



AMD-K6[®]

Processor

BIOS Design

Application Note

Publication # 21329 Rev: L Amendment/0
Issue Date: December 1999

© 1999 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. ("AMD") products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel or otherwise, to any intellectual property rights is granted by this publication. Except as set forth in AMD's Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD's products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD's product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD logo, K6, 3DNow!, and combinations thereof, K86, and AMD-K5 are trademarks, and AMD-K6 is a registered trademark of Advanced Micro Devices, Inc.

MMX is a trademark and Pentium is a registered trademark of Intel Corporation.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Contents

Contents	iii
List of Figures	v
List of Tables	vii
Revision History	ix
Audience	1
Processor Models and Steppings	2
BIOS Consideration Checklist	3
AMD-K6 [®] Processor Models 6, 7 and AMD-K6-2 Processor Model 8/[7:0]	3
CUID	3
CPU Speed Detection	3
Model-Specific Registers (MSRs)	3
Cache Testing	3
SMM Issues	4
AMD-K6 [®] -2 Processor Model 8/[F:8]	4
CUID	4
CPU Speed Detection	4
Model-Specific Registers (MSRs)	4
Cache Testing	5
SMM Issues	5
AMD-K6 [®] -III Processor Model 9	5
CUID	5
CPU Speed Detection	5
Model-Specific Registers (MSRs)	6
Cache Testing	6
SMM Issues	6
Register States After RESET and INIT	7
State of the Processor After INIT	8
CUID Identification Algorithms	9
Built-In Self-Test (BIST)	13

System Management Mode (SMM)	13
State-Save Map Differences	13
I/O Trap Dword Differences	13
Model-Specific Registers (MSRs)	14
Standard MSRs	15
Machine-Check Address Register (MCAR) and Machine-Check Type Register (MCTR)	15
Test Register 12 (TR12)	15
Time Stamp Counter (TSC)	15
AMD-K6 [®] Processor Models 6, 7 and AMD-K6-2 Processor Model 8/[7:0]	16
Extended Feature Enable Register (EFER)	16
Write Handling Control Register (WHCR)	17
SYSCALL/SYSRET Target Address Register (STAR)	19
AMD-K6 [®] -2 Processor Model 8/[F:8]	20
Extended Feature Enable Register (EFER)	20
Write Handling Control Register (WHCR)	23
SYSCALL/SYSRET Target Address Register (STAR)	26
UC/WC Cacheability Control Register (UWCCR)	26
Processor State Observability Register (PSOR)	30
Page Flush/Invalidate Register (PFIR)	31
AMD-K6 [®] -III Processor Model 9	33
Extended Feature Enable Register (EFER)	33
Write Handling Control Register (WHCR)	34
SYSCALL/SYSRET Target Address Register (STAR)	35
UC/WC Cacheability Control Register (UWCCR)	35
Processor State Observability Register (PSOR)	35
Page Flush/Invalidate Register (PFIR)	36
Level-2 Cache Array Access Register (L2AAR)	36
New AMD-K6 [®] Processor Instructions	41
Additional Considerations	41
Software Timing Dependencies Relative to Memory Controller Setup	41
Pipelining Support	42
Read-Only Memory	42

List of Figures

Figure 1.	CPUID Instruction Flow Chart	12
Figure 2.	Extended Feature Enable Register (EFER) —MSR C000_0080h (Models 6, 7, and 8/[7:0])	16
Figure 3.	Write Handling Control Register (WHCR) —MSR C000_0082h (Models 6, 7, and 8/[7:0])	18
Figure 4.	SYSCALL/SYSRET Target Address Register (STAR) —MSR C000_0081h (Models 8 and 9)	19
Figure 5.	Extended Feature Enable Register (EFER) —MSR C000_0080h (Model 8/[F:8])	21
Figure 6.	Write Handling Control Register (WHCR) —MSR C000_0082h (Model 8/[F:8])	24
Figure 7.	UC/WC Cacheability Control Register (UWCCR))—MSR C000_0085h (Model 8/[F:8])	27
Figure 8.	Processor State Observability Register (PSOR) —MSR C000_0087h (Model 8/[F:8])	30
Figure 9.	Page Flush/Invalidate Register (PFIR) —MSR C000_0088h (Model 8/[F:8])	31
Figure 10.	Extended Feature Enable Register (EFER) —MSR C000_0080h (Model 9)	34
Figure 11.	Processor State Observability Register (PSOR) —MSR C000_0087h (Model 9)	35
Figure 12.	L2 Cache Organization.	37
Figure 13.	L2 Cache Sector and Line Organization	37
Figure 14.	L2 Tag or Data Location - EDX.	38
Figure 15.	L2 Data - EAX.	39
Figure 16.	L2 Tag Information - EAX.	40
Figure 17.	LRU Byte.	40

List of Tables

Table 1.	State of the AMD-K6 [®] Processor Models 6, 7, and 8/[7:0] After RESET	7
Table 2.	State of the AMD-K6 [®] -2 Processor Model 8/[F:8] After RESET	7
Table 3.	State of the AMD-K6 [®] -III Processor Model 9 After RESET	8
Table 4.	Summary of AMD-K6 [®] Processor Models and BIOS Boot String	10
Table 5.	AMD-K6 [®] Processor I/O Trap Dword Configuration at Offset FFA4h	13
Table 6.	Summary of MSR Differences Within the AMD-K6 [®] Family	14
Table 7.	Extended Feature Enable Register (EFER) Definition (Models 6, 7, and 8/[7:0])	17
Table 8.	SYSCALL/SYSRET Target Address Register (STAR) Definition (Models 8 and 9).	20
Table 9.	Extended Feature Enable Register (EFER) Definition (Model 8/[F:8])	21
Table 10.	EWBEC Settings	23
Table 11.	WC/UC Memory Type	28
Table 12.	Valid Masks and Range Sizes	28
Table 13.	Processor-to-Bus Clock Ratios	31
Table 14.	Extended Feature Enable Register (EFER) Definition (Model 9)	34
Table 15.	Tag versus Data Selector	38

Revision History

Date	Rev	Description
Dec 1997	G	Added boot string information to the CPUID description on page 4.
Dec 1997	G	Added processor speeds, bus speeds, and boot string information to Table 4 on page 10.
Dec 1997	G	Added 3DNow!™ to “New AMD-K6® Processor Instructions” on page 41.
May 1998	H	Deleted AMD-K5™ processor information. Added URL to BIOS information for the AMD-K5 processor on page 1.
Nov 1998	H1	Added “Processor Models and Steppings” on page 2.
Nov 1998	H1	Added “AMD-K6®-2 Processor Model 8/[F:8]” on page 4 to the BIOS Consideration Checklist.
Nov 1998	H1	Added Table 2, “State of the AMD-K6®-2 Processor Model 8/[F:8] After RESET,” on page 7.
Nov 1998	H1	Added Table 6, “Summary of MSR Differences Within the AMD-K6® Family,” on page 14.
Nov 1998	H1	Moved paragraphs describing four MSRs into “Standard MSRs” on page 15.
Nov 1998	H1	Added “AMD-K6® Processor Models 6, 7 and AMD-K6-2 Processor Model 8/[7:0]” on page 16 and “AMD-K6®-2 Processor Model 8/[F:8]” on page 20 to the Model-Specific Register (MSRs) section.
Nov 1998	H1	Added “Pipelining Support” and “Read-Only Memory” on page 42.
July 1998	I	Released under a nondisclosure agreement.
Feb 1999	J	Added AMD-K6-III processor information to “Processor Models and Steppings” on page 2 and to “BIOS Consideration Checklist” on page 5.
Feb 1999	J	Added Table 3, “State of the AMD-K6®-III Processor Model 9 After RESET,” on page 8.
Feb 1999	J	Added to “CPUID Identification Algorithms” on page 9 a recommendation to change “L2” to “External” within the Summary/Configuration screen.
Feb 1999	J	Added new boot strings to Table 4, “Summary of AMD-K6® Processor Models and BIOS Boot String,” on page 10.
Feb 1999	J	Added extended function 8000_0006h to Figure 1, “CPUID Instruction Flow Chart,” on page 12.
Feb 1999	J	Added L2 unified cache to “Built-In Self-Test (BIST)” on page 13.
Feb 1999	J	Added AMD-K6-III processor information to Table 6, “Summary of MSR Differences Within the AMD-K6® Family,” on page 14.
Feb 1999	J	Added “AMD-K6®-III Processor Model 9” on page 33.
Sept 1999	K	Added note 2 to Table 1, “State of the AMD-K6® Processor Models 6, 7, and 8/[7:0] After RESET,” on page 7.
Sept 1999	K	Revised Figure 1, “CPUID Instruction Flow Chart,” on page 12.
Dec 1999	L	Added AMD-K6-2/533 to Table 4, “Summary of AMD-K6® Processor Models and BIOS Boot String,” on page 10.

Application Note

AMD-K6[®] **Processor** **BIOS Design**

This document highlights the BIOS modifications required to fully support the AMD-K6[®] processors. This document is a supplement to the *AMD K86™ Family BIOS and Software Tools Developers Guide*, order# 21062. Unless otherwise noted, the information in this application note pertains to all processors in the AMD-K6 family, which includes the AMD-K6 processor Models 6 and 7, the AMD-K6-2 processor Model 8, and the AMD-K6-III processor Model 9.

There can be more than one way to implement the functionality detailed in this document, and the information provided is for demonstration purposes.

All referenced AMD-K6 processor documents can be found on the AMD website at <http://www.amd.com/K6/k6docs>. For BIOS information on the AMD-K5™ processor, refer to the *AMD K86™ Family BIOS Design Application Note* that can be found at <http://www.amd.com/products/cpg/techdocs/archive.html>.

