DATACOM

DM-SV01 Server

BIOS User Manual

Revision 1.2 – Last updated Dec 2022

DATACOM

The information and specifications provided in this manual are subject to change without notice and are not recognized as any kind of contract. Datacom specifically disclaims any warranties, expressed or implied, of merchantability, fitness for any particular purpose and noninfringement, including those arising from a course of usage, dealing or trade practice.

Although every precaution has been taken in the preparation of this document, Datacom does not assume any liability for any errors or omissions as well as no obligation is assumed for damage resulting from the use of the information contained in this manual.

In no event will Datacom or its suppliers be liable for any direct, indirect, special, speculative, consequential or incidental, including, without limitation, lost profits or loss or damage to data or hardware arising out of the use or inability to use this product or this manual, even in case Datacom or its suppliers have been advised of the possibility of such damages. In particular, Datacom shall not have liability for the costs of replacing, repairing or recovering software, data or hardware related to the product or its use.

Table of Contents

1 Introduction

This document is intended to list and describe all the configurations available in the DM-SV01 BIOS/UEFI environment.

This document is compatible with BIOS releases starting from “**DM_SV01_R200”**.

2 BIOS Overview

2.1 Accessing BIOS/UEFI Menu

The BIOS/UEFI menu can be accessed by hitting the “F2” key during the boot process.

The boot menu can be accessed by hitting the “F5” key during the boot process. In the boot menu, it is possible to select the device to boot from in the current boot process.

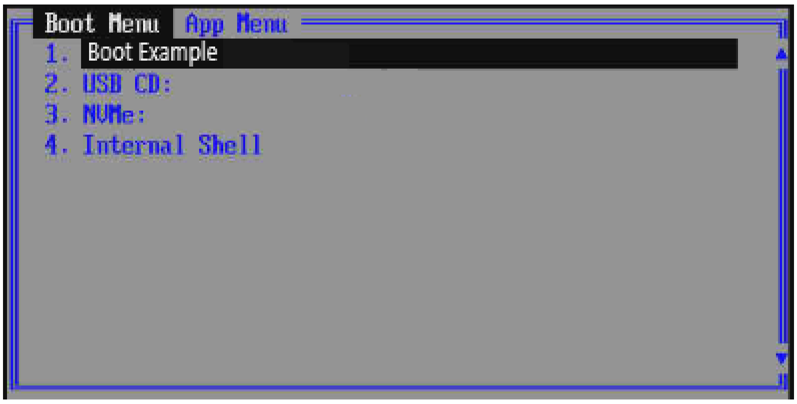

Figure 1: Boot Menu

In the boot menu it is possible to hit <TAB> and access the “App Menu”, which allows the user to enter the BIOS/UEFI menu or to display the Diagnostic Splash screen.

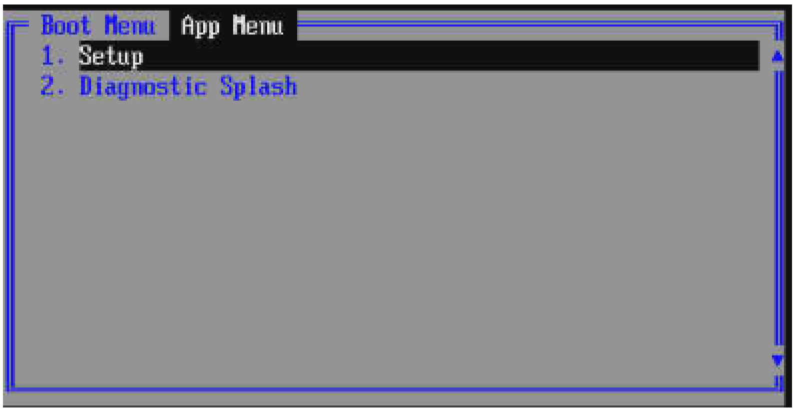

Figure 2: App Menu

2.2 Function Keys

The table below lists the main function keys used to navigate and control BIOS/UEFI settings.

Key	Description
F1	Help
Esc	Exit the current menu
Enter	Enter the selected menu / confirm selected entry
Up/Down arrows (↑ ↓)	Navigate through the menu items
Left/Right arrows (← →)	Navigate through the sub-menus
Plus/minus signs (+ -)	Change boot order in the “Boot” menu
F9	Restore setup defaults
F10	Save current configuration and Exit BIOS/UEFI

2.3 “Auto” Setting

There are several menu items which use “Auto” as one of the configuration options. The “Auto” setting means that the FW (BIOS/UEFI) is responsible for deciding which setting will be used, according to the current server configuration and the recommendations of the processors’ manufacturer.

2.4 Restore Default Settings

The user can restore all BIOS/UEFI configurations to the factory default settings by means of the following operations:

1. Using the BIOS menu “Load Setup Defaults” (please refer to section 3.7.4 Load Setup Defaults).
2. Performing a “clear CMOS” directly in the server’s hardware.

The clear CMOS procedure can be found in the document (7) “DM-SV01 Server - Product Manual”.

2.5 BIOS/UEFI Post Codes

When the BIOS/UEFI is being initialized, some codes are provided as checkpoints for the boot process. These codes can be checked by means of the DM-SV01 debug port, when connecting the DM-SV01 debug card. If the system presents some issue when booting, the user can check the last post code available in order to identify the point of failure.

3 BIOS/UEFI Setup

3.1 Main Menu

The figure 3 shows the BIOS Main menu screen.

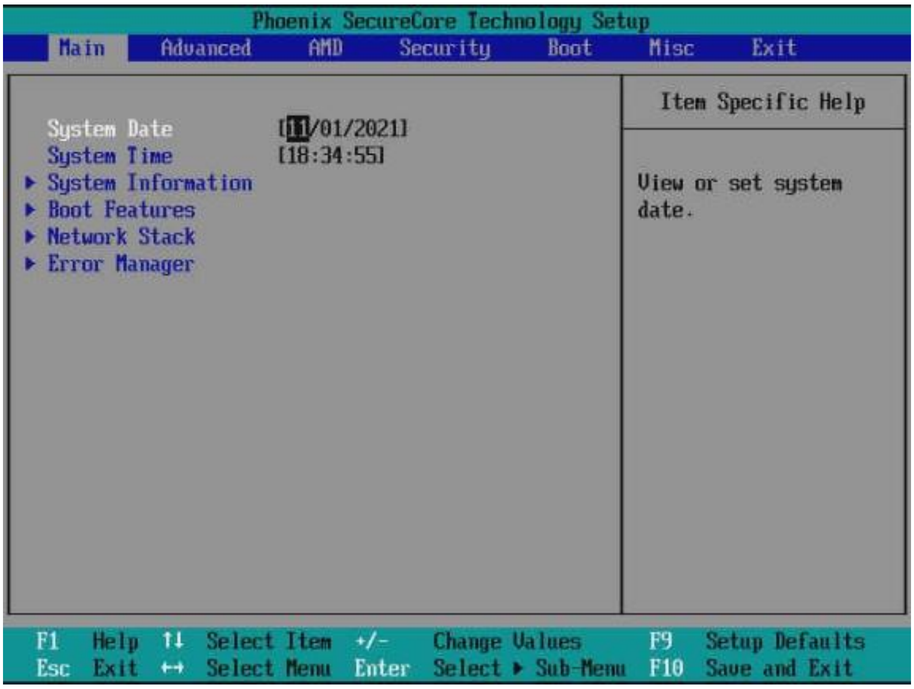

```
                 Phoenix SecureCore Technology Setup
  Main   Advanced   AMD   Security   Boot   Misc   Exit

   System Date          [11/01/2021]          Item Specific Help
   System Time          [18:34:55]
 ▸ System Information                        View or set system
 ▸ Boot Features                             date.
 ▸ Network Stack
 ▸ Error Manager

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 3: BIOS Main menu

3.1.1 System Date

To update system date, use the <arrow keys> to navigate to **System Date** setting. Use the <Enter> key to highlight the field to be updated and then type the desired value of that field.

Menu Item	Description	Default Config	Format
System Date	Configure system date.	Current system date	MM/DD/YYYY

3.1.2 System Time

To update system time, use the <arrow keys> to navigate to **System Time** setting. Use the <Enter> key to highlight the field to be updated and then type the desired value of that field.

Menu Item	Description	Default Config	Format
System Time	Configure system time.	Current system time	HH:MM:SS

3.1.3 System Information

The **System Information** menu is a read only option that shows a brief description of the system components installed in the equipment.

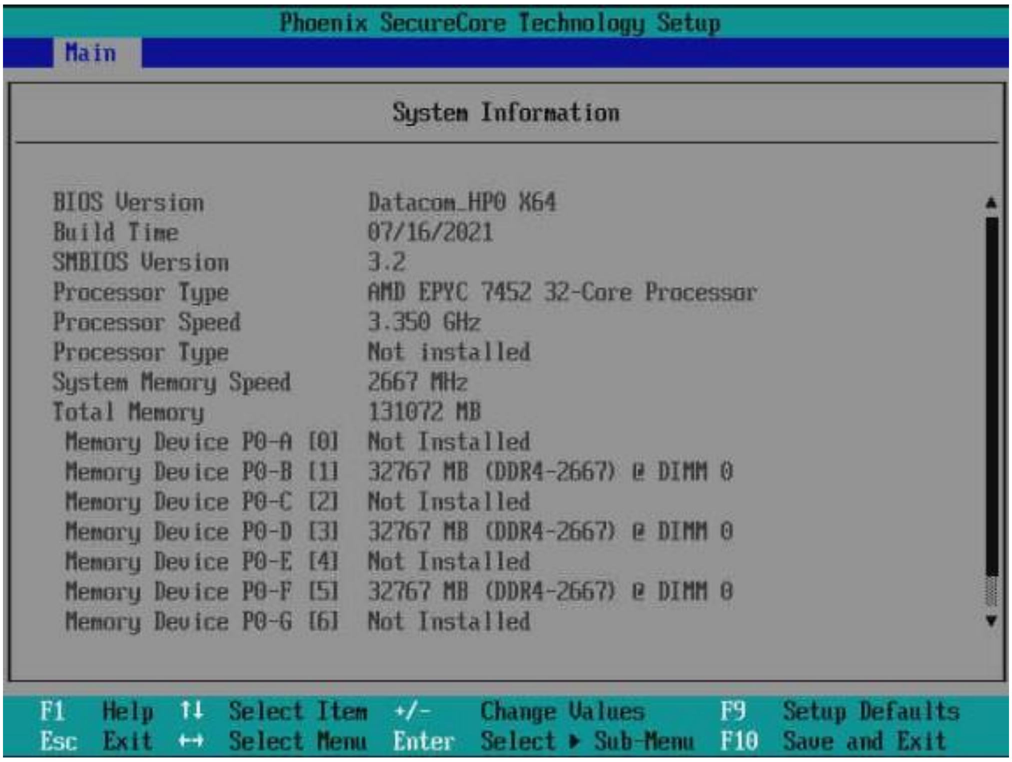

```
                    Phoenix SecureCore Technology Setup
  Main

                          System Information

  BIOS Version               Datacom_HP0 X64
  Build Time                 07/16/2021
  SMBIOS Version             3.2
  Processor Type             AMD EPYC 7452 32-Core Processor
  Processor Speed            3.350 GHz
  Processor Type             Not installed
  System Memory Speed        2667 MHz
  Total Memory               131072 MB
   Memory Device P0-A [0]    Not Installed
   Memory Device P0-B [1]    32767 MB (DDR4-2667) @ DIMM 0
   Memory Device P0-C [2]    Not Installed
   Memory Device P0-D [3]    32767 MB (DDR4-2667) @ DIMM 0
   Memory Device P0-E [4]    Not Installed
   Memory Device P0-F [5]    32767 MB (DDR4-2667) @ DIMM 0
   Memory Device P0-G [6]    Not Installed

  F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
  Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 4: System Information menu

The following information is displayed when accessing this menu.

