VC 1.x Groups
	SSE3	NX		
Willamette, Northwood, Foster, Prestonia, Gallatin	No	No	A	A
Nocona Cedar Mill	Yes	No	B	B
Xeon 50xx (Dempsey) Xeon 70xx (Paxville) Xeon 71xx (Tulsa) Prescott	Yes	Yes	C	

Compatibility Groups: Intel Core Family

Intel Core and Core 2 CPU Model	CPU Feature	VI3 Groups	VC 1.x Groups
	SSSE3		
Xeon LV (Sossaman) Yonah	No	A	A
Xeon 51xx (Woodcrest) Xeon 53xx (Clovertown) Merom Conroe	Yes	B	B

Compatibility Groups: AMD K8

CPU Model	CPU Feature			VI3 Groups	VC 1.x Groups
	FFXSR	SSE3	RDTSCP		
Rev C and earlier	No	No	No	A	A
Rev D	Yes	No	No	B	
Rev E	Yes	Yes	No	C	
Rev F	Yes	Yes	Yes	D	

Example: Dell PowerEdge 1950

- Dell PowerEdge 1950 servers come with different Intel Xeon CPUs:
 - Xeon 5000 Series is Pentium 4 family (Dempsey core)
 - Xeon 5100 Series is Intel Core 2 family (Woodcrest core)
- Must be very careful when choosing hardware mix!

SYSTEM OPTIONS

PowerEdge 1950

Dual Core Intel® Xeon® 5110, 4MB Cache, 1.60GHz,

[? Help Me Choose](#)

- Dual Core Intel® Xeon® 5050, 2x2MB Cache, 3.00GHz, 667M
- Dual Core Intel® Xeon® 5060, 2x2MB Cache, 3.20GHz, 1066
- Dual Core Intel® Xeon® 5080, 2x2MB Cache, 3.73GHz, 1066
- Dual Core Intel® Xeon® 5110, 4MB Cache, 1.60GHz, 1066MI
- Dual Core Intel® Xeon® 5130, 4MB Cache, 2.00GHz, 1333MI
- Dual Core Intel® Xeon® 5140, 4MB Cache, 2.33GHz, 1333MI
- Dual Core Intel® Xeon® 5150, 4MB Cache, 2.66GHz, 1333MI

Compatibility Matrix for Dell PowerEdge

VMotion compatibility matrix for VMware Virtual Infrastructure 3.0

The following table provides the VMotion compatibility matrix on Dell™ PowerEdge™ (PE) servers for VMware Virtual Infrastructure 3.0. Read the notes below for important VMotion requirements and information.

VMotion Compatibility across Dell PowerEdge Servers		Legacy	8th Generation PE Servers				9th Generation PE Servers with Intel Xeon 5000 series			9th Generation PE Servers with Intel Xeon 5100 series		
		PE 6650	PE 1850	PE 2850	PE 1855	PE 6850	PE 1950	PE 2950	PE 1955	PE 1950	PE 2950	PE 1955
Legacy	PE 6650	Yes	No	No	No	No	No	No	No	No	No	No
8th Generation PE Servers	PE 1850	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
	PE 2850	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
	PE 1855	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
	PE 6850	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
9th Generation PE Servers with Intel Xeon 5000 series	PE 1950	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
	PE 2950	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
	PE 1955	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
9th Generation PE Servers with Intel Xeon 5100 series	PE 1950	No	No	No	No	No	No	No	No	Yes	Yes	Yes
	PE 2950	No	No	No	No	No	No	No	No	Yes	Yes	Yes
	PE 1955	No	No	No	No	No	No	No	No	Yes	Yes	Yes

Legend:

Yes – VMotion works between the corresponding two servers/processor models.