Audience

It is assumed that the reader possesses the proper knowledge of the K86 processors, the x86 architecture, and programming requirements to understand the information presented in this document.

Processor Models and Steppings

There are four models within the AMD-K6 family of processors—Models 6, 7, 8, and 9:

- **Model 6 (AMD-K6 processor Model 6)**—This processor is manufactured in the 0.35-micron process. The Model 6 supports six Model-Specific Registers (MSRs).
- **Model 7 (AMD-K6 processor Model 7)**—This is the first processor manufactured in the 0.25-micron process. The Model 7 implements the same six MSRs as the Model 6, and the bits and fields within these six MSRs are defined identically.
- **Model 8 (AMD-K6-2 processor Model 8)**—This processor is also manufactured in the 0.25-micron process. Some important features supported by the Model 8 include the 3DNow![™] instruction set and support for a 100-MHz processor bus. This document covers the following two specific stepping ranges of the Model 8:
 - **Model 8/[7:0]** is any of eight possible model/steppings—Models 8/0, 8/1, 8/2, 8/3, 8/4, 8/5, 8/6, or 8/7. Model 8/[7:0] implements the same six MSRs as the Models 6 and 7, and the bits and fields within these six MSRs are defined identically. Model 8/[7:0] also implements the SYSCALL/SYSRET Target Address Register (STAR) MSR for a total of seven MSRs.
 - **Model 8/[F:8]** is any of eight possible model/steppings—Models 8/8, 8/9, 8/A, 8/B, 8/C, 8/D, 8/E, or 8/F. Model 8/[F:8] implements the same six MSRs as the Models 6, 7, and 8/[7:0], but the bits and fields within two of these MSRs are not defined identically. Also, Model 8/[F:8] supports the STAR MSR and three additional MSRs, for a total of ten MSRs.
- **Model 9 (AMD-K6-III processor Model 9)**—This processor is also manufactured in the 0.25-micron process. In addition to supporting the 3DNow! instruction set and a 100-MHz processor bus, the Model 9 contains a 256-Kbyte backside L2 cache. The Model 9 implements the same six MSRs as the Models 6, 7, and 8/[7:0], but the bits and fields within two of these MSRs are not defined identically. Also, the Model 9 supports the STAR MSR and four additional MSRs, for a total of eleven MSRs.

Table 6 on page 14 summarizes the MSR differences between the models and steppings of the AMD-K6 family of processors.

BIOS Consideration Checklist

AMD-K6[®] Processor Models 6, 7 and AMD-K6-2 Processor Model 8/[7:0]

The term *processor* in this section is defined as the AMD-K6 processor Models 6 and 7, and the AMD-K6-2 processor Model 8/[7:0].

CPUID

- Use the CPUID instruction to properly identify the processor. For information on the CPUID instruction, refer to the *AMD Processor Recognition Application Note*, order# 20734.
- Determine the processor model, stepping, and features using functions 0000_0001h and 8000_0001h of the CPUID instruction.
- Display BIOS boot strings: *AMD-K6(tm)/XXX* for Models 6 and 7, and *AMD-K6(tm)-2/XXX* for all steppings of the Model 8. For more information, see “CPUID Identification Algorithms” on page 9.

CPU Speed Detection

- Use speed detection algorithms that do not rely on repetitive instruction sequences.
- Use the Time Stamp Counter (TSC) to ‘clock’ a timed operation and compare the result to the Real Time Clock (RTC) to determine the operating frequency. See the *CPU Speed Determination Program* available on the AMD website at <http://www.amd.com/K6/k6docs/>.
- Display the recommended BIOS boot string as shown in Table 4, “Summary of AMD-K6[®] Processor Models and BIOS Boot String,” on page 10.

Model-Specific Registers (MSRs)

- Only access MSRs implemented in the processor.
- Enable Write Allocation by programming the Write Handling Control Register (WHCR). See “Write Handling Control Register (WHCR)” on page 17 and the *Implementation of Write Allocate in the K86™ Processors Application Note*, order# 21326 for more information.

Cache Testing

- The processor does not contain MSRs to allow for testing of the L1 cache.

SMM Issues

- The System Management Mode (SMM) functionality of the processor is the same as the Pentium[®] processor.
- Implement the processor SMM state-save area in a similar manner as Pentium processors except for the IDT Base and possibly Pentium processor-reserved areas. See “System Management Mode (SMM)” on page 13 for more information.

AMD-K6[®]-2 Processor Model 8/[F:8]

The term *processor* in this section is defined as the AMD-K6-2 processor Model 8/[F:8].

CPUID

- Use the CPUID instruction to properly identify the processor. For information on the CPUID instruction, refer to the *AMD Processor Recognition Application Note*, order# 20734.
- Determine the processor model, stepping and features using functions 0000_0001h and 8000_0001h of the CPUID instruction.
- Display BIOS boot string: *AMD-K6(tm)-2/XXX* for all steppings of the Model 8. For more information, see “CPUID Identification Algorithms” on page 9.

CPU Speed Detection

- Use speed detection algorithms that do not rely on repetitive instruction sequences.
- Use the Time Stamp Counter (TSC) to ‘clock’ a timed operation and compare the result to the Real Time Clock (RTC) to determine the operating frequency. See the *CPU Speed Determination Program* available on the AMD website at <http://www.amd.com/K6/k6docs/>.
- Display the recommended BIOS boot string as shown in Table 4, “Summary of AMD-K6[®] Processor Models and BIOS Boot String,” on page 10.

Model-Specific Registers (MSRs)

- Only access MSRs implemented in the processor.
- Enable Write Allocation by programming the Write Handling Control Register (WHCR). See “Write Handling Control Register (WHCR)” on page 23 and the *Implementation of Write Allocate in the K86™ Processors Application Note*, order# 21326 for more information.

Note: The WHCR register as defined in the Model 6, Model 7, and Model 8/[7:0] has changed in the Model 8/[F:8].

- Utilize the information provided in the Processor State Observability Register (PSOR) to display the correct processor bus frequency.

Cache Testing

- The processor does not contain MSRs to allow for testing of the L1 cache.

SMM Issues

- The System Management Mode (SMM) functionality of the processor is the same as the Pentium processor.
- Implement the processor SMM state-save area in a similar manner as Pentium processors except for the IDT Base and possibly Pentium processor-reserved areas. See “System Management Mode (SMM)” on page 13 for more information.

AMD-K6[®]-III Processor Model 9

The term *processor* in this section is defined as the AMD-K6-III processor Model 9.

CPUID

- Use the CPUID instruction to properly identify the processor. For information on the CPUID instruction, refer to the *AMD Processor Recognition Application Note*, order# 20734.
- Determine the processor model, stepping and features using functions 0000_0001h and 8000_0001h of the CPUID instruction.
- Display BIOS boot string: *AMD-K6(tm)-III/XXX* for all steppings of the Model 9. For more information, see “CPUID Identification Algorithms” on page 9.

CPU Speed Detection

- Use speed detection algorithms that do not rely on repetitive instruction sequences.
- Use the Time Stamp Counter (TSC) to ‘clock’ a timed operation and compare the result to the Real Time Clock (RTC) to determine the operating frequency. See the *CPU Speed Determination Program* available on the AMD website at <http://www.amd.com/K6/k6docs/>.
- Display the recommended BIOS boot string as shown in Table 4, “Summary of AMD-K6[®] Processor Models and BIOS Boot String,” on page 10.

Model-Specific Registers (MSRs)

- Only access MSRs implemented in the processor.
- Enable Write Allocation by programming the Write Handling Control Register (WHCR). See “Write Handling Control Register (WHCR)” on page 34 and the *Implementation of Write Allocate in the K86™ Processors Application Note*, order# 21326 for more information.

Note: The WHCR register as defined in the Model 6, Model 7, and Model 8/[7:0] has changed in the Model 9.

- Utilize the information provided in the Processor State Observability Register (PSOR) to display the correct processor bus frequency.

Cache Testing

- The AMD-K6-III processor does not contain MSRs to allow for testing of the L1 cache. However, the AMD-K6-III processor does contain a MSR that allows for testing of its L2 cache. This MSR is called L2AAR and it is described in “Level-2 Cache Array Access Register (L2AAR)” on page 36.

SMM Issues

- The System Management Mode (SMM) functionality of the processor is the same as the Pentium processor.
- Implement the processor SMM state-save area in a similar manner as Pentium processors except for the IDT Base and possibly Pentium processor-reserved areas. See “System Management Mode (SMM)” on page 13 for more information.

Register States After RESET and INIT

After the processor has completed its initialization following the recognition of an asserted RESET or INIT signal, the states of all architecture registers and MSRs are compatible with those of Pentium processors. Differences are listed in Table 1 through Table 3.

Table 1. State of the AMD-K6[®] Processor Models 6, 7, and 8/[7:0] After RESET

Register	RESET State	Notes
EDX	0000_05MSh	1
EFER	0000_0000_0000_0000h	
STAR	0000_0000_0000_0000h	2
WHCR	0000_0000_0000_0000h	
Notes:		
<ol style="list-style-type: none"> 1. "M" represents the Model and "S" represents the Stepping. 2. AMD-K6 processor Models 6 and 7 do not support STAR. 		

Table 2. State of the AMD-K6[®]-2 Processor Model 8/[F:8] After RESET

Register	RESET State	Notes
EDX	0000_058Sh	1
EFER	0000_0000_0000_0002h	
PFIR	0000_0000_0000_0000h	
PSOR	0000_0000_0000_01SBh	1, 2
STAR	0000_0000_0000_0000h	
UWCCR	0000_0000_0000_0000h	
WHCR	0000_0000_0000_0000h	
Notes:		
<ol style="list-style-type: none"> 1. "S" represents the Stepping. 2. "B" represents PSOR[3:0], where PSOR[3] equals 0, and PSOR[2:0] is equal to the value of the BF[2:0] signals sampled during the falling transition of RESET. 		