Menu Item	**Description**	**Format**
BIOS Version	BIOS release version	String
Build Time	BIOS release build time	MM/DD/YYYY
SMBIOS Version	SMBIOS Current Version	String
Processor Type	Information about CPU unit installed at socket P0	String
Processor Speed	Maximum Clock Speed (Boost) of CPU unit installed at socket P0 (in GHz)	String
Processor Type	Information about CPU unit installed at socket P1	String
Processor Speed	Maximum Clock Speed (Boost) of CPU unit installed at socket P1 (in GHz). This field is hidden if no CPU is present at socket P1.	String
System Memory Speed	Current Configured DRAM Speed	String
Total Memory	Total DRAM memory capacity installed in the system	String
Memory Device P0-A [0]	Information about DDR memory installed at socket P0, slot A	String

Memory Device P0-B [1]	Information about DDR memory installed at socket P0, slot B	String
Memory Device P0-C [2]	Information about DDR memory installed at socket P0, slot C	String
Memory Device P0-D [3]	Information about DDR memory installed at socket P0, slot D	String
Memory Device P0-E [4]	Information about DDR memory installed at socket P0, slot E	String
Memory Device P0-F [5]	Information about DDR memory installed at socket P0, slot F	String
Memory Device P0-G [6]	Information about DDR memory installed at socket P0, slot G	String
Memory Device P0-H [7]	Information about DDR memory installed at socket P0, slot H	String
Memory Device P1-A [8]	Information about DDR memory installed at socket P1, slot A	String
Memory Device P1-B [9]	Information about DDR memory installed at socket P1, slot B	String
Memory Device P1-C [10]	Information about DDR memory installed at socket P1, slot C	String
Memory Device P1-D [11]	Information about DDR memory installed at socket P1, slot D	String
Memory Device P1-E [12]	Information about DDR memory installed at socket P1, slot E	String
Memory Device P1-F [13]	Information about DDR memory installed at socket P1, slot F	String
Memory Device P1-G [14]	Information about DDR memory installed at socket P1, slot G	String
Memory Device P1-H [15]	Information about DDR memory installed at socket P1, slot H	String

3.1.4 Boot Features

The Boot Features menu is used to configure some basic settings related to the behavior of the system during the boot process.

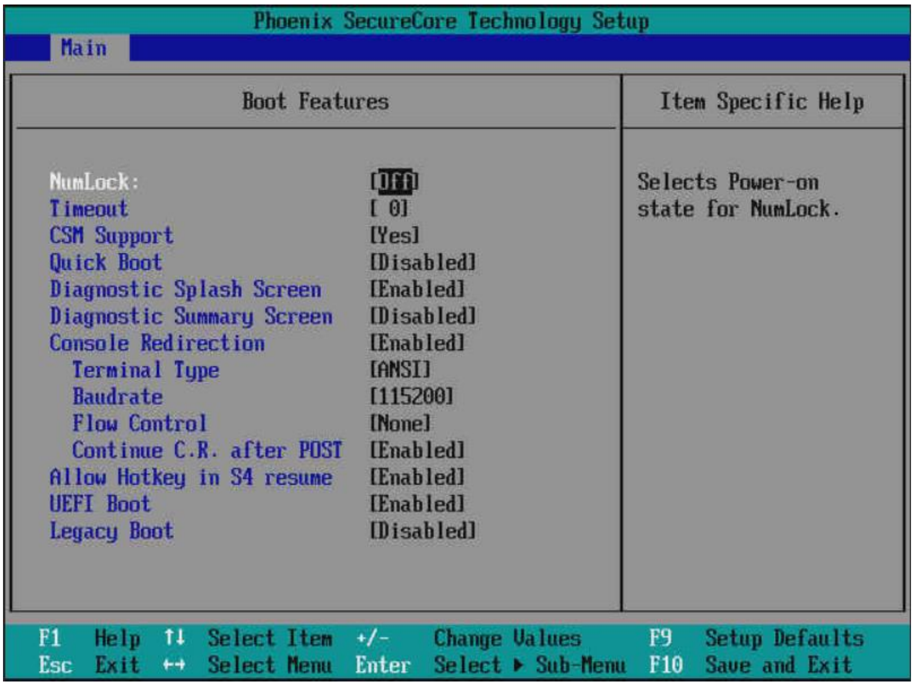

```
                     Phoenix SecureCore Technology Setup
 Main
            Boot Features                         Item Specific Help

 NumLock:                      [Off]              Selects Power-on
 Timeout                       [ 0]               state for NumLock.
 CSM Support                   [Yes]
 Quick Boot                    [Disabled]
 Diagnostic Splash Screen      [Enabled]
 Diagnostic Summary Screen     [Disabled]
 Console Redirection           [Enabled]
   Terminal Type               [ANSI]
   Baudrate                    [115200]
   Flow Control                [None]
   Continue C.R. after POST    [Enabled]
 Allow Hotkey in S4 resume     [Enabled]
 UEFI Boot                     [Enabled]
 Legacy Boot                   [Disabled]

 F1   Help  ↑↓  Select Item   +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu   Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 5: Boot Features menu

3.1.4.1 NumLock

Menu Item	Description	Default Config	Options
NumLock	Defines if the NumLock state is ON or OFF when Powering the system on.	Off	[Off] [On]

3.1.4.2 Timeout

Menu Item	Description	Default Config	Options
Timeout	Number of seconds that POST will wait for the user input before booting.	0	[<value in seconds>]

When the **timeout** value is higher than 0, the BIOS will show the screen presented in figure 6. The screen keeps showing for the amount of time set in the timeout BIOS option and, during this time, the user can hit F2 key to access BIOS setup. When the time expires without the user hitting the F2 key, the boot process proceeds normally.

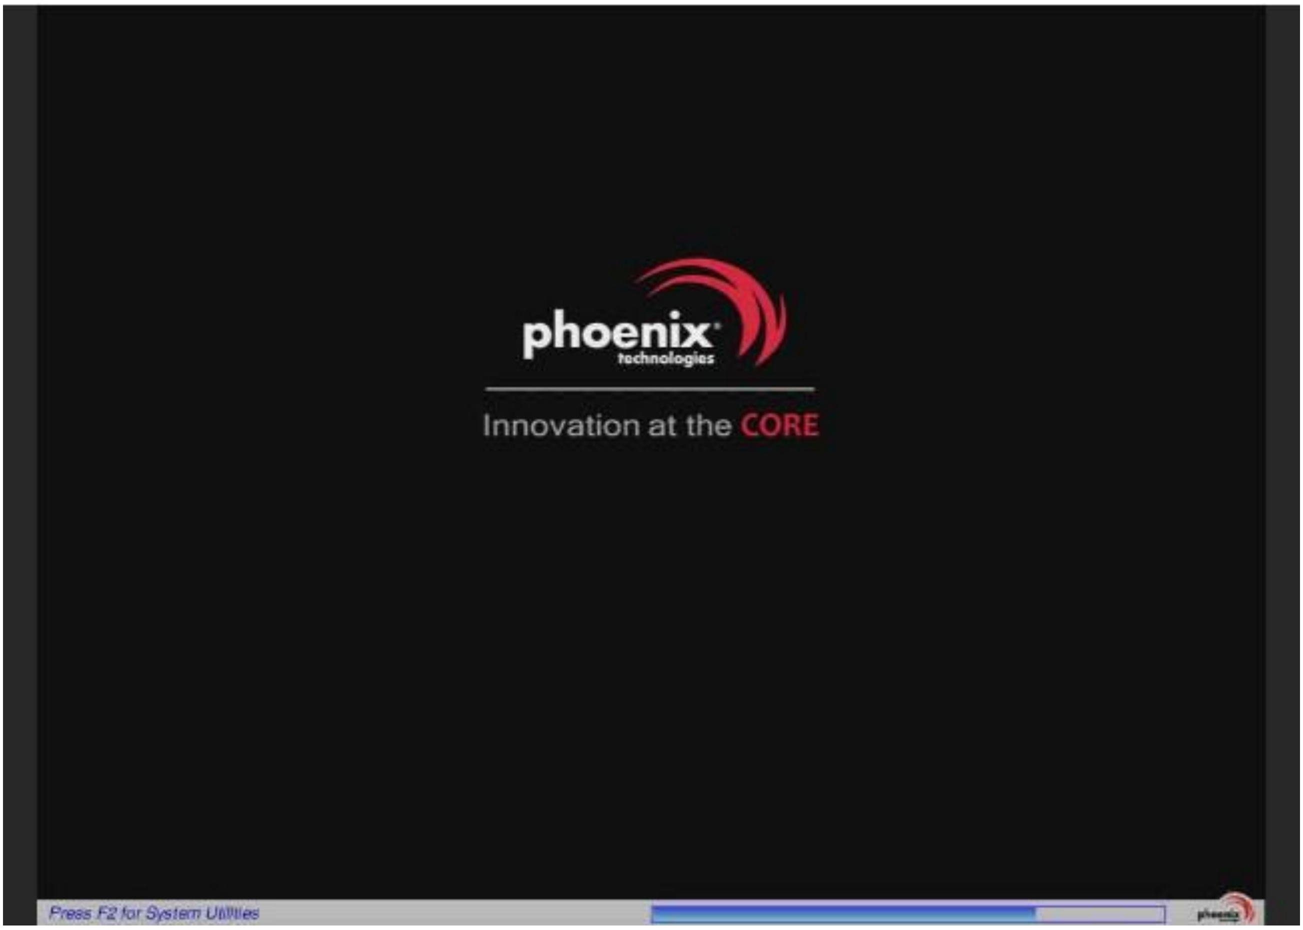

Press F2 for System Utilities

Figure 6: Screen shown when timeout value is higher than 0

3.1.4.3 CSM Support

The CSM (Compatibility Support Mode) Support option is used to activate the legacy boot mode. When the CSM support is set to “Yes”, the Legacy Boot configuration is enabled (please refer to section 3.1.4.10 Legacy Boot).

Note: The DM-SV01 does not support booting in Legacy Mode. Although some BIOS releases may have this option available, the user must keep CSM Support with default settings.

Menu Item	Description	Default Config	Options
CSM Support	Enables/Disables the CSM (Compatibility Support Mode) to support legacy functions.	No	[Yes] [No]

3.1.4.4 Quick Boot

The Quick Boot feature, when enabled, configures the system to skip some hardware verification procedures (specially RAM checking) during POST, making the boot process faster.

Menu Item	Description	Default Config	Options
Quick Boot	Enable/Disable Quick Boot.	Disabled	[Enabled] [Disabled]

3.1.4.5 Diagnostic Splash Screen

Menu Item	Description	Default Config	Options
Diagnostic Splash Screen	Enable/Disable Diagnostic Splash Screen. If Disabled, the Diagnostic Splash Screen will not display unless you press the hotkey* during boot.	Enabled	[Enabled] [Disabled]

The option above Enables or disables the Splash Screen display during the boot process. The figure 7 below shows an example of a splash screen being displayed. The screen is shown for a few seconds and it brings some basic information about the system, such as BIOS release and build time, CPU and RAM information, etc. During the time the splash screen is being displayed, the user can hit F2 key to enable access to the BIOS menu or F5 to access the boot menu.

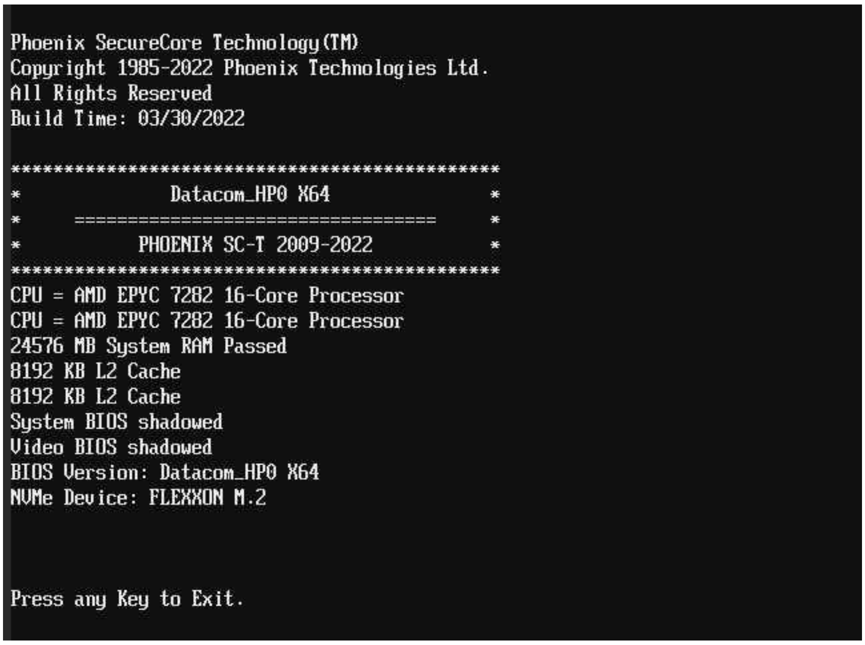