No – VMotion fails with CPU incompatibility error message

© 2006 Dell Inc. (www.dell.com/downloads/global/solutions/vmotion_compatibility_matix.pdf)

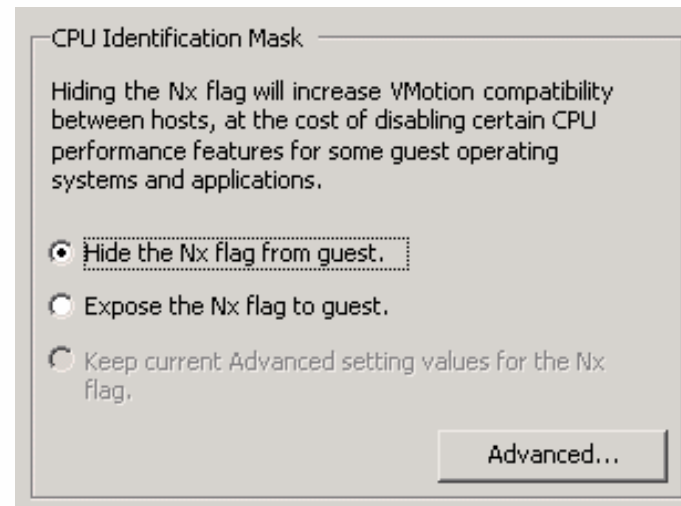
Compatibility Matrix for Other Vendors

- No other vendor has so far published VMotion compatibility data

Compatibility Relaxation: NX/XD

- VMotion between CPUs with and without NX allowed if NX bit is hidden from the guest OS
- In VI3, the NX CPU feature can be hidden from the guest
- This relaxation is supported by VMware

- OSes that use NX if present:
 - Windows XP SP2
 - Windows Server 2003
 - Windows Vista
 - RedHat RHEL 3 Upd 3, RHEL 4
 - SUSE Linux 9.2, 10
 - Solaris 10

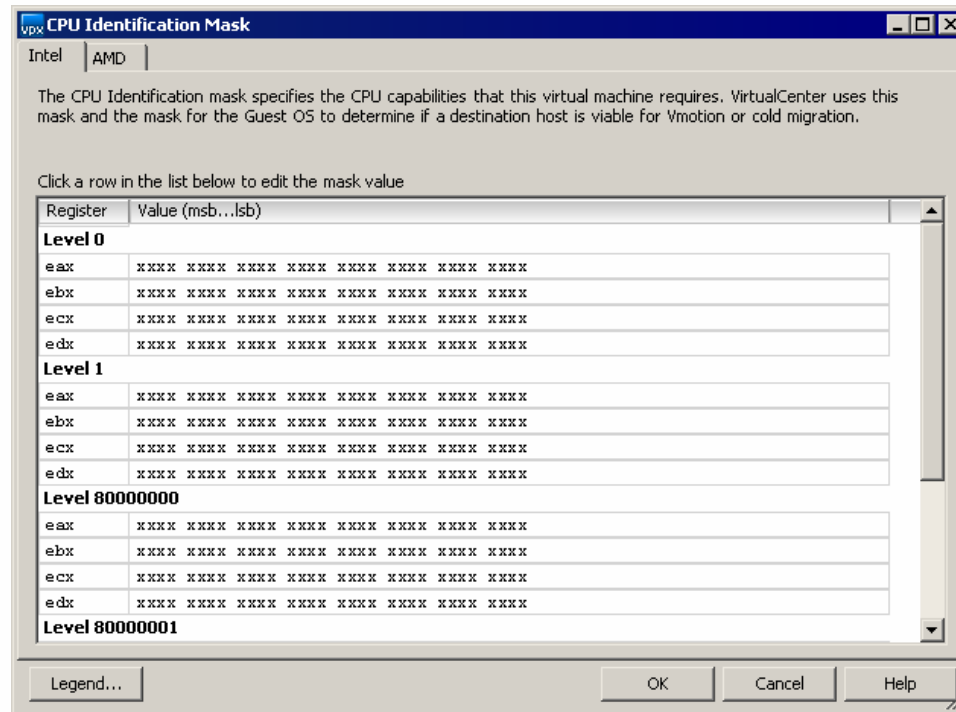


Pentium 4 Groups with NX Relaxation

Pentium 4 CPU Model	CPU Feature		VI3 Groups	VC 1.x Groups
	SSE3	NX		
Willamette, Northwood, Foster, Prestonia, Gallatin	No	No	A	A
Nocona Cedar Mill	Yes	No	B	B
Xeon 50xx (Dempsey) Xeon 70xx (Paxville) Xeon 71xx (Tulsa) Prescott	Yes	Yes		