Table 3. State of the AMD-K6[®]-III Processor Model 9 After RESET

Register	RESET State	Notes
EDX	0000_059Sh	1
EFER	0000_0000_0000_0002h	2
L2AAR	0000_0000_0000_0000h	
PFIR	0000_0000_0000_0000h	
PSOR	0000_0000_0000_00SBh	1, 3
STAR	0000_0000_0000_0000h	
UWCCR	0000_0000_0000_0000h	
WHCR	0000_0000_0000_0000h	
Notes: <ol style="list-style-type: none"> 1. "S" represents the Stepping. 2. Because EFER[4] equals 0 after RESET, the L2 cache is enabled by default after RESET. 3. "B" represents PSOR[3:0], where PSOR[3] equals 0, and PSOR[2:0] is equal to the value of the BF[2:0] signals sampled during the falling transition of RESET. 		

State of the Processor After INIT

The assertion of INIT causes the processor to empty its pipelines, initialize most of its internal state, and branch to address FFFF_FFF0h—the same instruction execution starting point used after RESET. Unlike RESET, the processor preserves the contents of its caches, the floating-point state, the SMM base, MSRs, and the CD and NW bits of the CR0 register.

The edge-sensitive interrupts FLUSH# and SMI# are sampled and preserved during the INIT process and are handled accordingly after the initialization is complete. However, the processor resets any pending NMI interrupt upon sampling INIT asserted.

INIT can be used as an accelerator for 80286 code that requires a reset to exit from Protected mode back to Real mode.

CPUID Identification Algorithms

The CPUID instruction provides information about the processor (vendor, type, name, etc.) and its capabilities (features). After detecting the processor and its capabilities, software can be accurately tuned to the system for maximum performance and benefit to users. For example, game software can test the performance level available from a particular processor by detecting the type of processor. If the performance level is high enough, the software can enable additional capabilities or more advanced algorithms. Another example involves testing if the processor supports 3DNow! instructions. If the software finds this functionality present when it checks the feature bits, it can utilize these more powerful extensions for better performance on new multimedia software.

For more detailed information, refer to the *AMD Processor Recognition Application Note*, order# 20734.

Table 4 on page 10 shows the recommended BIOS boot strings for the AMD-K6 processor. The recommended boot strings are:

- *AMD-K6(tm)/XXX* for Models 6 and 7
- *AMD-K6(tm)-2/XXX* for Model 8 (all steppings)
- *AMD-K6(tm)-III/XXX* for Model 9

The value for XXX is determined by calculating the core frequency of the processor. Use the Time Stamp Counter (TSC) to ‘clock’ a timed operation and compare the result to the Real Time Clock (RTC) to determine the operating frequency.

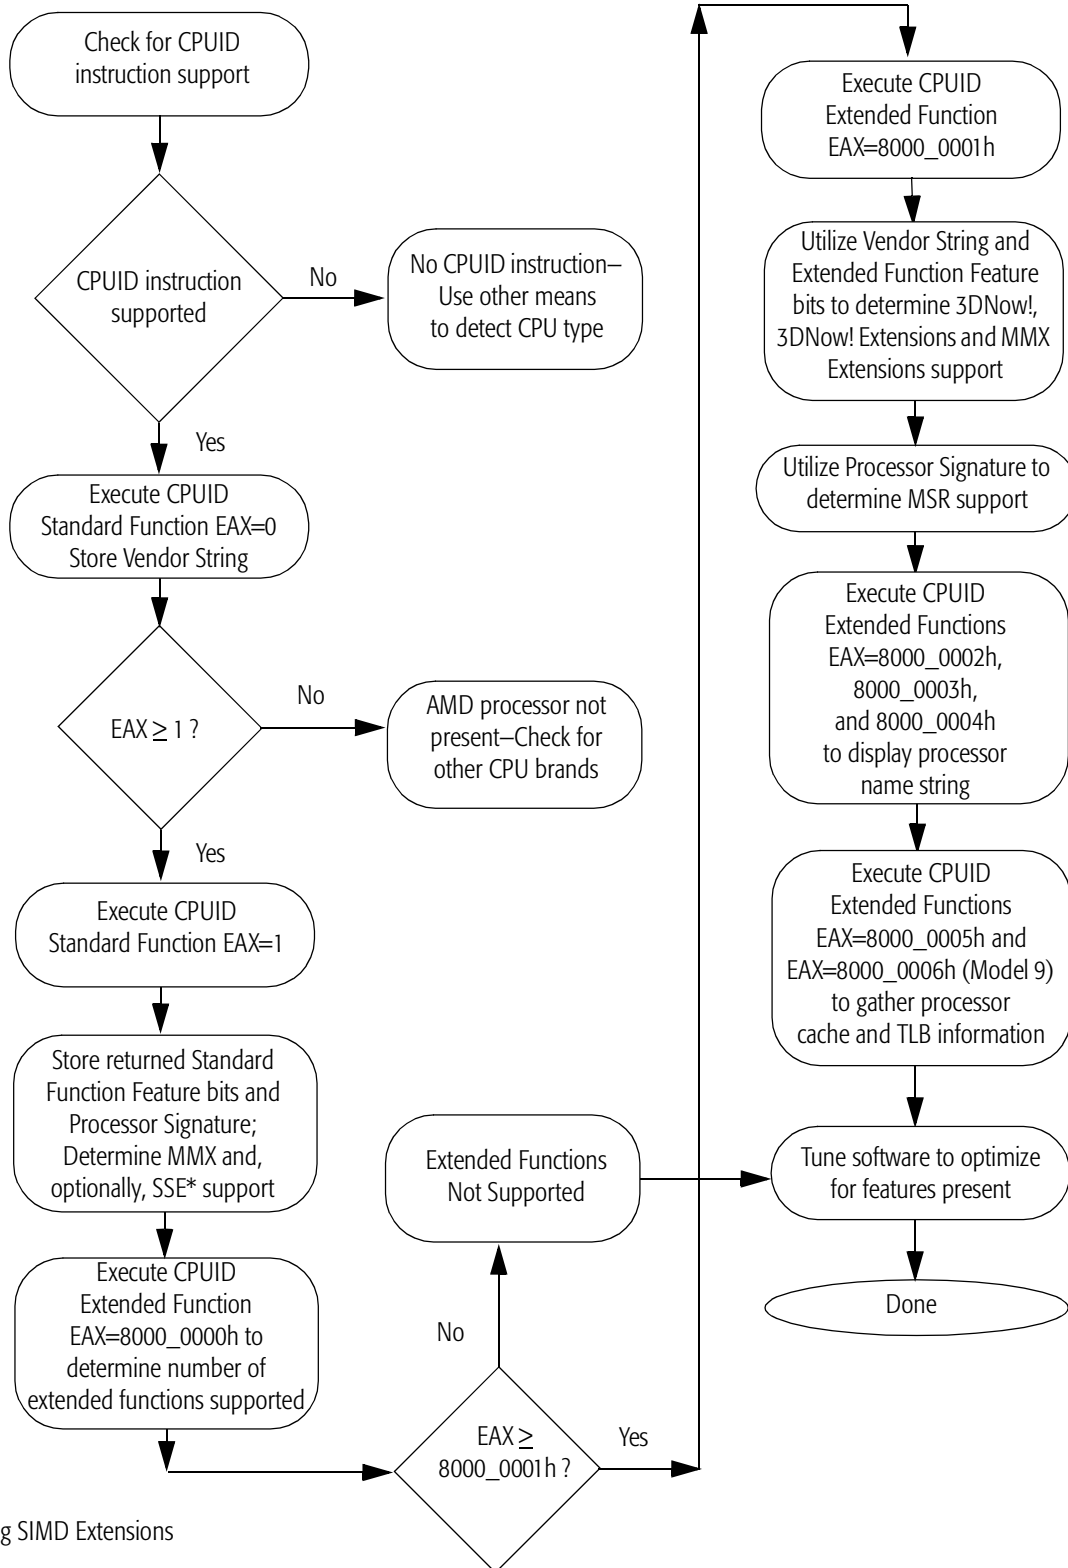
In addition to displaying the recommended boot string, BIOS code that normally indicates the presence of a L2 cache within the Summary/Configuration screen should change all occurrences of “L2” to “External.”

Note: *Table 4 contains information intended to prepare the infrastructure for potential future products. These products at these speed grades may or may not be announced, but BIOS software should be prepared to support these options.*

Table 4. Summary of AMD-K6[®] Processor Models and BIOS Boot String

Instruction Family	Model	CPU Speed (MHz)	CPU Bus Speed (MHz)	Recommended BIOS Boot-String Display
5 (AMD-K6 [®] Processor)	6	166	66	AMD-K6(tm)/166
		200	66	AMD-K6(tm)/200
		233	66	AMD-K6(tm)/233
	7	200	66	AMD-K6(tm)/200
		233	66	AMD-K6(tm)/233
		266	66	AMD-K6(tm)/266
		300	66	AMD-K6(tm)/300
	8	233	66	AMD-K6(tm)-2/233
		266	66	AMD-K6(tm)-2/266
		300	66	AMD-K6(tm)-2/300
		333	66	AMD-K6(tm)-2/333
		366	66	AMD-K6(tm)-2/366
		400	66	AMD-K6(tm)-2/400
		333	95	AMD-K6(tm)-2/333
		380	95	AMD-K6(tm)-2/380
		300	100	AMD-K6(tm)-2/300
		350	100	AMD-K6(tm)-2/350
		400	100	AMD-K6(tm)-2/400
		450	100	AMD-K6(tm)-2/450
		475	95	AMD-K6(tm)-2/475
		500	100	AMD-K6(tm)-2/500
		533	97	AMD-K6(tm)-2/533
		550	100	AMD-K6(tm)-2/550
		600	100	AMD-K6(tm)-2/600
	9	350	100	AMD-K6(tm)-III/350
		400	100	AMD-K6(tm)-III/400
		450	100	AMD-K6(tm)-III/450
		475	95	AMD-K6(tm)-III/475
		500	100	AMD-K6(tm)-III/500
		550	100	AMD-K6(tm)-III/550
		600	100	AMD-K6(tm)-III/600

Figure 1 shows a flow chart for the CPUID instruction. Use this chart to implement a CPUID algorithm.



