



AMD Extended Migration

Elsie Wahlig
8/30/2007

Migration 201

Migration is the ability to move a VM to another physical machine under the control of the Hypervisor

- Offline or Cold Migration: suspend or hibernate the guest
- Online or Live Migration: move the guest while it is running

Once a VM is booted, it has expectations that must be maintained or emulated by software

- ISA expectations: ISA must not change out from under the guest
- Device Dependencies

Virtualize the ISA to provide a consistent view for VMs

AMD **has already** implemented technology, called AMD-V™ Extended Migration, that will enable ISA compatibility for features between the range of AMD Opteron processors (C, D, E, F, Barcelona and beyond)

CPUID Feature Identifier Registers

OS and Applications use existing CPU Feature Identifier registers to find out the processor's instruction set capabilities

- CPUID Functions 0000_0001 and 8000_0001 describe the instruction set features in ECX and EDX

The main issue for migration is user-accessible instructions that are available on one generation but not another; we use CPUID masking to hide these features

- OS-only instructions are generally not a problem
- Well-behaved user code checks the features reported via CPUID
- Try and Catch to detect the opcode behavior
 - SSE1 and Monitor/MWait
 - SSE1 is available on all AMD Opteron™ processors and Monitor/MWait can be disabled.

Hypervisor must enforce migrations to platforms that can safely support the Guest applications' ISA requirements

Configuration Software

Hypervisor provides the least common denominator available in the migration set

- *Recognizes users' compatibility preferences*
 - *Simple Compatibility: Machine level*
 - *Advanced Compatibility: Feature level*
- *Translates designated machines into appropriate compatibility level*
 - *Hypervisor understands Defaults for Family of Machine identified*
 - *Hypervisor translates to feature level compatibility*

Make the Admin's responsibilities simple

- Select machines for compatibility set
- Assign the appropriate VM compatibility mode

AMD Opteron™ Processor Evolution – ISA and System Features



Family Fh – Rev C

AMD Opteron™ 840, 840 EE, 842, 846 HE, 850 ...
 AMD Opteron™ 240, 242, 244, 246, 246 HE, 248, ...
 AMD Opteron™ 140, 140EE, 146HE, 148, 150, ...

SSE

MMX

SSE2

Family Fh – Rev E

AMD Opteron™ 860, 860 HE, 865, 870, 875, 890, ...
 AMD Opteron™ 260 HE, 265, 275 HE, 280, 290, ...
 AMD Opteron™ 165, 185, ...

SSE3

Family Fh – Rev F

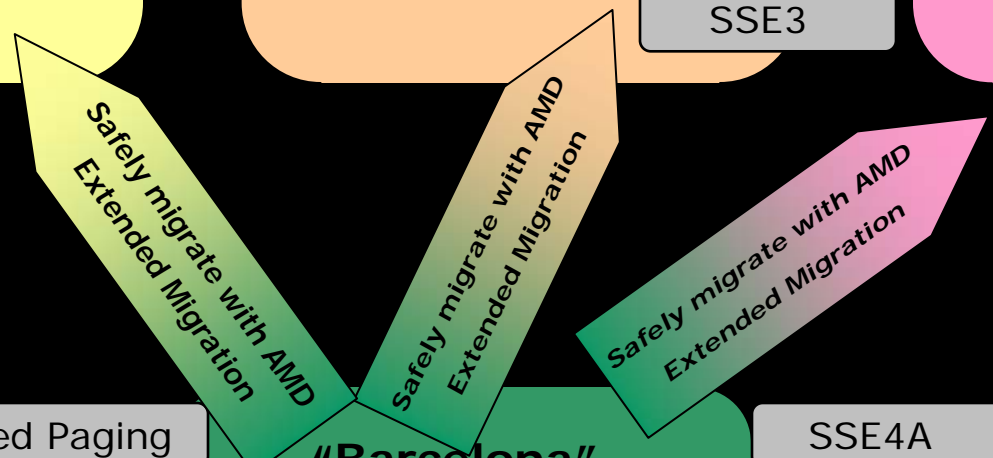
AMD Opteron™ 8212, 8220, 8218, 8222 SE, ...
 AMD Opteron™ 2220, 2212, 2214, 2210, ...
 AMD Opteron™ 1222 SE, 1214, 1210, ...

SVM*

CMPXCHG16B

RDTSCP*

SKINIT*



"Barcelona"

Quad-Core AMD Opteron™ processor
 Family 10h

Nested Paging

IBS

OSVW

1 GB Paging

SSE4A

POPCNT

LZCNT

MONITOR/MWAIT*

*SVM Intercepts

Policy and Mechanism

Hypervisor Manager sets migration policy based on a datacenter policy

- Implementation is product-independent
- Hypervisor has knowledge of all systems in its migration pool
- VM attributes go with the VM when it migrates

Hypervisor uses attributes to reprogram the new target machine with the feature set expected by the guest

- MSR C001_1004 and C001_1005
- Override full feature set reported
- Read CPU's default values and mask out bits for compatibility

Example setting a Family 10h to Family 0Fh rev F

```
/*
 * Setup the EDX and EAX for MSR write to put new
 * values for CPUID Fn 0000_0001 (ECX and EDX)
 */
mov ecx, 0xC0011004
mov edx, 0x00002001
mov eax, 0x178FBFBF
wrmsr
/*
 * Setup the EDX and EAX for MSR write to put new
 * values for CPUID Fn 8000_0001 (ECX and EDX)
 */
mov ecx, 0xC0011005
mov edx, 0x0000001F
mov eax, 0xEBD3FBFF
wrmsr
```

What's next?

Know your guests and applications the customers want to run

- Are you consolidating apps from older hardware?
- Do they have requirements on new hardware features

Make the right server pools

Put options in the your configuration software to enable safe migration

Use AMD-V™ Extended Migration to combine virtual machines with similar requirements together

Links

White paper posted on developer.amd.com

<http://developer.amd.com/assets/Live%20Virtual%20Machine%20Migration%20on%20AMD%20processors.pdf>

Stepping for cpus available at <http://www.amdcompare.com>

Thank You !

Trademark Attribution

AMD, the AMD Arrow logo, AMD Opteron, AMD-V and combinations thereof are trademarks of Advanced Micro Devices, Inc. in the United States and/or other jurisdictions. Other names used in this presentation are for identification purposes only and may be trademarks of their respective owners.

©2007 Advanced Micro Devices, Inc. All rights reserved.