Relaxing Other CPU Constraints

- CPUID values and requirements can be changed in the Advanced section of the CPU Identification Mask dialog box
- For each CPUID level a mask for each of the 4 registers can be specified



CPUID Mask Values

Mask Character	X	0	1	F	T	H	R
Required host value				0	1		
Guest sees	HOST	0	1	HOST (0)	HOST (1)	HOST	0
Checked on hot migration?	No	No	No	Yes	Yes	Yes	Yes

“HOST” means: host value at VM startup time

CPUID Masks in the VM Configuration File

- VM configuration file can specify CPUID values for the guest

- Format:

`cpuid.<level>.<register> = <mask>`

`cpuid.<level>.<register>.amd = <mask>`

- Mask is of the form: “----:----:----:----:----:----:----:----”

- 32 mask characters: -, X, 0, 1, F, T, H, R

- Dash “-” means: leave default mask value unchanged

- Example: hide NX bit from guest

(Remember: NX/XD is CPUID level 80000001h, reg EDX, bit 20)

`cpuid.80000001.edx = "----:----:---0:----:----:----:----:----"`

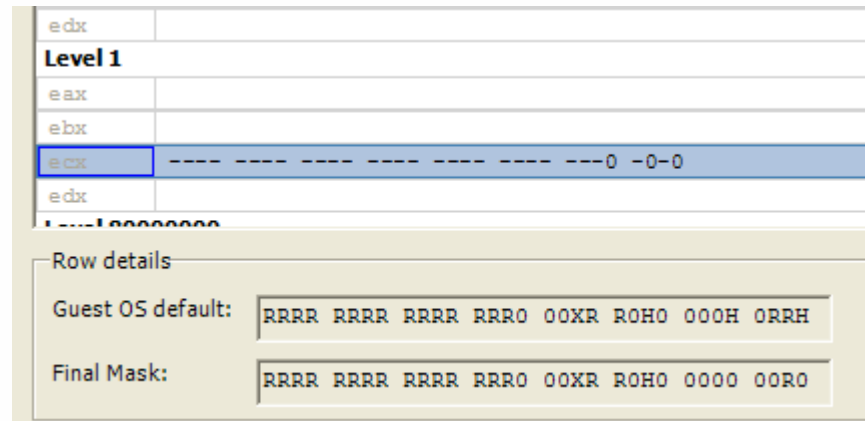
Guest sees bit 20 of register edx as 0 on Intel and AMD cpus

Relaxing SSE3 Requirement on Pentium 4

- Goal: Enable VMotion between SSE3 and non-SSE3 CPUs
- Might be useful to enable VMotion between older and newer Pentium 4 based CPUs
- Warning: relaxing SSE3 compatibility requirement is **not supported** by VMware
- Works only if neither OS nor any application uses SSE3 instructions
- Why not supported?
 - SSE3 feature can be used by user-level code (applications)
 - Mask **does not work for user-level code** (i.e. applications)
 - In user-level code, CPUID is executed directly on hardware and is not intercepted by VMware.
 - Thus, VM cannot reliably hide SSE3 from an application
- Better: cold migration


Relaxing SSE3 Requirement on Pentium 4

- Method 1:
Using the VC2 UI



- Method 2: Config file setting
cpuid.1.ecx = "-----:-----:-----:-----:-----:-----:---0:-0-0"
- Guest will see SSE3, reserved bit 2, and bit 4 (DS-CPL) as zero
 - ▶ Bit 2 (reserved) and bit 4 (DS-CPL) are 1 on all Intel CPUs which support SSE3