```
Phoenix SecureCore Technology(TM)
Copyright 1985-2022 Phoenix Technologies Ltd.
All Rights Reserved
Build Time: 03/30/2022

**********************************************
*              Datacom_HP0 X64               *
*     =================================      *
*           PHOENIX SC-T 2009-2022           *
**********************************************
CPU = AMD EPYC 7282 16-Core Processor
CPU = AMD EPYC 7282 16-Core Processor
24576 MB System RAM Passed
8192 KB L2 Cache
8192 KB L2 Cache
System BIOS shadowed
Video BIOS shadowed
BIOS Version: Datacom_HP0 X64
NVMe Device: FLEXXON M.2


Press any Key to Exit.
```

Figure 7: Splash Screen Example

* If the user wants to display the splash screen manually, the procedure below can be followed.

1) Press <F5> key during the boot process.
2) The screen with Boot Menu and App menu will be shown.
3) Hit <TAB> to select the App Menu
4) Use the arrow keys to navigate to the “Diagnostic Splash” option.
5) Hit <ENTER> and the Diagnostic Splash screen will be displayed.

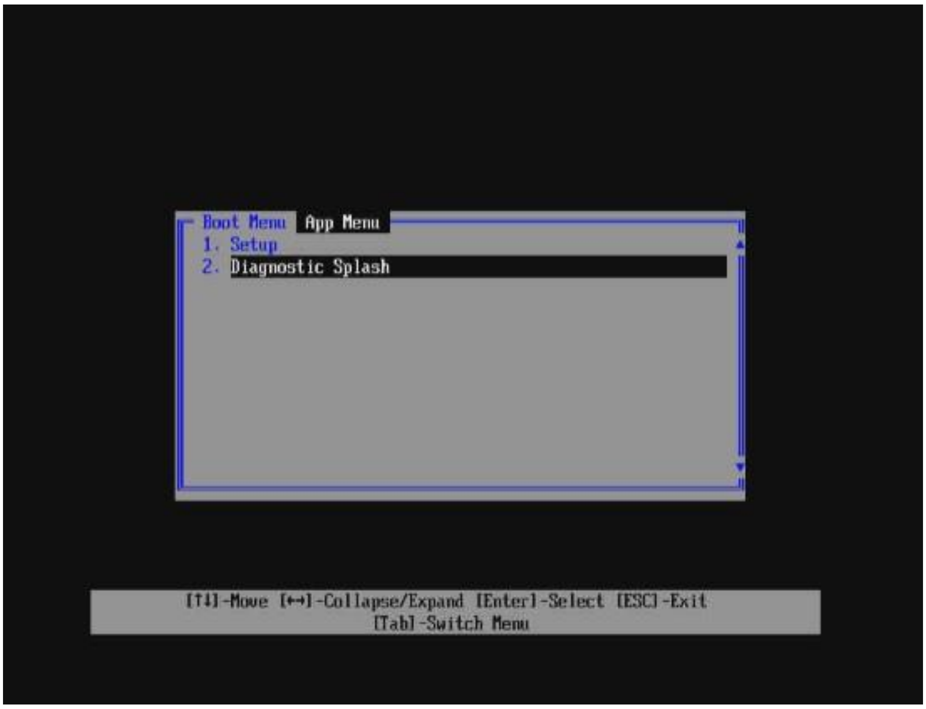

Figure 8: Diagnostic Splash screen option

3.1.4.6 Diagnostic Summary Screen

Menu Item	Description	Default Config	Options
Diagnostic Summary Screen	Enable/Disable Diagnostic Summary Screen display when booting. The figure 9 shows an example of a summary screen.	Disabled	[Enabled] [Disabled]

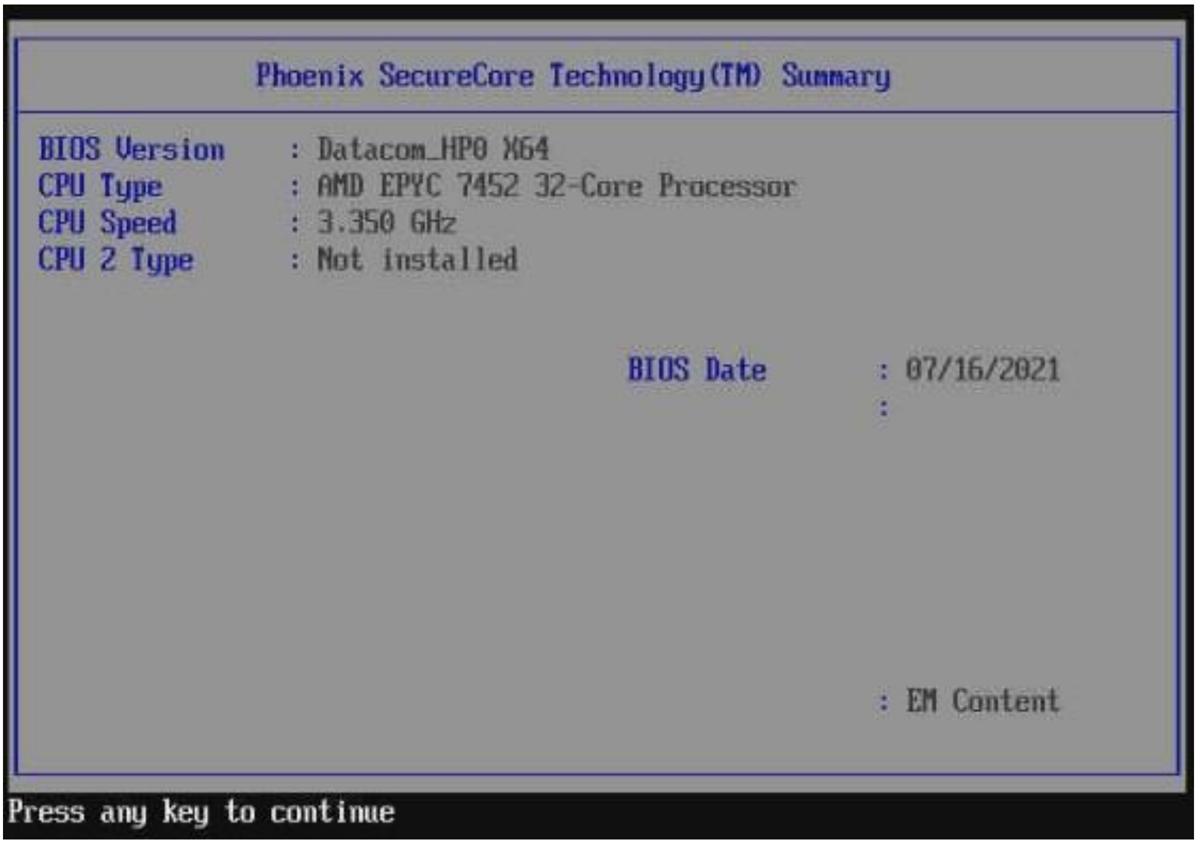