* Streaming SIMD Extensions

Figure 1. CPUID Instruction Flow Chart

Built-In Self-Test (BIST)

For all models of the AMD-K6 processor, BIST is run unconditionally following the falling transition of RESET. The results of the test are contained in the general purpose register EAX. If EAX contains 0000_0000h, then BIST was successful. If the contents of EAX are non-zero, the BIST failed. The internal resources tested during BIST include the following:

- L1 instruction and data caches
- L2 unified cache (Model 9 only)
- Instruction and Data Translation Lookaside Buffers (TLBs)

System Management Mode (SMM)

This section documents the SMM differences between specified models of the AMD-K6 processor and the Pentium processor. For more information on SMM implementation in the K86 processors, see the *AMD K86™ Family BIOS and Software Tools Developers Guide*, order# 21062.

State-Save Map Differences

The SMM implemented in the AMD-K6 processor differs from the SMM implemented in the Pentium processor in one way. The IDT Base location in the AMD-K6 processor is located at offset FF90h. Pentium has the IDT Base located at offset FF94h.

I/O Trap Dword Differences

The I/O trap dword is located at offset FFA4h. This state-save area, which is reserved in Pentium, contains information regarding an I/O instruction that may have been trapped by an SMI# assertion.

Table 5. AMD-K6[®] Processor I/O Trap Dword Configuration at Offset FFA4h

Bits 31–16	Bits 15–4	Bit 3	Bit 2	Bit 1	Bit 0
I/O Port Address	Reserved	Rep String Operation	I/O String Operation	Valid I/O Instruction	Input or Output

Model-Specific Registers (MSRs)

All models and steppings of the AMD-K6 processor family support the following four standard MSRs, and the bits and fields within each of these MSRs are defined identically:

- Machine-Check Address Register (MCAR)—ECX = 00h
- Machine-Check Type Register (MCTR)—ECX = 01h
- Test Register 12 (TR12)—ECX = 0Eh
- Time Stamp Counter (TSC)—ECX = 10h

All models and steppings of the AMD-K6 processor family also support the following two MSRs, but the fields within each of these MSRs are defined differently as shown in Table 6 (an ‘X’ indicates support for a register or field):

- Extended Feature Enable Register (EFER)—ECX = C000_0080h
- Write Handling Control Register (WHCR)—ECX = C000_0082h

Table 6. Summary of MSR Differences Within the AMD-K6[®] Family

Model	Stepping	Standard ¹ MSRs	EFER ²				WHCR ³		UWCCR	PSOR	PFIR	L2AAR
			L2D	EWBEC	DPE	SCE	508 MB	4092 MB				
6	All	X				X	X					
7	All	X				X	X					
8	7:0	X				X	X					
8	F:8	X		X	X	X		X	X	X	X	
9	All	X	X	X	X	X		X	X	X	X	X

Notes:

1. There are four MSRs that every model and stepping of the AMD-K6 family of processors support identically—MCAR, MCTR, TR12, and TSC.
2. L2D, EWBEC, and DPE are bits/fields supported in EFER for the indicated models/steppings. All models/steppings support the System Call Extension (SCE) bit in EFER, even if the corresponding SYSCALL and SYSRET instructions and the STAR register are not supported.
3. Indicates whether the WAELIM field supports 508 Mbytes or 4092 Mbytes of memory. The location of the WAE15M bit and the WAELIM field within the WHCR register differs between the models/steppings that support 508 Mbytes of memory and those that support 4092 Mbytes of memory.

Standard MSRs

This section describes the four standard MSRs that every model and stepping of the AMD-K6 family of processors support identically.

Machine-Check Address Register (MCAR) and Machine-Check Type Register (MCTR)

The processor does not support the generation of a machine check exception, but does provide a 64-bit Machine Check Address Register (MCAR) and a 64-bit Machine Check Type Register (MCTR) for software compatibility. Because the processor does not support machine check exceptions, the contents of the MCAR and MCTR are only affected by the WRMSR instruction and by RESET being sampled asserted (where all bits in each register are reset to 0).

The processor also provides the Machine Check Exception (MCE) bit in Control Register 4 (CR4, bit 6) as a read-write bit. However, the state of this bit has no effect on the operation of the processor.

Test Register 12 (TR12)

The processor provides the 64-bit Test Register 12 (TR12), but only the Cache Inhibit (CI) bit (bit 3 of TR12) is supported. All other bits in TR12 have no effect on the processor's operation. The I/O Trap Restart function (bit 9 of TR12) is always enabled on the AMD-K6.

Time Stamp Counter (TSC)

With each processor clock cycle, the processor increments a 64-bit time stamp counter (TSC) MSR. The counter can be written or read using the WRMSR or RDMSR instructions when the ECX register contains the value 10h and CPL = 0. The counter can also be read using the RDTSC instruction, but the required privilege level for this instruction is determined by the Time Stamp Disable (TSD) bit in CR4. With either of these instructions, the EDX and EAX registers hold the upper and lower dwords of the 64-bit value to be written to or read from the TSC, as follows:

- *EDX*—Upper 32 bits of TSC
- *EAX*—Lower 32 bits of TSC

The TSC can be loaded with any arbitrary value. This feature is compatible with the Pentium processor.

AMD-K6[®] Processor Models 6, 7 and AMD-K6-2 Processor Model 8/[7:0]

The AMD-K6 processor Models 6 and 7 and the AMD-K6-2 processor Model 8/[7:0] provide the following MSR. The first four MSRs are described in “Standard MSRs” on page 15. The contents of ECX selects the MSR to be addressed by the RDMSR and WRMSR instruction.

- Machine-Check Address Register (MCAR)—ECX = 00h
- Machine-Check Type Register (MCTR)—ECX = 01h
- Test Register 12 (TR12)—ECX = 0Eh
- Time Stamp Counter (TSC)—ECX = 10h
- Extended Feature Enable Register (EFER)—ECX = C000_0080h
- Write Handling Control Register (WHCR)—ECX = C000_0082h

All steppings (F:0) of the AMD-K6-2 processor Model 8 support the following MSR, and the bits and fields within this MSR are defined identically:

- SYSCALL/SYSRET Target Address Register (STAR)—ECX = C000_0081h

Extended Feature Enable Register (EFER)

The Extended Feature Enable Register (EFER) contains the control bits that enable the extended features of the AMD-K6 processor. Figure 2 shows the format of the EFER register, and Table 7 defines the function of each bit of the EFER register. The EFER register is MSR C000_0080h.




Figure 2. Extended Feature Enable Register (EFER)—MSR C000_0080h (Models 6, 7, and 8/[7:0])

Table 7. Extended Feature Enable Register (EFER) Definition (Models 6, 7, and 8/[7:0])

Bit	Description	R/W	Function
63–1	Reserved	R	Writing a 1 to any reserved bit causes a general protection fault to occur. All reserved bits are always read as 0.
0	System Call Extension (SCE)	R/W	SCE must be set to 1 to enable the usage of the SYSCALL and SYSRET instructions.

Note: *The AMD-K6 processor Models 6 and 7 provide the SCE bit in the EFER register, but this bit does not affect processor operation because the SYSCALL and SYSRET instructions and the STAR register are not supported in these models.*

Write Handling Control Register (WHCR)

The processor contains a split level-1 (L1) 64-Kbyte writeback cache organized as a separate 32-Kbyte instruction cache and a 32-Kbyte data cache with two-way set associativity. The cache line size is 32 bytes and lines are read from memory using an efficient pipelined burst read cycle. Further performance gains are achieved by the implementation of a write allocation scheme.

A write allocate, if enabled, occurs when the processor has a pending memory write cycle to a cacheable line and the line does not currently reside in the L1 cache. For more information on write allocate, see the *Implementation of Write Allocate in the K86™ Processors Application Note*, order# 21326, and the Cache Organization section of the *AMD-K6[®] Processor Data Sheet*, order# 20695 or the *AMD-K6[®]-2 Processor Data Sheet*, order# 21850.

This section describes two programmable mechanisms used by the processor to determine when to perform write allocate. When either of these mechanisms indicates that a pending write is to a cacheable area of memory, a write allocate is performed.

Before enabling write allocate or changing memory cacheability/writeability, the BIOS must writeback and invalidate the internal cache by using the WBINVD instruction. In addition, write allocate should be enabled only after performing any memory sizing or typing algorithms.

The Write Handling Control Register (WHCR) is a MSR that contains three fields—the WCDE bit, the Write Allocate Enable Limit (WAELIM) field, and the Write Allocate Enable 15-to-16-Mbyte (WAE15M) bit (see Figure 3).




Figure 3. Write Handling Control Register (WHCR)—MSR C000_0082h (Models 6, 7, and 8/[7:0])

WCDE. For proper functionality, always program bit 8 of WHCR to 0. See “Pipelining Support” on page 42 for more information on WCDE.