Pentium 4 Groups with SSE3 Relaxation

Pentium 4 CPU Model	CPU Feature		VI3 Groups	VC 1.x Groups
	SSE3	NX		
Willamette, Northwood, Foster, Prestonia, Gallatin	No	No		A
Nocona Cedar Mill	Yes	No		B
Xeon 50xx (Dempsey) Xeon 70xx (Paxville) Xeon 71xx (Tulsa) Prescott	Yes	Yes	C	

Supported Relaxations

- Only OS-level CPU features can be relaxed safely and are supported by VMware.

- Supported Relaxations
 - NX/XD
 - RDTSCP

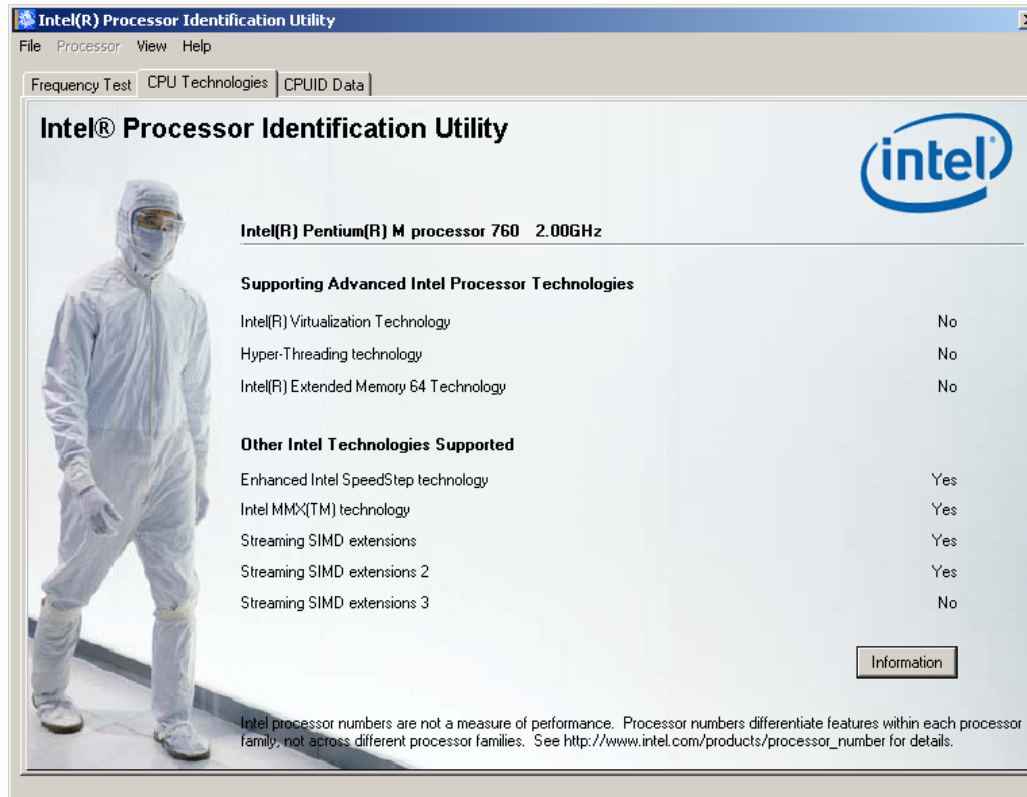
- Unsupported Relaxations
 - SSE3
 - SSSE3
 - FFXSR
 - CMPXCHG16B

Cold Migration

- Cold migration: power off VM and power on on different host
- No CPU compatibility checks
- 64 bit guest will most likely fault on 32 bit CPU
 - (ESX 3.0.1: power-on will not be permitted in this case)
- Vendor is not checked
 - Known case that does not work:
 - Installing Linux on an AMD system, then later power on on Intel CPU will cause VM to hang.
 - Cause: AMD-specific “prefetch” instruction (used in kernel fault handler code) leads to recursive faults.