```
                  Phoenix SecureCore Technology(TM) Summary

 BIOS Version       : Datacom_HP0 X64
 CPU Type           : AMD EPYC 7452 32-Core Processor
 CPU Speed          : 3.350 GHz
 CPU 2 Type         : Not installed

                                     BIOS Date          : 07/16/2021
                                                        :



                                                        : EM Content

Press any key to continue
```

Figure 9: Diagnostic Summary Screen Example

3.1.4.7 Console Redirection

The console redirection functionality is used to attach the CPU serial communication to the BMC. Using this feature, it is possible to access the CPU’s console from BMC web interface, by means of the “Serial over LAN console” option. Although the serial terminal configurations are available in the Console Redirection sub-menus, it is recommended to keep the default configuration for this feature to work properly with BMC.

The console redirection feature is disabled by default. Enabling it may slow down the navigation in the BIOS menu and the access to the virtual media. So it is recommended to keep it disabled unless it is necessary to use this functionality.

Menu Item	Description	Default Config	Options
Console Redirection	Enable/Disable Console Redirection. When enabled, it allows the UEFI console redirection.	Disabled	[Enabled] [Disabled]
Terminal Type	Select the type of terminal to be used in the serial connection.	ANSI	[ANSI] [VT100] [VT100+] [UTF8] [TTY] [Linux] [XTerm] [VT400] [SC0]
Baudrate	Configure the baudrate of the serial communication.	115200	[9600] [19200] [38400] [57600] [115200]
Flow Control	Configure flow control of the serial communication.	None	[None] [RTS/CTS] [XON/XOFF]
Continue C.R. after POST	Enable or disable the console redirection functionality after the Operating System has already been loaded.	Disabled	[Enabled] [Disabled]

3.1.4.8 Allow Hotkey in S4 resume

Menu Item	Description	Default Config	Options
Allow Hotkey in S4 resume	Enable or disable hotkey detection when the system is resuming from Hibernate state.	Enabled	[Enabled] [Disabled]

3.1.4.9 UEFI Boot

The option below enables or disables the boot of the system in UEFI mode. Once the DM-SV01 server does not support boot in legacy mode, the UEFI Boot must be always enabled.

Menu Item	Description	Default Config	Options
UEFI Boot	Enables/Disables the UEFI boot mode.	Enabled	[Enabled] [Disabled]

3.1.4.10 Legacy Boot

The Legacy Boot mode can be enabled or disabled by means of the option below. In order for the legacy boot option to become available, it is necessary to activate the CSM (Compatibility Support Mode) option as explained in section 3.1.4.3 CSM Support.

Note: The DM-SV01 does not support booting in Legacy Mode. Although some BIOS releases may have this option available, the user must keep Legacy Boot with default settings.

Menu Item	Description	Default Config	Options
Legacy Boot	Enables/Disables the Legacy Boot mode.	Disabled	[Enabled] [Disabled]

3.1.4.11 Boot in Legacy Video Mode

The option below, when enabled, forces the video to run in Text Mode 3 at the end of BIOS POST. The option is available only if the Legacy Boot Mode is enabled (please refer to section 3.1.4.10 Legacy Boot).

Note: The DM-SV01 does not support booting in Legacy Mode. Although some BIOS releases may have this option available, the user must keep Boot in Legacy Video Mode with default settings.

Menu Item	Description	Default Config	Options
Boot in Legacy Video Mode	Enables/Disables the Legacy Boot mode for video.	Disabled	[Enabled] [Disabled]

3.1.4.12 Load OPROM

The option below is used to configure the option ROMs (OPROMs) loading behavior when legacy boot is enabled. It is possible to configure the loading of all OPROMs or the loading of “on demand” OPROMs, according to the boot device. The option is available only if the Legacy Boot Mode is enabled (please refer to section 3.1.4.10 Legacy Boot).

Note: The DM-SV01 does not support booting in Legacy Mode. Although some BIOS releases may have this option available, the user must keep Load OPROM with default settings.

Menu Item	Description	Default Config	Options
Load OPROM	Configures the OPROMs loading as “All” or “On Demand”.	On Demand	[All] [On Demand]

3.1.5 Network Stack (PXE Boot)

The Network Stack Option is used to allow the Operating System to boot through the network (PXE Boot).

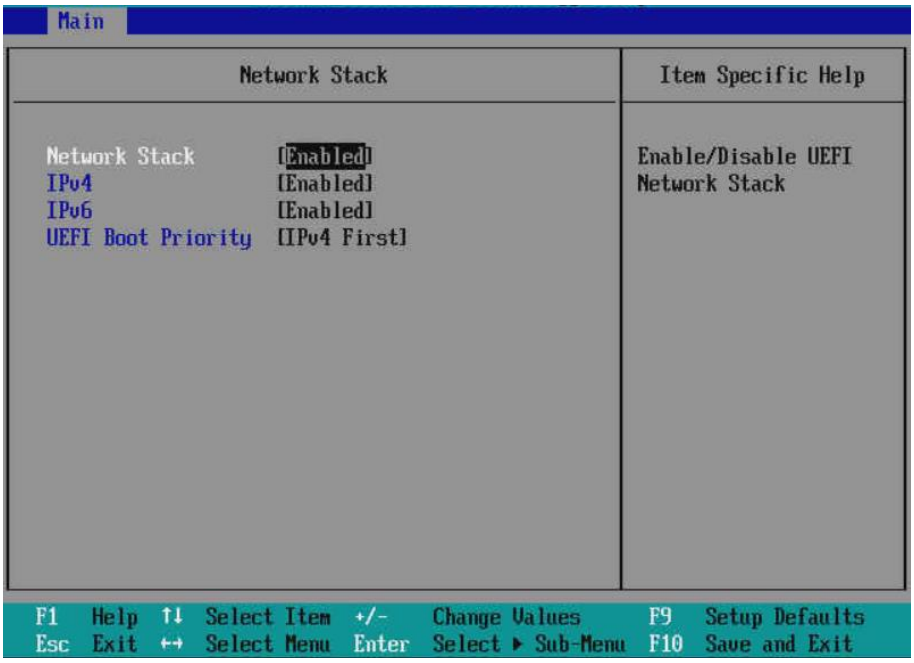
Main
Network Stack
Item Specific Help
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

```
Network Stack        [Enabled]
IPv4                 [Enabled]
IPv6                 [Enabled]
UEFI Boot Priority   [IPv4 First]
```

Enable/Disable UEFI Network Stack

Figure 10: Network Stack menu

Menu Item	Description	Default Config	Options
Network Stack	Enable/Disable UEFI Network Stack	Disabled	[Enabled] [Disabled]
IPv4	Enable/Disable IPv4 support for Network Stack.	N/A	[Enabled] [Disabled]
IPv6	Enable/Disable IPv6 support for Network Stack.	N/A	[Enabled] [Disabled]
UEFI Boot Priority	If both IPv4 and IPv6 protocols are Enabled for Network Stack, select which of them is the preferred one to be used at the boot process.	N/A	[IPv4 First] [IPv6 First]

3.1.6 Error Manager

The Error Manager menu allows the user to access the BIOS/UEFI error log information. The FW records information about errors identified in the system and makes them available to the user by means of the “View Error Manager Log” menu. Additionally, the user is able to clear the error log through the option “Clear Error Manager Log”.

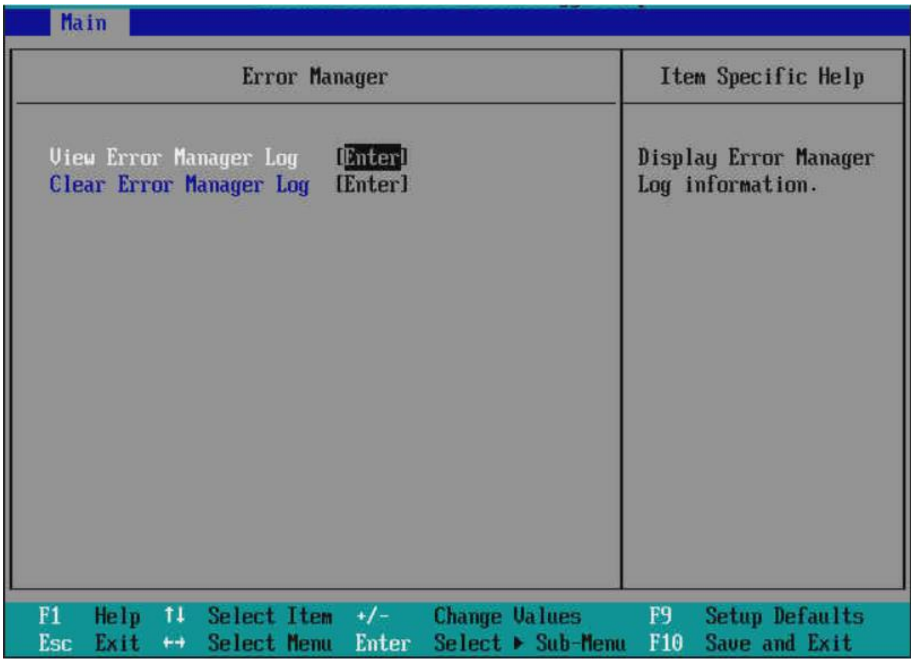

```
 Main
                 Error Manager                      Item Specific Help

   View Error Manager Log    [Enter]           Display Error Manager
   Clear Error Manager Log   [Enter]           Log information.