Write Allocate Enable Limit. The WAE15M field is 7 bits wide. This field, multiplied by 4 Mbytes, defines an upper memory limit. Any pending write cycle that misses the L1 cache and that addresses memory below this limit causes the processor to perform a write allocate (assuming the address is not within a range where write allocates are disallowed). Write allocate is disabled for memory accesses at and above this limit unless the processor determines a pending write cycle is cacheable by means of one of the other write allocate mechanisms—“Write to a Cacheable Page” and “Write to a Sector” (for more information, see the “Cache Organization” chapter in the AMD-K6[®] Processor Data Sheet, order# 20695 or the AMD-K6[®]-2 Processor Data Sheet, order# 21850). The maximum value of this limit is $((2^7-1) \cdot 4 \text{ Mbytes}) = 508 \text{ Mbytes}$. When all the bits in this field are set to 0, all memory is above this limit and the write allocate mechanism is disabled (even if all bits in the WAE15M field are set to 0, write allocates can still occur due to the “Write to a Cacheable Page” and “Write to a Sector” mechanisms).

Once the BIOS determines the amount of RAM installed in the system, this number should also be used to program the WAE15M field. For example, a system with 32 Mbytes of RAM would program the WAE15M field with the value 0001000b. This value (8), when multiplied by 4 Mbytes, yields 32 Mbytes as the write allocate limit.

Write Allocate Enable 15-to-16-Mbyte. The WAE15M bit is used to enable write allocations for the memory write cycles that address the 1 Mbyte of memory between 15 Mbytes and 16 Mbytes. This bit must be set to 1 to allow write allocates in this memory area. This sub-mechanism of the WAELIM provides a memory hole to prevent write allocates. This memory hole is provided to account for a small number of uncommon memory-mapped I/O adapters that use this particular memory address space. If the system contains one of these peripherals, the bit should be set to 0 (even if the WAE15M bit is set to 0, write allocates can still occur between 15 Mbytes and 16 Mbytes due to the “Write to a Cacheable Page” and “Write to a Sector” mechanisms). The WAE15M bit is ignored if the value in the WAELIM field is set to less than 16 Mbytes.

By definition, write allocations are not performed in the memory area between 640 Kbytes and 1 Mbyte unless the processor determines a pending write cycle is cacheable by means of “Write to a Cacheable Page” or “Write to a Sector.” It is not safe to perform write allocations between 640 Kbytes and 1 Mbyte (000A_0000h to 000F_FFFFh) because it is considered a noncacheable region of memory.

SYSCALL/SYSRET Target Address Register (STAR)

All steppings (F:0) of the AMD-K6-2 processor Model 8 and the AMD-K6-III processor Model 9 implement the STAR register. This register contains the target EIP address used by the SYSCALL instruction and the 16-bit code and stack segment selector bases used by the SYSCALL and SYSRET instructions. Figure 4 shows the format of the STAR register, and Table 8 defines the function of each field of the STAR register. The STAR register is MSR C000_0081h. For more information about SYSCALL/SYSRET, see the *AMD-K6[®] Processor SYSCALL and SYSRET Instruction Specification Application Note*, order# 21086.




Figure 4. SYSCALL/SYSRET Target Address Register (STAR)—MSR C000_0081h (Models 8 and 9)

Table 8. SYSCALL/SYSRET Target Address Register (STAR) Definition (Models 8 and 9)

Bit	Description	R/W	Function
63–48	SYSRET CS and SS Selector Base	R/W	During the SYSRET instruction, this field is copied into the CS register and the contents of this field, plus 1000b, are copied into the SS register.
47–32	SYSCALL CS and SS Selector Base	R/W	During the SYSCALL instruction, this field is copied into the CS register and the contents of this field, plus 1000b, are copied into the SS register.
31–0	Target EIP Address	R/W	During the SYSCALL instruction, this address is copied into the EIP and points to the new starting address.

AMD-K6[®]-2 Processor Model 8/[F:8]

The AMD-K6-2 processor Model 8/[F:8] provides the following ten MSRs. The first four MSRs are described in “Standard MSRs” on page 15. The contents of ECX selects the MSR to be addressed by the RDMSR and WRMSR instruction.

- Machine-Check Address Register (MCAR)—ECX = 00h
- Machine-Check Type Register (MCTR)—ECX = 01h
- Test Register 12 (TR12)—ECX = 0Eh
- Time Stamp Counter (TSC)—ECX = 10h
- Extended Feature Enable Register (EFER)—ECX = C000_0080h
- Write Handling Control Register (WHCR)—ECX = C000_0082h
- SYSCALL/SYSRET Target Address Register (STAR)—ECX = C000_0081h
- UC/WC Cacheability Control Register (UWCCR)—ECX = C000_0085h
- Processor State Observability Register (PSOR)—ECX = C000_0087h
- Page Flush/Invalidate Register (PFIR)—ECX = C000_0088h

Extended Feature Enable Register (EFER)

The Extended Feature Enable Register (EFER) contains the control bits that enable the extended features of the processor. Figure 5 shows the format of the EFER register, and Table 9 on page 21 defines the function of each bit of the EFER register. The EFER register is MSR C000_0080h.




Figure 5. Extended Feature Enable Register (EFER)—MSR C000_0080h (Model 8/[F:8])

Table 9. Extended Feature Enable Register (EFER) Definition (Model 8/[F:8])

Bit	Description	R/W	Function
63-4	Reserved	R	Writing a 1 to any reserved bit causes a general protection fault to occur. All reserved bits are always read as 0.
3-2	EWBE Control (EWBEC)	R/W	This 2-bit field controls the behavior of the processor with respect to the ordering of write cycles and the EWBE# signal. EFER[3] and EFER[2] are Global EWBE Disable (GEWBED) and Speculative EWBE Disable (SEWBED), respectively.
1	Data Prefetch Enable (DPE)	R/W	DPE must be set to 1 to enable data prefetching (this is the default setting following reset). If enabled, cache misses initiated by a memory read within a 32-byte cache line are conditionally followed by cache-line fetches of the other line in the 64-byte sector.
0	System Call Extension (SCE)	R/W	SCE must be set to 1 to enable the usage of the SYSCALL and SYSRET instructions.

EWBE Control. The AMD-K6-2 processor Model 8/[F:8] contains an 8-byte write merge buffer that allows the processor to conditionally combine data from multiple noncacheable write cycles into this merge buffer. The merge buffer operates in conjunction with the Memory Type Range Registers (MTRRs). Refer to “UC/WC Cacheability Control Register (UWCCR)” on page 26 for a description of the MTRRs.

Merging multiple write cycles into a single write cycle reduces processor bus utilization and processor stalls, thereby increasing the overall system performance.

The presence of the merge buffer creates the potential to perform out-of-order write cycles relative to the processor’s

cache. In general, the ordering of write cycles that are driven externally on the system bus and those that hit the processor's cache can be controlled by the EWBE# signal. If EWBE# is sampled negated, the processor delays the commitment of write cycles to cache lines in the modified state or exclusive state in the processor's cache. Therefore, the system logic can enforce strong ordering by negating EWBE# until the external write cycle is complete, thereby ensuring that a subsequent write cycle that hits the cache does not complete ahead of the external write cycle.

However, the addition of the write merge buffer introduces the potential for out-of-order write cycles to occur between writes to the merge buffer and writes to the processor's cache. Because these writes occur entirely within the processor and are not sent out to the processor bus, the system logic is not able to enforce strong ordering with the EWBE# signal.

The EWBE control (EWBEC) bits provide a mechanism for enforcing three different levels of write ordering in the presence of the write merge buffer:

- EFER[3] is defined as the Global EWBE Disable (GEWBED). When GEWBED equals 1, the processor does not attempt to enforce any write ordering internally or externally (the EWBE# signal is ignored). This is the maximum performance setting.
- EFER[2] is defined as the Speculative EWBE Disable (SEWBED). SEWBED only affects the processor when GEWBED equals 0. If GEWBED equals 0 and SEWBED equals 1, the processor enforces strong ordering for all internal write cycles with the exception of write cycles addressed to a range of memory defined as uncacheable (UC) or write-combining (WC) by the MTRRs. In addition, the processor samples the EWBE# signal. If EWBE# is sampled negated, the processor delays the commitment of write cycles to processor cache lines in the modified state or exclusive state until EWBE# is sampled asserted.

This setting provides performance comparable to, but slightly less than, the performance obtained when GEWBED equals 1 because some degree of write ordering is maintained.

- If GEWBED equals 0 and SEWBED equals 0, the processor enforces strong ordering for all internal and external write cycles. In this setting, the processor assumes, or *speculates*, that strong order must be maintained between writes to the merge buffer and writes that hit the processor's cache. Once the merge buffer is written out to the processor's bus, the EWBE# signal is sampled. If EWBE# is sampled negated, the processor delays the commitment of write cycles to processor cache lines in the modified state or exclusive state until EWBE# is sampled asserted.

This setting is the default after RESET and provides the lowest performance of the three settings because full write ordering is maintained.

Table 10 summarizes the three settings of the EWBE# field for the EFER register, along with the effect of write ordering and performance.

Table 10. EWBE# Settings

EFER[3] (GEWBED)	EFER[2] (SEWBED)	Write Ordering	Performance
1	0 or 1	None	Best
0	1	All except UC/WC	Close-to-Best
0	0	All	Slowest

Enforcing complete write ordering in a uniprocessor system is usually not necessary. In order to achieve the highest level of performance while still maintaining support for the EWBE# signal, AMD recommends that the BIOS set EFER[3:2] to 01b (close-to-best performance). Many uniprocessor systems do not support the EWBE# signal, in which case AMD recommends that the BIOS set EFER[3:2] to 10b or 11b (best performance).