Determining CPU Models and Features

■ Intel Processor Identification Utility



Intel(R) Processor Identification Utility

File Processor View Help

Frequency Test CPU Technologies CPUID Data

Intel® Processor Identification Utility

Intel(R) Pentium(R) M processor 760 2.00GHz

Supporting Advanced Intel Processor Technologies

Intel(R) Virtualization Technology	No
Hyper-Threading technology	No
Intel(R) Extended Memory 64 Technology	No

Other Intel Technologies Supported

Enhanced Intel SpeedStep technology	Yes
Intel MMX(TM) technology	Yes
Streaming SIMD extensions	Yes
Streaming SIMD extensions 2	Yes
Streaming SIMD extensions 3	No

Information

Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details.

Determining CPU Models and Features

- VMware CPUID program on bootable floppy or CDROM image

```
TEST: 56983: CPUID CHANGE: 340063
Reporting CPUID for 1 logical CPU...

  Family: 0f Model: 04 Stepping: 1

  ID1ECX      ID1EDX      ID81ECX      ID81EDX
  0x0000659d  0xbfefbfff  0000000000  0x20100000

Vendor       : Intel
Processor Cores : 1
SSE Support  : SSE3
Supports NX / ED : Yes
Hyperthreading : Yes
Supports 64-bit Longmode : Yes
Supports 64-bit VMWare : No

Copyright (c) 1998-2006 VMware, Inc. All rights reserved.
Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847,
6,704,925, 6,711,672, 6,725,289 and 6,735,601; patents pending.
VMware, the VMware "boxes" logo, GSX Server, ESX Server,
Virtual SMP and VMotion are trademarks of VMware, Inc.

PASS: Test 56983: CPUID
Press any key to reboot.
```

VC 1.3 versus VI3

- ESX 2.5/VC 1.3 does not support 64 bit guests
- ESX 2.5/VC 1.3 hides these bits from guests by default:
NX, FFXSR, RDTSCP
- CPU compatibility factors:
 - CPU Vendor
 - CPU Family
 - SSE3, SSSE3
- Relaxation masks have different format
 - See VMware KB article 1993 for details

Upgrade VMotion

- Special case: one-way VMotion from ESX 2.5.x to ESX 3.0.1 or later
- Virtual disks moved from VMFS2 to VMFS3 volume
- CPUID masks will automatically be set to match ESX 2.5.x behavior
 - NX, FFXSR, RDTSCP bits hidden from guest
- Some requirements (persistent disks, consecutive vNIC numbers)
- After successful Upgrade VMotion:
 - VMware Tools can be upgraded in running VM
 - Still need to shut down VM to upgrade virtual hardware version

References

- VMware VI3 documentation
- VMware KB articles 1991, 1992, 1993
- Intel: Application Note 485: Processor Identification and CPUID Instruction
- AMD: Publication 25481: CPUID Specification
- Dell: VMware VMotion Compatibility Matrix for Dell PowerEdge Servers
- Intel web page “Intel Processor Numbers”
(http://www.intel.com/products/processor_number/index.htm)
- Wikipedia, Keywords:
 - > Xeon, Pentium 4, Intel Core, Intel Core 2
 - > Opteron, Athlon
 - > SSE3, SSSE3, NX bit

Thank You

- The VMware Customer Support team says Thank You for your attention



Questions

Presentation Download

Please remember to complete your
session evaluation form
and return it to the room monitors
as you exit the session

The presentation for this session can be downloaded at
<http://www.vmware.com/vmtn/vmworld/sessions/>

Enter the following to download (case-sensitive):

Username: cbv_rep
Password: cbvfor9v9r

Some or all of the features in this document may be representative of feature areas under development. Feature commitments must not be included in contracts, purchase orders, or sales agreements of any kind. Technical feasibility and market demand will affect final delivery.

VMWORLD 2006