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 11: Error Manager Menu

Menu Item	Description	Default Config	Options
View Error Manager Log	Access the Error Manager to view the errors log.	N/A	[Enter]
Clear Error Manager Log	Clear all the Error Manager logs.	N/A	[Enter]

3.2 Advanced Menu

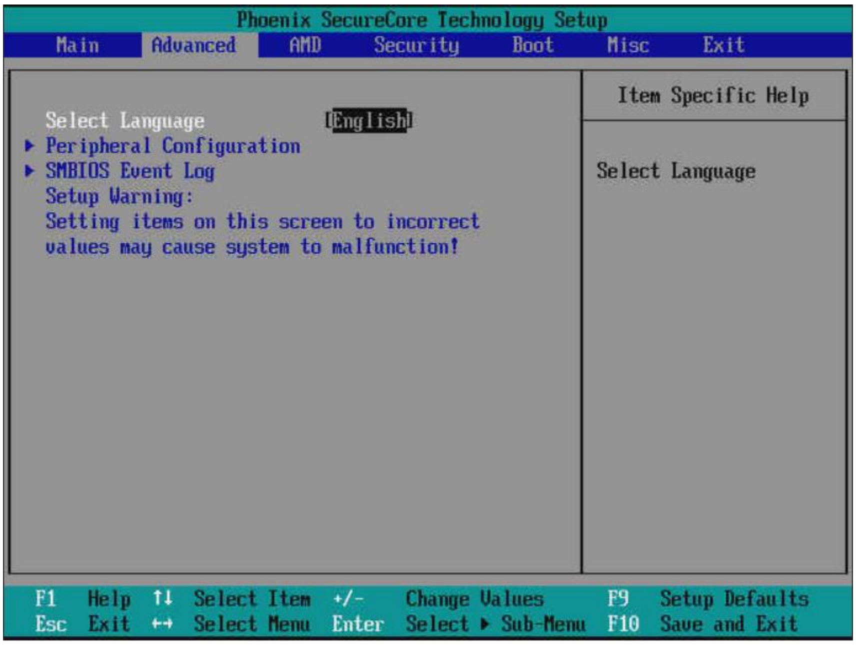

```
                     Phoenix SecureCore Technology Setup
   Main    Advanced    AMD    Security    Boot    Misc    Exit

    Select Language              [English]        Item Specific Help
  ▸ Peripheral Configuration
  ▸ SMBIOS Event Log                              Select Language
    Setup Warning:
    Setting items on this screen to incorrect
    values may cause system to malfunction!

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 12: Advanced Menu

3.2.1 Peripheral Configuration

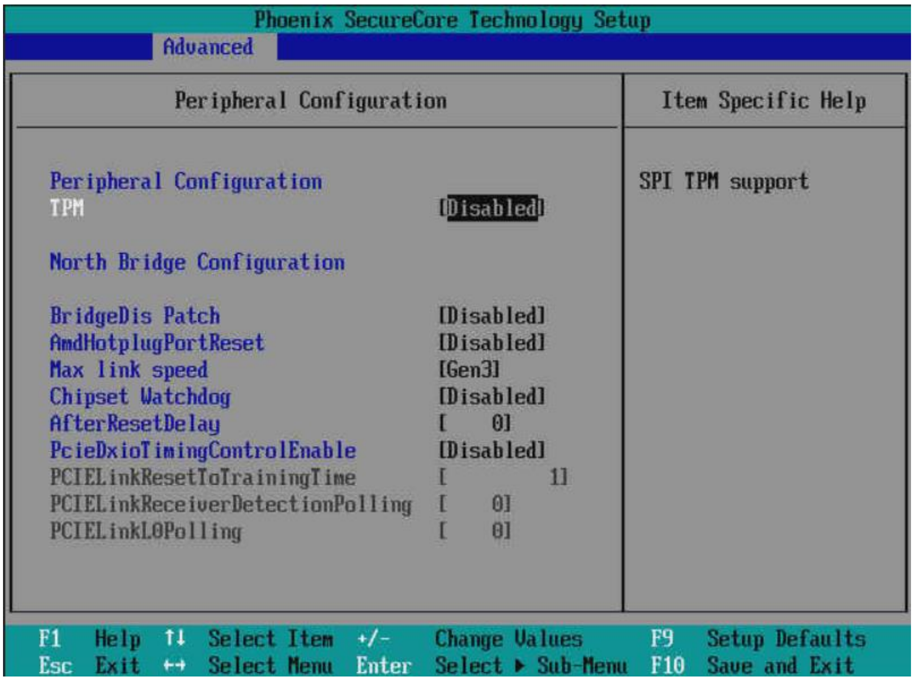

```
                    Phoenix SecureCore Technology Setup
              Advanced
  Peripheral Configuration                          Item Specific Help

  Peripheral Configuration                          SPI TPM support
  TPM                              [Disabled]

  North Bridge Configuration

  BridgeDis Patch                  [Disabled]
  AmdHotplugPortReset              [Disabled]
  Max link speed                   [Gen3]
  Chipset Watchdog                 [Disabled]
  AfterResetDelay                  [    0]
  PcieDxioTimingControlEnable      [Disabled]
  PCIELinkResetToTrainingTime      [          1]
  PCIELinkReceiverDetectionPolling [    0]
  PCIELinkL0Polling                [    0]

F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 13: Peripheral Configuration Menu

3.2.1.1 TPM

The TPM Option allows the user to enable or disable the SPI bus that communicates with the TPM device. This option must be set to “SPI TPM” to enable the use of the TPM for security configurations, as explained in section 3.4.8 Trusted Platform Module (TPM).

Menu Item	**Description**	**Default Config**	**Options**
TPM	Enable/Disable TPM’s SPI bus.	Disabled	[SPI TPM] [Disabled]

3.2.2 North Bridge Configuration

3.2.2.1 BridgeDis Patch

This option is for debug purposes only, please keep it with default settings.

Menu Item	**Description**	**Default Config**	**Options**
BridgeDis Patch	Enable/Disable the Bridge Disable Patch. Please keep it with default settings.	Disabled	[Enabled] [Disabled]

3.2.2.2 AmdHotplugPortReset

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
AmdHotplugPort Reset	Enable/Disable the Hotplug Port Reset. Please keep it with default settings.	Disabled	[Enabled] [Disabled]

3.2.2.3 Max link speed

Configures the maximum link speed capability for all PCIe devices of the server. There are two configurations available:

- Gen3: configures PCIe 3.0 (8GT/s) as the **maximum** speed for all PCIe slots. The user can insert cards specified as PCIe gen1, gen2 or gen3. When a lower speed card is inserted, the PCIe bus automatically negotiates and speeds the link down to match the card's speed capability. When the user inserts a PCIe gen4 card, it will be automatically limited down to gen3 link speed.
- Gen4: configures PCIe 4.0 (16GT/s) as the **maximum** speed for all PCIe slots. The user can insert cards specified as PCIe gen1, gen2, gen3 or gen4. When a lower speed card is inserted, the PCIe bus automatically negotiates and speeds the link down to match the card's speed capability.

Whenever possible, it is recommended to set the Max Link Speed to Gen3 in order to save power. The transmission and reception algorithms for each PCIe lane may consume more power when Gen4 speed is used. Since the processor package has a power dissipation limit, using Gen4 may result in less power available for the processing cores.

Note: DM-SV01 server does not support PCIe gen4 for the E1.S disk slots of the 2xE1.S adapter card.

Menu Item	Description	Default Config	Options
Max link speed	Configures the maximum link speed capability for all PCIe devices of the server.	Gen3	[Gen3] [Gen4]

3.2.2.4 Chipset Watchdog

This option is used to disable the watchdog, in order to allow the installation of the Oracle Linux operating system.

Menu Item	Description	Default Config	Options
Chipset Watchdog	Enable/Disable the watchdog, in order to allow the installation of the Oracle Linux operating system.	Disabled	[Enabled] [Disabled]

3.2.2.5 AfterResetDelay

Configures the delay between the release of the PCIe reset signal and the start of the communication, in milliseconds. It is recommended to keep this field with the default value.

Menu Item	Description	Default Config	Options
AfterResetDelay	Time in milliseconds between the release of the reset signal and the start of the communication of the PCIe bus. Acceptable range is from 0 up to 65535ms.	0	[<value in milliseconds>]

3.2.2.6 PCieDxioTimingControlEnable

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
PcieDxioTimingControlEnable	Enable/Disable the menus for configuring PCIe timing parameters.	Disabled	[Enabled] [Disabled]

3.2.2.7 PCIELinkResetToTrainingTime

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
PCIELinkResetToTrainingTime	Time in microseconds between the release of the reset signal and the start of the PCIe bus training. Acceptable range is from 0 up to 2,147,483,647us.	Disabled	[<value in microseconds>]

3.2.2.8 PCIELinkReceiverDetectionPolling

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
PCIELinkReceiverDetectionPolling	Maximum time from the beginning of the PCIe training until the device is detected successfully or timeout. Acceptable range is from 0 up to 10,000ms.	Disabled	[<value in milliseconds>]

3.2.2.9 PCIELinkL0Polling

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
PCIELinkL0Polling	Maximum time for the PCIe device to be in L0 state after detection has been completed. Acceptable range is from 0 up to 10,000ms.	Disabled	[<value in milliseconds>]

3.2.3 SMBIOS Event Log

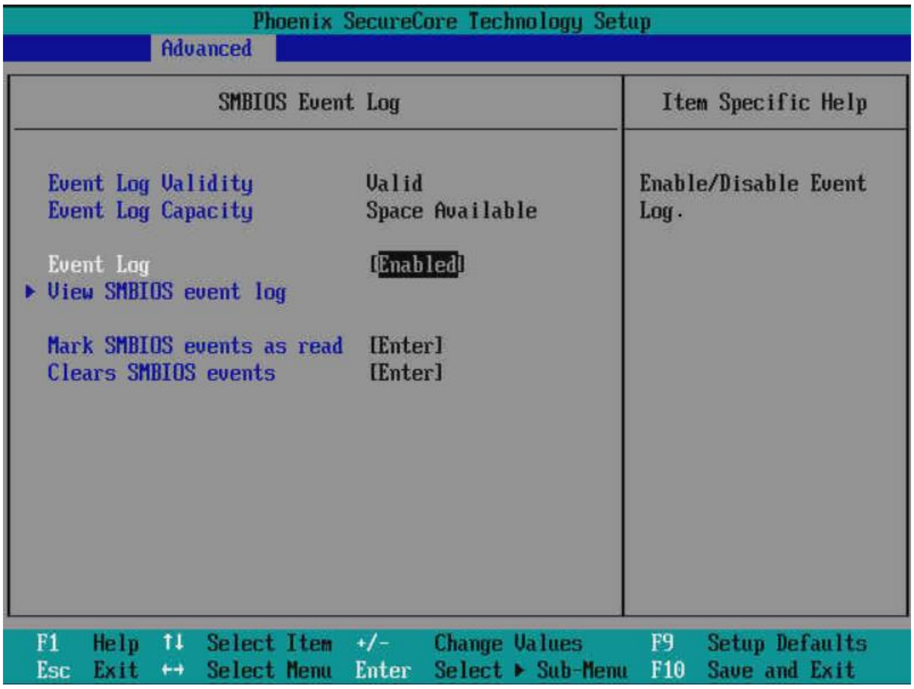

```
                    Phoenix SecureCore Technology Setup
                      Advanced

            SMBIOS Event Log                     Item Specific Help

   Event Log Validity          Valid             Enable/Disable Event
   Event Log Capacity          Space Available   Log.

   Event Log                   [Enabled]
 ▸ View SMBIOS event log

   Mark SMBIOS events as read  [Enter]
   Clears SMBIOS events        [Enter]

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 14: SMBIOS Event Log Menu

3.2.3.1 View SMBIOS event log

Access this menu to view all the recorded SMBIOS event logs.

Menu Item	Description	Default Config	Options
View SMBIOS event log	View all the recorded SMBIOS event logs.	N/A	N/A

3.2.3.2 Mark SMBIOS events as read

Mark all SMBIOS logs as read. After confirming this operation, the events already marked as read will not be displayed in the event log anymore.

Menu Item	Description	Default Config	Options
Mark SMBIOS events as read	Mark all SMBIOS events as read.	N/A	Confirm [Yes] [No]

3.2.3.3 Clear SMBIOS events

Clear all SMBIOS logs. After confirming this operation, the events will be completely erased from the event log.

Menu Item	Description	Default Config	Options
Clear SMBIOS events	Clear all the recorded SMBIOS event logs.	N/A	Confirm [Yes] [No]

3.3 AMD Menu

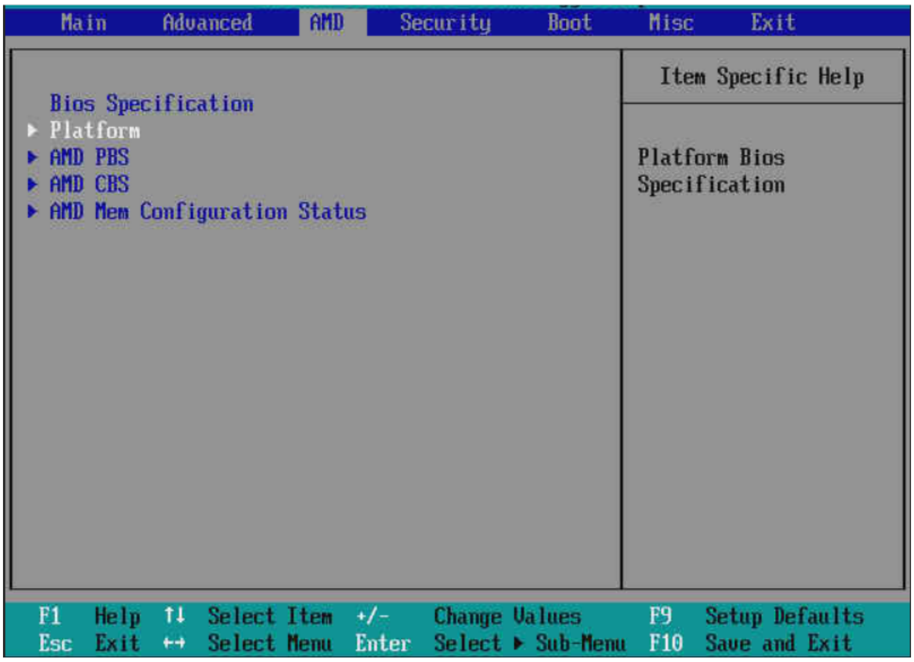