Write Handling Control Register (WHCR)

The AMD-K6-2 processor Model 8/[F:8] contains a split level-1 (L1) 64-Kbyte writeback cache organized as a separate 32-Kbyte instruction cache and a 32-Kbyte data cache with two-way set associativity. The cache line size is 32 bytes, and lines are read from memory using an efficient pipelined burst read cycle. Further performance gains are achieved by the implementation of a write allocation scheme.


A write allocate, if enabled, occurs when the processor has a pending memory write cycle to a cacheable line and the line does not currently reside in the L1 cache. For more information on write allocate, see the *Implementation of Write Allocate in the K86™ Processors Application Note*, order# 21326 and the Cache Organization section of the *AMD-K6[®]-2 Processor Data Sheet*, order# 21850 or the *AMD-K6[®]-III Processor Data Sheet*, order# 21918.

This section describes two programmable mechanisms used by the AMD-K6-2 processor Model 8/[F:8] to determine when to perform write allocate. When either of these mechanisms indicates that a pending write is to a cacheable area of memory, a write allocate is performed.

Before enabling write allocate or changing memory cacheability, the BIOS must write back and invalidate the internal cache by using the WBINVD instruction. In addition, write allocate should be enabled only after performing any memory sizing or typing algorithms.

The Write Handling Control Register (WHCR) is a MSR that contains two fields—the Write Allocate Enable Limit (WAELIM) field, and the Write Allocate Enable 15-to-16-Mbyte (WAE15M) bit (see Figure 6).

Note: *The WHCR register as defined in the Model 6, Model 7, and Model 8/[7:0] has changed in the Model 8/[F:8].*



Note: *Hardware RESET initializes this MSR to all zeros.*

Figure 6. Write Handling Control Register (WHCR)—MSR C000_0082h (Model 8/[F:8])

Write Allocate Enable Limit. The WAELIM field is 10 bits wide. This field, multiplied by 4 Mbytes, defines an upper memory limit. Any pending write cycle that misses the L1 cache and that addresses memory below this limit causes the processor to perform a write allocate (assuming the address is not within a range where write allocates are disallowed). Write allocate is disabled for memory accesses at and above this limit unless the processor determines a pending write cycle is cacheable by means of one of the other write allocate mechanisms—“Write to a Cacheable Page” and “Write to a Sector” (for more information, see the “Cache Organization” chapter in the *AMD-K6[®]-2 Processor Data Sheet*, order# 21850 or the *AMD-K6[®]-III Processor Data Sheet*, order# 21918). The maximum value of this limit is $((2^{10}-1) \cdot 4 \text{ Mbytes}) = 4092 \text{ Mbytes}$. When all the bits in this field are set to 0, all memory is above this limit and the write allocate mechanism is disabled (even if all bits in the WAELIM field are set to 0, write allocates can still occur due to the “Write to a Cacheable Page” and “Write to a Sector” mechanisms).

Once the BIOS determines the amount of RAM installed in the system, this number should also be used to program the WAELIM field. For example, a system with 32 Mbytes of RAM would program the WAELIM field with the value 00_0000_1000b. This value (8), when multiplied by 4 Mbytes, yields 32 Mbytes as the write allocate limit.

Write Allocate Enable 15-to-16-Mbyte. The WAE15M bit is used to enable write allocations for the memory write cycles that address the 1 Mbyte of memory between 15 Mbytes and 16 Mbytes. This bit must be set to 1 to allow write allocates in this memory area. This sub-mechanism of the WAELIM provides a memory hole to prevent write allocates. This memory hole is provided to account for a small number of uncommon memory-mapped I/O adapters that use this particular memory address space. If the system contains one of these peripherals, the bit should be set to 0 (even if the WAE15M bit is set to 0, write allocates can still occur between 15 Mbytes and 16 Mbytes due to the “Write to a Cacheable Page” and “Write to a Sector” mechanisms). The WAE15M bit is ignored if the value in the WAELIM field is set to less than 16 Mbytes.

By definition, write allocations are not performed in the memory area between 640 Kbytes and 1 Mbyte unless the processor determines a pending write cycle is cacheable by

means of “Write to a Cacheable Page” or “Write to a Sector.” It is not safe to perform write allocations between 640 Kbytes and 1 Mbyte (000A_0000h to 000F_FFFFh) because it is considered a noncacheable region of memory. Additionally, if a memory region is defined as write-combinable or uncacheable by a MTRR, write allocations are not performed in that region.

**SYSCALL/SYSRET
Target Address
Register (STAR)**

The STAR register in the AMD-K6-2 processor Model 8/[F:8] is identical to the implementation of this register in the Model 8/[7:0]. See “SYSCALL/SYSRET Target Address Register (STAR)” on page 19.

**UC/WC Cacheability
Control Register
(UWCCR)**

The AMD-K6-2 processor Model 8/[F:8] provides two variable-range Memory Type Range Registers (MTRRs)—MTRR0 and MTRR1—that each specify a range of memory. Each range can be defined as one of the following memory types:

- Uncacheable (UC) memory—Memory read cycles are sourced directly from the specified memory address and the processor does not allocate a cache line. Memory write cycles are targeted at the specified memory address and a write allocation does not occur.
- Write-Combining (WC) memory—Memory read cycles are sourced directly from the specified memory address and the processor does not allocate a cache line. The processor conditionally combines data from multiple noncacheable write cycles that are addressed within this range into a merge buffer. Merging multiple write cycles into a single write cycle reduces processor bus utilization and processor stalls, thereby increasing the overall system performance. This memory type is applicable for linear video frame buffers.

Note: The MTRRs defined in this document are not software compatible to the MTRRs defined by the Pentium Pro and Pentium II processors.

The programmer accesses the MTRRs by addressing the 64-bit MSR known as the UC/WC Cacheability Control Register (UWCCR). The MSR address of the UWCCR is C000_0085h. Following reset, all bits in the UWCCR register are set to 0. MTRR0 (lower 32 bits of the UWCCR register) defines the size and memory type of range 0 and MTRR1 (upper 32 bits) defines the size and memory type of range 1 (see Figure 7).

Prior to programming write-combining or uncacheable areas of memory in the UWCCR, the software must disable the processor's cache, then flush the cache. This can be achieved by setting the CD bit in CR0 to 1 and executing the WBINVD instruction. Following the programming of the UWCCR, the processor's cache must be enabled by setting the CD bit in CR0 to 0.

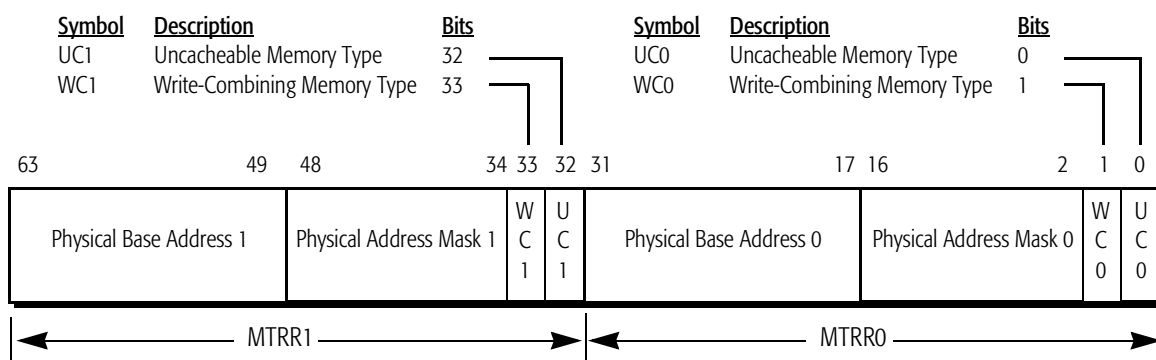


Figure 7. UC/WC Cacheability Control Register (UWCCR)–MSR C000_0085h (Model 8/[F:8])

Physical Base Address n (n=0, 1). This address is the 15 most-significant bits of the physical base address of the memory range. The least-significant 17 bits of the base address are not needed because the base address is by definition always aligned on a 128-Kbyte boundary.

Physical Address Mask n (n=0, 1). This value is the 15 most-significant bits of a physical address mask that is used to define the size of the memory range. This mask is logically ANDed with both the physical base address field of the UWCCR register and the physical address generated by the processor. If the results of the two AND operations are equal, then the generated physical address is considered within the range. That is, if:

$$\text{Mask \& Physical Base Address} = \text{Mask \& Physical Address Generated}$$

then the physical address generated by the processor is in the range.

WCn (n=0, 1). When set to 1, this memory range is defined as write-combinable (refer to Table 11). Write-combinable memory is uncacheable.

UCn (n=0, 1). When set to 1, this memory range is defined as uncacheable (refer to Table 11).

Table 11. WC/UC Memory Type

WCn	UCn	Memory Type
0	0	No effect on cacheability or write-combining
1	0	Write-combining memory range (uncacheable)
0 or 1	1	Uncacheable memory range

Memory-Range Restrictions. The following rules regarding the address alignment and size of each range must be adhered to when programming the physical base address and physical address mask fields of the UWCCR register:

- The minimum size of each range is 128 Kbytes.
- The physical base address must be aligned on a 128-Kbyte boundary.
- The physical base address must be *range-size aligned*. For example, if the size of the range is 1 Mbyte, then the physical base address must be aligned on a 1-Mbyte boundary.
- All bits set to 1 in the physical address mask must be contiguous. Likewise, all bits set to 0 in the physical address mask must be contiguous. For example:

111_1111_1100_0000b is a valid physical address mask

111_1111_1101_0000b is invalid

Table 12 lists the valid physical address masks and the resulting range sizes that can be programmed in the UWCCR register.