```
  Main     Advanced    AMD    Security    Boot    Misc    Exit

                                           Item Specific Help
   Bios Specification
 ▶ Platform                               Platform Bios
 ▶ AMD PBS                                Specification
 ▶ AMD CBS
 ▶ AMD Mem Configuration Status

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu  F10  Save and Exit
```

Figure 15: AMD Menu

3.3.1 Platform

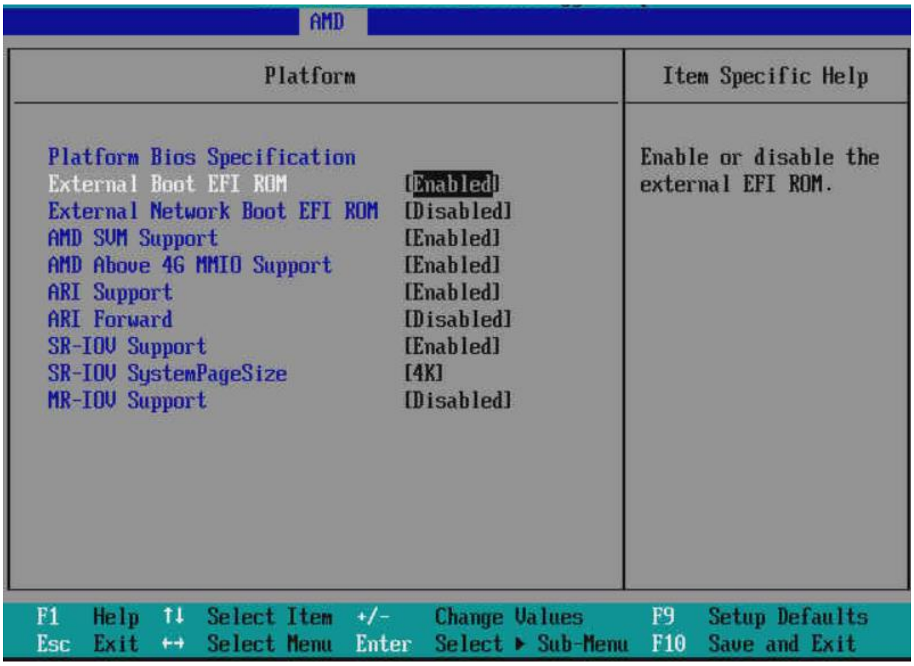

```
AMD

Platform                                   Item Specific Help

Platform Bios Specification                Enable or disable the
External Boot EFI ROM          [Enabled]   external EFI ROM.
External Network Boot EFI ROM  [Disabled]
AMD SVM Support                [Enabled]
AMD Above 4G MMIO Support      [Enabled]
ARI Support                    [Enabled]
ARI Forward                    [Disabled]
SR-IOV Support                 [Enabled]
SR-IOV SystemPageSize          [4K]
MR-IOV Support                 [Disabled]

F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 16: Platform Menu

3.3.1.1 External Boot EFI ROM

This option is used to enable or disable the initialization of EFI Option ROMs during boot. If set to disabled, the UEFI will not load nor execute any Option ROM code from plugin devices.

Menu Item	Description	Default Config	Options
External Boot EFI ROM	Enables/Disables the initialization of external EFI Option ROMs during the boot process.	Enabled	[Enabled] [Disabled]

3.3.1.2 External Network Boot EFI ROM

This option is used to enable or disable the initialization of EFI LAN Option ROMs during boot. If set to disabled, the UEFI will not load nor execute any network-related Option ROM code from plugin devices.

Menu Item	Description	Default Config	Options
External Network Boot EFI ROM	Enables/Disables the initialization of external LAN EFI Option ROMs during the boot process.	Disabled	[Enabled] [Disabled]

3.3.1.3 AMD SVM Support

The AMD SVM (Secure Virtual Machine) is the proprietary virtualization technology from AMD. It allows the system to run virtualization environments (hypervisors and virtual machines). This option may be used to enable or disable the virtualization technology.

For additional details about AMD SVM, please refer to the AMD doc (8) “AMD64 Architecture Programmer’s Manual Volume 2: System Programming”.

Menu Item	Description	Default Config	Options
AMD SVM Support	Enables/Disables the SVM (Secure Virtual Machine) virtualization features.	Enabled	[Enabled] [Disabled]

3.3.1.4 AMD Above 4G MMIO Support

The “Above 4G MMIO Support” option allows the system to allocate more than 4GB of memory address space to MMIO for PCIe devices. The 4GB maximum address space is a legacy limitation from 32-bit devices or Operating Systems (2^{32} = 4GiB), so keeping this option enabled allows the PCIe device to access the 64 bit address space, above the 4GB limitation.

Menu Item	Description	Default Config	Options
AMD Above 4G MMIO Support	Allows the system to enable/disable the allocation of more than 4GB of memory address space to MMIO for PCIe devices	Enabled	[Enabled] [Disabled]

3.3.1.5 ARI Support

The ARI (Alternative Routing ID) allows the PCIe devices to use 8-bit function numbers, instead of the default 3-bit function numbers. This allows the PCIe device to provide more than the default value of 8 Virtual Functions (VFs), so that it can be virtualized to a higher number of virtual guests.

The ARI Support must be enabled when the user wants to activate the SR-IOV functionality in virtualization environments. Please refer to section 3.3.1.7 SR-IOV Support for additional details about the SR-IOV functionality.

Menu Item	Description	Default Config	Options
ARI Support	Enables/Disables the ARI (Alternative Routing ID) feature.	Enabled	[Enabled] [Disabled]

3.3.1.6 ARI Forward

The ARI Forward option allows the ARI feature to be forwarded to endpoints after passing through a PCIe switch. This menu option can enable/disable this feature.

Menu Item	Description	Default Config	Options
ARI Forward	Enables/Disables the forwarding of the ARI feature to endpoints residing after a PCIe switch.	Disabled	[Enabled] [Disabled]

3.3.1.7 SR-IOV Support

The Single Root I/O Virtualization (SR-IOV) feature allows a single PCIe endpoint device to be divided into multiple virtual functions (VF), which can be accessed by several virtual machines from a host in a virtualization environment. The virtual machines can use the virtual functions (VFs) to have access to specific hardware functionalities of the endpoint PCIe device. The main purpose of the use of VFs is to allow the virtual machine to send and receive data directly to/from the PCIe device, without the need to make use of the hypervisor software layers. This menu option enables or disables this feature.

Note: in order to allow a PCIe endpoint device to use SR-IOV, the user must also enable the Alternative Routing-ID Interpretation (ARI) feature, as explained in section 3.3.1.5 ARI Support.

Note: each PCIe endpoint device may have its own SR-IOV enable/disable configuration. Therefore, in order to enable the SR-IOV feature, the user needs to check the specific PCIe card setting, by accessing its configuration options in the miscellaneous menu (please refer to section 3.6 Misc Menu for details).

Menu Item	Description	Default Config	Options
SR-IOV Support	Enables/Disables the Single Root I/O Virtualization (SR-IOV) feature.	Enabled	[Enabled] [Disabled]

3.3.1.8 SR-IOV SystemPageSize

This option defines the page size the system will use to map the PCIe memory addresses of the Virtual Functions (VFs) when the SR-IOV functionalities are enabled.

Menu Item	Description	Default Config	Options
SR-IOV SystemPageSize	Defines the system page size for PCIe devices using SR-IOV functionality.	4K	[4K] [8K] [64K] [256K] [1M] [4M]

3.3.1.9 MR-IOV Support

The MR-IOV (Multi-Root I/O Virtualization) is a feature similar to the SR-IOV (3.3.1.7 SR-IOV Support). However, the MR-IOV feature allows a single PCIe endpoint device to be used as multiple virtual functions (VF) by virtual machines residing in different hosts, and not only in the host which has the PCIe device attached to it, as in the SR-IOV feature. The virtual machines can then use the virtual functions (VFs) to access specific hardware functionalities of the endpoint PCIe device.

This menu option enables or disables this feature.

Note: in order to allow a PCIe endpoint device to use MR-IOV, the user must also enable the Alternative Routing-ID Interpretation (ARI) feature, as explained in section 3.3.1.5 ARI Support, and ARI Forward, as shown in section 3.3.1.6 ARI Forward.

Note: each PCIe endpoint device may have its own MR-IOV enable/disable configuration. Therefore, in order to enable the MR-IOV feature, the user needs to check the specific PCIe card setting, by accessing its configuration options in the miscellaneous menu (please refer to section 3.6 Misc Menu for details).

Menu Item	Description	Default Config	Options
MR-IOV Support	Enables/Disables the Multi-Root I/O Virtualization (MR-IOV) feature.	Disabled	[Enabled] [Disabled]

3.3.2 AMD PBS

3.3.2.1 PCIe slot bifurcation

The DM-SV01 allows the user to configure bifurcation options for the PCIe slots, so that it is possible to use PCIe cards which takes advantage of the slot split. An example is a standard PCIe to NVMe adapter card, which is connected to a single PCIe slot and can give access to two or four NVMe drives by using the slot bifurcation.

The bifurcation options available in the PBS menu depends on the riser card type assembled in the DM-SV01 server, once each riser card provides different PCIe slots configuration and, therefore, different bifurcation options. The bifurcation options available for each riser card are shown in the table below.

Risercard type	Slot	Default Config*	Options*
Riser card x16+x8	3	1x16	[1x16] [2x8] [4x4]
	2	1x8	[1x8] [2x4]
	1	Unavailable	-

Riser card 3x8	3	1x8	[1x8] [2x4]
	2	1x8	[1x8] [2x4]
	1	1x8	[1x8] [2x4]

* Please see below the meaning of each config option:

- 1x16: 1 PCIe slot, width = x16 (no bifurcation)
- 1x8: 1 PCIe slot, width = x8 (no bifurcation)
- 2x8: 2 PCIe slots, width = x8
- 2x4: 2 PCIe slots, width = x4
- 4x4: 4 PCIe slots, width = x4

The DM-SV01 BIOS/UEFI is capable of detecting which riser card type is connected to the system and automatically adapts the PBS menu with the information of that specific card and the corresponding bifurcation options/default settings. When no riser card is equipped in the system, the bifurcation options are not shown.

The figures below show the PBS menu screen for each riser card assembly possibility. The Figure 17 shows the PBS menu screen when no riser card is equipped in the DM-SV01 server; the Figure 18 shows the PBS menu screen when the riser card x16+x8 is equipped in the DM-SV01 server; and the Figure 19 shows the PBS menu screen when the riser card 3x8 is equipped in the DM-SV01 server.

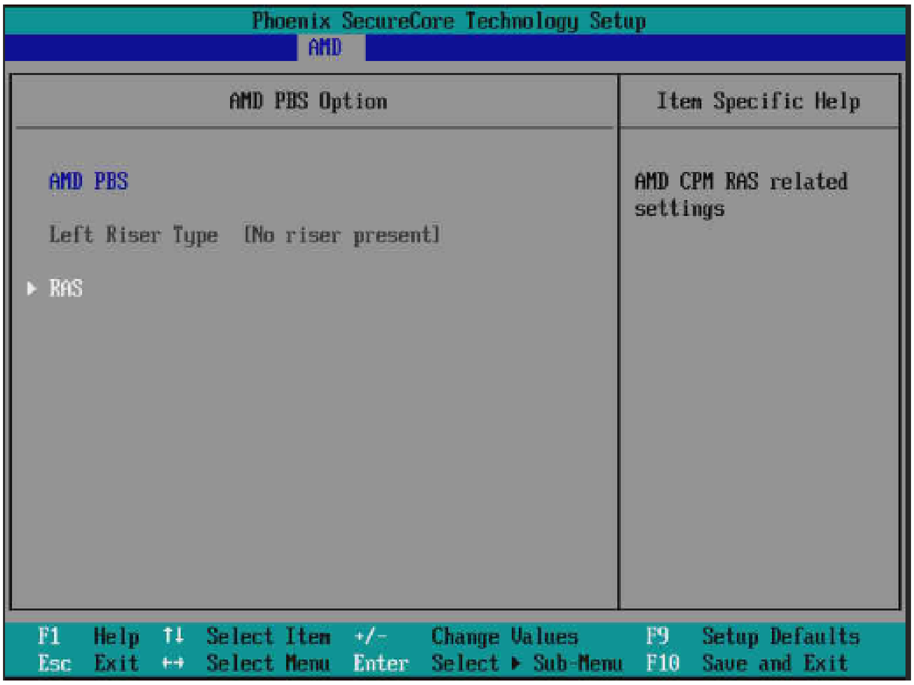
Phoenix SecureCore Technology Setup
AMD
AMD PBS Option
Item Specific Help
AMD PBS
▸ RAS
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Left Riser Type [No riser present]

AMD CPM RAS related settings

Figure 17: PBS bifurcation setting when no riser card is equipped

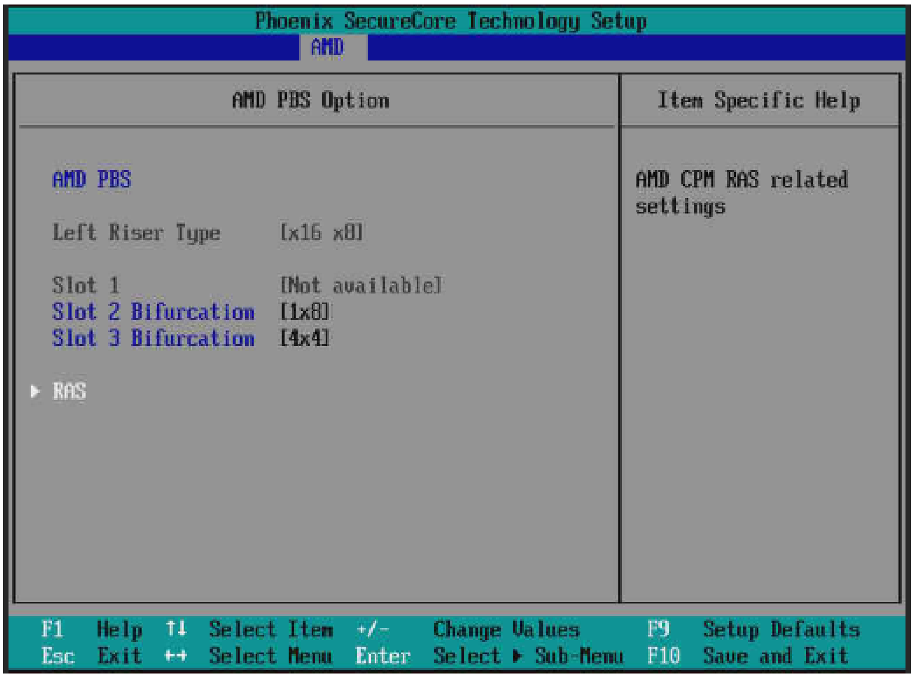
Phoenix SecureCore Technology Setup
AMD
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

```
AMD PBS Option                            Item Specific Help

AMD PBS                                   AMD CPM RAS related
                                          settings
Left Riser Type       [x16 x8]

Slot 1                [Not available]
Slot 2 Bifurcation    [1x8]
Slot 3 Bifurcation    [4x4]

▸ RAS
```

Figure 18: PBS bifurcation setting when riser card x16+x8 is equipped

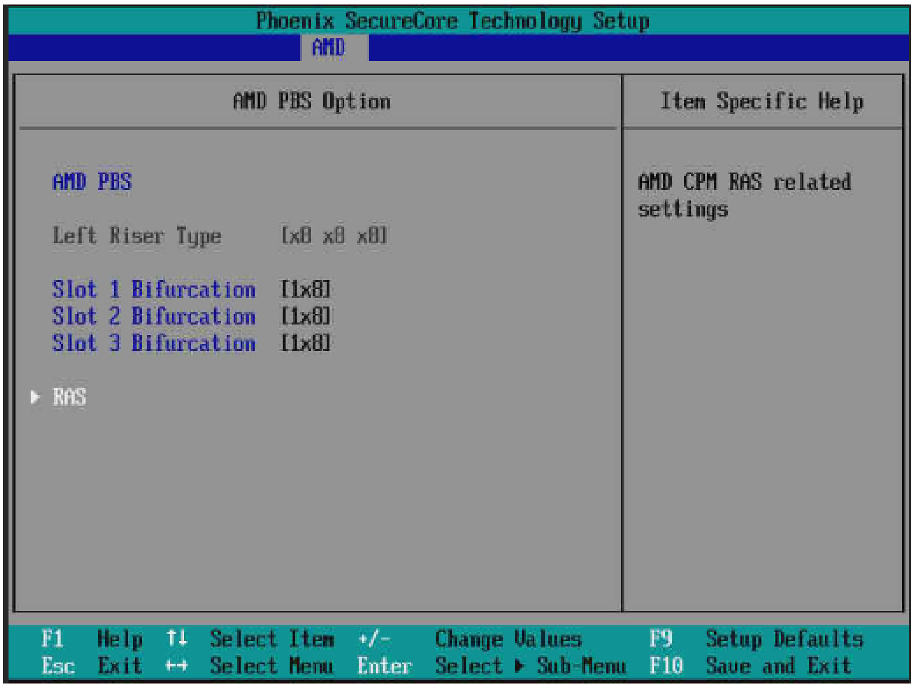
Phoenix SecureCore Technology Setup
AMD
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

```
AMD PBS Option                            Item Specific Help

AMD PBS                                   AMD CPM RAS related
                                          settings
Left Riser Type       [x8 x8 x8]

Slot 1 Bifurcation    [1x8]
Slot 2 Bifurcation    [1x8]
Slot 3 Bifurcation    [1x8]

▸ RAS
```

Figure 19: PBS bifurcation setting when riser card 3x8 is equipped

3.3.2.1.1 PCIe slot bifurcation when using the DM-SV01 2xE1.S Adapter Card

The DM-SV01 2xE1.S Adapter card is a module that connects to the riser card PCIe slot (x16 or x8) and provides two E1.S slots to connect NVMe disks. The 2xE1.S adapter card requires the proper bifurcation configured in the PCIe slot in order to correctly distribute the PCIe lanes between the two E1.S ports.

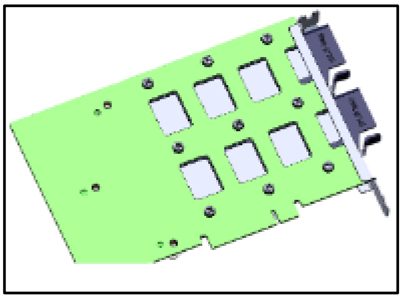

Figure 20: 2xE1.S PCIe Adapter Card

When the DM-SV01 system is equipped with the DM-SV01 2xE1.S Adapter Card, the BIOS/UEFI is capable of detecting this card and automatically configures the PCIe slot bifurcation for the corresponding slot. When this automatic configuration is performed by BIOS/UEFI, the bifurcation setting for the corresponding slot becomes grayed, indicating it is disabled for user intervention.