Table 12. Valid Masks and Range Sizes

Masks	Size
111_1111_1111_1111b	128 Kbytes
111_1111_1111_1110b	256 Kbytes
111_1111_1111_1100b	512 Kbytes
111_1111_1111_1000b	1 Mbyte
111_1111_1111_0000b	2 Mbytes

Table 12. Valid Masks and Range Sizes (continued)

Masks	Size
111_1111_1110_0000b	4 Mbytes
111_1111_1100_0000b	8 Mbytes
111_1111_1000_0000b	16 Mbytes
111_1111_0000_0000b	32 Mbytes
111_1110_0000_0000b	64 Mbytes
111_1100_0000_0000b	128 Mbytes
111_1000_0000_0000b	256 Mbytes
111_0000_0000_0000b	512 Mbytes
110_0000_0000_0000b	1 Gbyte
100_0000_0000_0000b	2 Gbytes
000_0000_0000_0000b	4 Gbytes

Example. Suppose that the range of memory from 16 Mbytes to 32 Mbytes is uncacheable, and the 8-Mbyte range of memory on top of 1 Gbyte is write-combinable. Range 0 is defined as the uncacheable range, and range 1 is defined as the write-combining range.

Extracting the 15 most-significant bits of the 32-bit physical base address that corresponds to 16 Mbytes (0100_0000h) yields a physical base address 0 field of 000_0000_1000_0000b. Because the uncacheable range size is 16 Mbytes, the physical mask value 0 field is 111_1111_1000_0000b, according to Table 12. Bit 1 of the UWCCR register (WC0) is set to 0 and bit 0 of the UWCCR register is set to 1 (UC0).

Extracting the 15 most-significant bits of the 32-bit physical base address that corresponds to 1 Gbyte (4000_0000h) yields a physical base address 1 field of 010_0000_0000_0000b. Because the write-combining range size is 8 Mbytes, the physical mask value 1 field is 111_1111_1100_0000b, according to Table 12. Bit 33 of the UWCCR register (WC1) is set to 1 and bit 32 of the UWCCR register is set to 0 (UC1).

Processor State Observability Register (PSOR)

The AMD-K6-2 processor Model 8/[F:8] provides the Processor State Observability Register (PSOR) (see Figure 8).




Figure 8. Processor State Observability Register (PSOR)—MSR C000_0087h (Model 8/[F:8])

NOL2. This read-only bit indicates whether the processor contains a L2 cache. This bit is always set to 1 for Model 8/[F:8].

STEP. This read-only field contains the stepping ID. This is identical to the value returned by CPUID standard function 1 in EAX[3:0].

BF. This read-only field contains the value of the BF signals sampled by the processor during the falling transition of RESET, which allows the BIOS to determine the frequency of the host bus. The core frequency must first be determined using the Time Stamp Counter method (See “Time Stamp Counter (TSC)” on page 15). The core frequency is then divided by the processor-clock to bus-clock ratio as determined by the BF field of the PSOR register (see Table 13). The result is the frequency of the processor bus.

Table 13. Processor-to-Bus Clock Ratios

State of BF[2:0]	Processor-Clock to Bus-Clock Ratio
100b	2.5x
101b	3.0x
110b	6.0x*
111b	3.5x
000b	4.5x
001b	5.0x
010b	4.0x
011b	5.5x

Note:
 * The 2.0x ratio that was supported on Models 6, 7, and 8/[7:0] is no longer supported on Model 8/[F:8] or Model 9. Instead, if BF[2:0] equals 110b, a ratio of 6.0x is selected.

Page Flush/Invalidate Register (PFIR)

The AMD-K6-2 processor Model 8/[F:8] contains the Page Flush/Invalidate Register (PFIR) (see Figure 9 on page 31) that allows cache invalidation and optional flushing of a specific 4-Kbyte page from the linear address space. The total amount of L1 cache in the Model 8/[F:8] is 64 Kbytes. Using this register can result in a much lower cycle count for flushing particular pages versus flushing the entire cache. When the PFIR is written to (using the WRMSR instruction), the invalidation and, optionally, the flushing begins.




Figure 9. Page Flush/Invalidate Register (PFIR)—MSR C000_0088h (Model 8/[F:8])

LINPAGE. This 20-bit field must be written with bits 31:12 of the linear address of the 4-Kbyte page that is to be invalidated and optionally flushed from the L1 cache.

PF. If an attempt to invalidate or flush a page results in a page fault, the processor sets the PF bit to 1, and the invalidate or flush operation is not performed (even though invalidate operations do not normally generate page faults). In this case, an actual page fault exception is not generated. If the PF bit equals 0 after an invalidate or flush operation, then the operation executed successfully. The PF bit must be read after every write to the PFIR register to determine if the invalidate or flush operation executed successfully.

F/I. This bit is used to control the type of action that occurs to the specified linear page. If a 0 is written to this bit, the operation is a flush, in which case all cache lines in the modified state within the specified page are written back to memory, after which the entire page is invalidated. If a 1 is written to this bit, the operation is an invalidation, in which case the entire page is invalidated without the occurrence of any writebacks.

AMD-K6[®]-III Processor Model 9

The AMD-K6-III processor Model 9 provides the following eleven MSRs. The first four MSRs are described in “Standard MSRs” on page 15. The contents of ECX selects the MSR to be addressed by the RDMSR and WRMSR instruction.

- Machine-Check Address Register (MCAR)—ECX = 00h
- Machine-Check Type Register (MCTR)—ECX = 01h
- Test Register 12 (TR12)—ECX = 0Eh
- Time Stamp Counter (TSC)—ECX = 10h
- Extended Feature Enable Register (EFER)—ECX = C000_0080h
- Write Handling Control Register (WHCR)—ECX = C000_0082h
- SYSCALL/SYSRET Target Address Register (STAR)—ECX = C000_0081h
- UC/WC Cacheability Control Register (UWCCR)—ECX = C000_0085h
- Processor State Observability Register (PSOR)—ECX = C000_0087h
- Page Flush/Invalidate Register (PFIR)—ECX = C000_0088h
- Level-2 Cache Array Access Register (L2AAR)—ECX = C000_0089h

Extended Feature Enable Register (EFER)

Bits 3:0 of the EFER register in the AMD-K6-III processor Model 9 are identical to the implementation of these bits in the Model 8/[F:8]. See “Extended Feature Enable Register (EFER)” on page 20. The L2 Disable bit (L2D)—EFER[4]—is an addition to the EFER register in the Model 9.

Figure 10 shows the format of the EFER register, and Table 14 defines the function of each bit of the EFER register. The EFER register is MSR C000_0080h.




Figure 10. Extended Feature Enable Register (EFER)—MSR C000_0080h (Model 9)

Table 14. Extended Feature Enable Register (EFER) Definition (Model 9)

Bit	Description	R/W	Function
63–5	Reserved	R	Writing a 1 to any reserved bit causes a general protection fault to occur. All reserved bits are always read as 0.
4	L2 Disable (L2D)	R/W	If L2D is set to 1, the L2 cache is completely disabled. This bit is provided for debug and testing purposes. For normal operation and maximum performance, this bit must be set to 0 (this is the default setting following reset).
3–2	EWBE Control (EWBEC)	R/W	This 2-bit field controls the behavior of the processor with respect to the ordering of write cycles and the EWBE# signal. EFER[3] and EFER[2] are Global EWBE Disable (GEWBED) and Speculative EWBE Disable (SEWBED), respectively.
1	Data Prefetch Enable (DPE)	R/W	DPE must be set to 1 to enable data prefetching (this is the default setting following reset). If enabled, cache misses initiated by a memory read within a 32-byte cache line are conditionally followed by cache-line fetches of the other line in the 64-byte sector.
0	System Call Extension (SCE)	R/W	SCE must be set to 1 to enable the usage of the SYSCALL and SYSRET instructions.

Note: Setting L2D to 1 does not guarantee cache coherency. To ensure coherency, the processor’s caches must be disabled (by setting the CD bit of the CR0 register to 1), then flushed prior to setting L2D to 1.

Write Handling Control Register (WHCR)

The AMD-K6-III processor contains a split level-1 (L1) 64-Kbyte writeback cache organized as a separate 32-Kbyte instruction cache and a 32-Kbyte data cache with two-way set associativity. The cache line size is 32 bytes, and lines are read from memory

using an efficient pipelined burst read cycle. In addition, the AMD-K6-III processor also contains a 256-Kbyte, 4-way set associative, unified level-2 (L2) cache. Further performance gains are achieved by the implementation of a write allocation scheme.

The WHCR register in the AMD-K6-III processor Model 9 is identical to the implementation of this register in the Model 8/[F:8]. See “Write Handling Control Register (WHCR)” on page 23.

Note: The WHCR register as defined in the Model 6, Model 7, and Model 8/[7:0] has changed in the Model 9.

SYSCALL/SYSRET Target Address Register (STAR)

The STAR register in the AMD-K6-III processor Model 9 is identical to the implementation of this register in the Model 8/[7:0]. See “SYSCALL/SYSRET Target Address Register (STAR)” on page 19.

UC/WC Cacheability Control Register (UWCCR)

The UWCCR register in the AMD-K6-III processor Model 9 is identical to the implementation of this register in the Model 8/[F:8]. See “UC/WC Cacheability Control Register (UWCCR)” on page 26.

Processor State Observability Register (PSOR)

The AMD-K6-III processor Model 9 provides the Processor State Observability Register (PSOR) (see Figure 11).




Figure 11. Processor State Observability Register (PSOR)—MSR C000_0087h (Model 9)

NOL2. This read-only bit indicates whether the processor contains an L2 cache. This bit is always set to 0 for Model 9.

STEP. This read-only field contains the stepping ID. This is identical to the value returned by CPUID standard function 1 in EAX[3:0].

BF. This read-only field contains the value of the BF signals sampled by the processor during the falling transition of RESET, which allows the BIOS to determine the frequency of the host bus. The core frequency must first be determined using the Time Stamp Counter method (See “Time Stamp Counter (TSC)” on page 15). The core frequency is then divided by the processor-clock to bus-clock ratio as determined by the BF field of the PSOR register (see Table 13 on page 31). The result is the frequency of the processor bus.

Page Flush/Invalidate Register (PFIR)

The PFIR register in the AMD-K6-III processor Model 9 is identical to the implementation of this register in the Model 8/[F:8]. See “Page Flush/Invalidate Register (PFIR)” on page 31. The invalidate and flush operations affect the processor’s L1 and L2 caches on the Model 9.

Level-2 Cache Array Access Register (L2AAR)

The AMD-K6-III processor Model 9 provides the L2AAR register that allows for direct access to the L2 cache and L2 tag arrays. The 256-Kbyte L2 cache in the AMD-K6-III processor is organized as shown in Figure 12:

- Four 64-Kbyte ways
- Each way contains 1024 sets
- Each set contains four 64-byte sectors (one sector in each way)
- Each sector contains two 32-byte cache lines
- Each cache line contains four 8-byte octets
- Each octet contains an upper and lower dword (4 bytes)

Each line within a sector contains its own MESI state bits, and associated with each sector is a tag and LRU (Least Recently Used) information.




Figure 12. L2 Cache Organization

Figure 13 shows the L2 cache sector and line organization. If bit 5 of the address of a cache line equals 1, then this cache line is stored in Line 1 of a sector. Similarly, if bit 5 of the address of a cache line equals 0, then this cache line is stored in Line 0 of a sector.

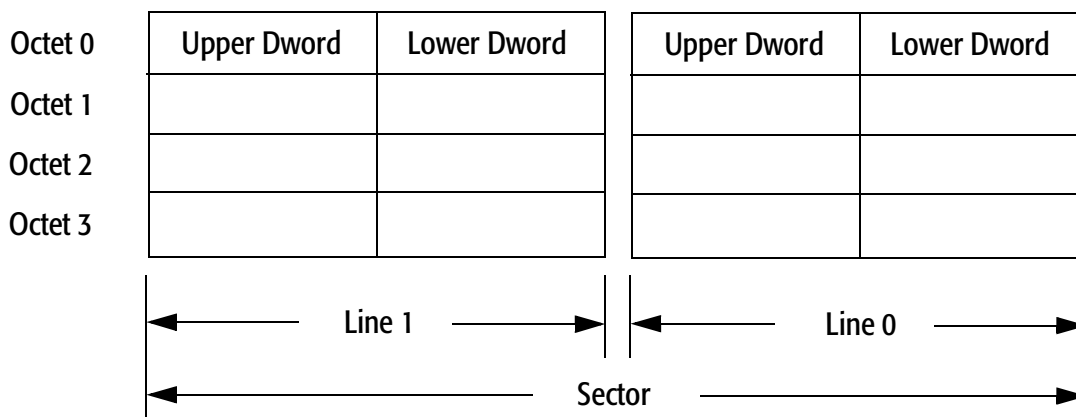


Figure 13. L2 Cache Sector and Line Organization

The L2AAR register is MSR C000_0089h. The operation that is performed on the L2 cache is a function of the instruction executed—RDMSR or WRMSR—and the contents of the EDX register. The EDX register specifies the location of the access, and whether the access is to the L2 cache data or tags (refer to Figure 14). Bit 20 of EDX (T/D) determines whether the access is to the L2 cache data or tag. Table 15 describes the operation that is performed based on the instruction and the T/D bit.




Figure 14. L2 Tag or Data Location - EDX

Table 15. Tag versus Data Selector

Instruction	T/D (EDX[20])	Operation
RDMSR	0	Read dword from L2 data array into EAX. Dword location is specified by EDX.
RDMSR	1	Read tag, line state and LRU information from L2 tag array into EAX. Location of tag is specified by EDX.
WRMSR	0	Write dword to the L2 data array using data in EAX. Dword location is specified by EDX.
WRMSR	1	Write tag, line state and LRU information into L2 tag array from EAX. Location of tag is specified by EDX.

When the L2AAR is read or written, EDX is left unchanged. This facilitates multiple accesses when testing the entire cache/tag array.

If the L2 cache data is read (as opposed to reading the tag information), the result (dword) is placed in EAX in the format as illustrated in Figure 15. Similarly, if the L2 cache data is written, the write data is taken from EAX.



Figure 15. L2 Data - EAX

If the L2 tag is read (as opposed to reading the cache data), the result is placed in EAX in the format as illustrated in Figure 16. Similarly, if the L2 tag is written, the write data is taken from EAX.

When writing to the L2 tag, special consideration must be given to the least significant bit of the Tag field of the EAX register—EAX[15]. The length of the L2 tag required to support the 256-Kbyte L2 cache on the Model 9 is 16 bits, which corresponds to bits 31:16 of the EAX register. However, the processor provides a total of 17 bits for storing the L2 tag—that is, 16 bits for the tag (EAX[31:16]), plus an additional bit for internal purposes (EAX[15]). During normal operation, the processor ensures that this additional bit (bit 15) always corresponds to the set in which the tag resides. Note that bits 15:6 of the address determine the set, in which case bit 15 equal to 0 addresses sets 0 through 511, and bit 15 equal to 1 addresses sets 512 through 1023.

In order to set the full 17-bit L2 tag properly when using the L2AAR register, EAX[15] must likewise correspond to the set in which the tag is being written—that is, EAX[15] must be equal to EDX[15] (refer to Figure 14 and Figure 16).

It is important to note that this special consideration is only required if the processor will subsequently be expected to properly execute instructions or access data from the L2 cache following the setup of the L2 cache by means of the L2AAR register. If the intent of using the L2AAR register is solely to test or debug the L2 cache without the subsequent intent of executing instructions or accessing data from the L2 cache, then this consideration is not required.

When accessing the L2 tag, the Line, Octet, and Dword fields of the EDX register are ignored.





Figure 16. L2 Tag Information - EAX

LRU (Least Recently Used). For the 4-way set associative L2 cache, each way has a 2-bit LRU field for each sector. Values for the LRU field are 00b, 01b, 10b, and 11b, where 00b indicates that the sector is “most recently used,” and 11b indicates that the sector is “least recently used” (see Figure 17). EAX[7:6] indicate LRU information for Way 0, EAX[5:4] for Way 1, EAX[3:2] for Way 2, and EAX[1:0] for Way 3.



LRU Values

- 00b Most Recently Used
- 01b Used More Recent Than 10b, But Less Recent Than 00b
- 10b Used More Recent Than 11b, But Less Recent Than 01b
- 11b Least Recently Used

Figure 17. LRU Byte

New AMD-K6[®] Processor Instructions

All models of the AMD-K6 processor implement the following new instruction set:

- MMX[™] Instructions—57 new instructions for multimedia software. See the *AMD-K6[®] Processor Multimedia Technology Manual*, order# 20726 for more information.

AMD-K6-2 processor Models 8 and above, and the AMD-K6-III processor Model 9 implement the following new instructions:

- 3DNow![™] Instructions—21 new instructions for multimedia software. See the *3DNow![™] Technology Manual*, order# 21928 for more information.
- SYSCALL and SYSRET—See the *SYSCALL and SYSRET Instruction Specification Application Note*, order# 21086 for more information.

Additional Considerations

Software Timing Dependencies Relative to Memory Controller Setup

Processors in the K86 family differ from other processors with regards to instruction latencies and the order or priority of processor bus cycles. Timing-dependent software that relies on the specific latencies of other processors should be re-tested for proper operation with the K86 processor. In addition, re-testing should be performed on components with variable timing (such as, memory modules, oscillators, and timers).

Particular attention should be paid to memory-setup subroutines that determine the type of DRAM in the system. Some chipsets may not tolerate a DRAM mode change (such as, EDO to SDRAM) on the same clock as a DRAM refresh cycle. For example some chipsets do not tolerate having its memory refresh enabled prior to changing memory mode types. Refresh should only be enabled after the memory type has been determined.

Note: *The BIOS for the K86 family of processors should enable the write allocate mechanisms only after performing any memory sizing or typing algorithms.*

Pipelining Support

All production models and steppings of the AMD-K6 processor support the WAELIM form of write allocate, which is the only form of write allocate that should be enabled. AMD does not recommend enabling the obsolete form of write allocate (WCDE) because system performance can be degraded by doing so.

Early implementations of the AMD-K6 processor did not support the WHCR register and therefore did not support the WAELIM form of write allocate. WCDE was the only form of write allocate supported, which required the chipset to assert KEN# for cacheable memory write cycles. Because KEN# is sampled by the processor on the clock edge on which the first BRDY# or NA# is sampled asserted, some chipsets that supported the WCDE form of write allocate did not assert NA# during write cycles in order to prevent the processor from sampling KEN# before it was valid (in this case, BRDY# was used by the processor to sample KEN#). If NA# is not asserted during memory write cycles, then the processor does not fully take advantage of the potential performance gains that bus pipelining can achieve.

For proper functionality, always program the WCDE bit to 0 for Models 6, 7, and 8/[7:0]. Models 8/[F:8] and 9 do not support the WCDE bit.

Read-Only Memory

The processor's caches must be flushed prior to defining any area of memory as cacheable and read only. (The BIOS is typically "shadowed" into main memory and defined as cacheable and read only.) If the caches are not flushed, then a line that resides in the processor's cache that falls within a read-only area of memory can be written to, which would place the cache line in the modified state. If this modified line is subsequently replaced and written back to memory, then the system may hang (or other unpredictable effects may occur) because the writeback is directed to an area of memory defined as read only by the chipset.