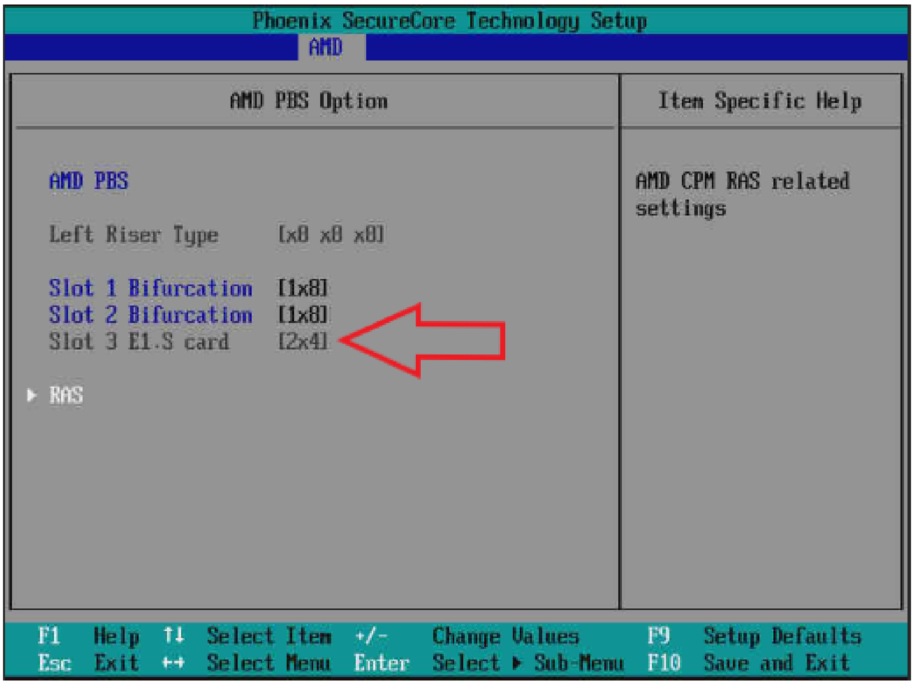

```
Phoenix SecureCore Technology Setup
AMD

AMD PBS Option                          Item Specific Help

AMD PBS                                 AMD CPM RAS related
                                        settings
Left Riser Type        [x8 x8 x8]

Slot 1 Bifurcation     [1x8]
Slot 2 Bifurcation     [1x8]
Slot 3 E1.S card       [2x4]

▸ RAS

F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 21: bifurcation option automatically configured by BIOS/UEFI

3.3.2.2 RAS (Reliability, Availability and Serviceability)

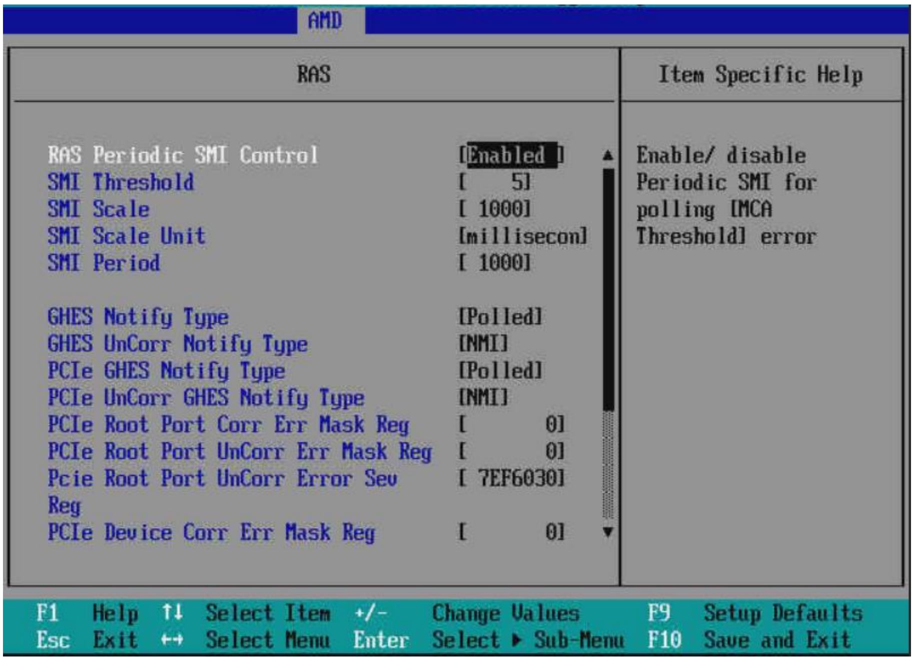

```
                              AMD
                  RAS                          Item Specific Help

   RAS Periodic SMI Control          [Enabled ]   Enable/ disable
   SMI Threshold                     [     5]     Periodic SMI for
   SMI Scale                         [ 1000]      polling [MCA
   SMI Scale Unit                    [millisecon]  Threshold] error
   SMI Period                        [ 1000]

   GHES Notify Type                  [Polled]
   GHES UnCorr Notify Type           [NMI]
   PCIe GHES Notify Type             [Polled]
   PCIe UnCorr GHES Notify Type      [NMI]
   PCIe Root Port Corr Err Mask Reg  [       0]
   PCIe Root Port UnCorr Err Mask Reg [      0]
   Pcie Root Port UnCorr Error Sev   [ 7EF6030]
   Reg
   PCIe Device Corr Err Mask Reg     [       0]

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 22: AMD PBS RAS Menu

3.3.2.2.1 RAS Periodic SMI Control

The SMI (System Management Interrupt) is triggered whenever the CPU detects an error event in the MCA (Machine Check Architecture), which is an interface responsible for reporting errors related to CPU or hardware components. When PFEH (section 3.3.3.1.5 Platform First Error Handling) is enabled, the CPU will send the SMI to the platform firmware CPM (Common Platform Module). The CPM is the AMD proprietary FW code residing in BIOS/UEFI and capable of handling these types of error events.

Using this mechanism, the SMI will be triggered for every single error detected by the CPU architecture. However, if a faulty HW is constantly experiencing certain error types, like deferred or correctable errors, this causes the SMI to be triggered at a high rate, causing the system performance to degrade substantially.

The RAS Periodic SMI Control option, when enabled, allows the CPU to define a threshold of SMIs allowed during a predefined period of time. If the threshold is reached, then the SMIs stop being sent for every single error, and start to be generated once for a period instead (polling mode).

The Menu items below are used to configure these thresholds for periodic SMI generation. The “RAS Periodic SMI Control” enables or disables this feature. The “SMI Threshold” is the maximum number of SMIs that can be generated in the time window defined in the “SMI Scale” option, before triggering the periodic SMI generation. The “SMI Period”, on the other hand, is the time interval that the SMI will be sent after entering in the polling mode after periodic SMI control has been previously triggered.

Menu Item	Description	Default Config	Options
RAS Periodic SMI Control	Enables/Disables the Periodic SMI Control feature.	Enabled	[Enabled] [Disabled]

SMI Threshold	Defines the maximum number of times the SMI can be generated within the time window defined in the “SMI Scale” option before triggering the polling mode from periodic SMI Control. Acceptable range: 0 up to 65535.	5	[<value>]
SMI Scale	Defines the time window over which the “SMI Threshold” option will be operating. Acceptable range: 0 up to 32767.	1000	[<value>]
SMI Scale Unit	Defines the time unit for the value defined in the “SMI Scale” option.	millisecond	[millisecond] [second] [minute]
SMI Period	Defines the period of time that the SMI will be generated after entering in polling mode through the Periodic SMI Control feature. Acceptable range: 0ms up to 32767ms.	1000ms	[<value in milliseconds>]

The GHES (Generic Hardware Error Source) is a structure from BIOS/UEFI FW to provide additional information to the system about an error occurring in hardware. When the hardware detects an error, it reports this error to the firmware, which performs additional verifications and then outputs the complete error information to the OS by means of the GHES.

The options below configure some parameters related to the GHES, and define its behavior for some specific situations.

Menu Item	Description	Default Config	Options
GHES Notify Type	Defines the notification type for deferred or corrected errors: polled or using System Control Interrupt (SCI).	Polled	[Polled] [SCI]
GHES UnCorr Notify Type	Defines the notification type for uncorrected errors: polled or using Non-Maskable Interrupt (NMI).	NMI	[Polled] [NMI]
PCIe GHES Notify Type	Defines the notification type for PCIe corrected errors: polled or using System Control Interrupt (SCI).	Polled	[Polled] [SCI]
PCIe UnCorr GHES Notify Type	Defines the notification type for PCIe uncorrected errors: polled or using Non-Maskable Interrupt (NMI).	NMI	[Polled] [NMI]

The items below configure masks for PCIe AER (Advanced Error Reporting) by writing directly in some AMD CPU internal registers. The content of each register is beyond the context of this manual and these registers are used only for debugging purposes, so that it is not expected that the user changes these values.

Menu Item	Description	Default Config	Options
PCIe Root Port Corr Err Mask Reg	Mask register for PCIe root port AER corrected errors.	0	[<value in hex>]

PCIe Root Port UnCorr Err Mask Reg	Mask register for PCIe root port AER uncorrected errors.	0	[<value in hex>]
PCIe Root Port UnCorr Err Sev Reg	Mask register for PCIe root port AER uncorrected errors severity setting (fatal or non-fatal error).	7EF6030	[<value in hex>]
PCIe Device Corr Err Mask Reg	Mask register for PCIe device AER corrected errors.	0	[<value in hex>]
PCIe Device UnCorr Err Mask Reg	Mask register for PCIe device AER uncorrected errors.	100000	[<value in hex>]
PCIe Device UnCorr Err Sev Reg	Mask register for PCIe device AER uncorrected errors severity setting (fatal or non-fatal error).	7EF6030	[<value in hex>]

The CCIX (Cache Coherent Interconnect for Accelerators) is an interface responsible for providing an interconnection between CPUs and hardware accelerators, using cache coherence to control memory sharing between them.

The options below configure some parameters related to the CCIX and define its behavior for some specific situations. A deferred error indicates that the data has been corrupted but it was not consumed by the system.

Menu Item	**Description**	**Default Config**	**Options**
CCIX GHES Deferred Err Notify Type	Defines the notification type for CCIX deferred errors: polled or using System Control Interrupt (SCI).	Polled	[Polled] [SCI]
CCIX GHES Corrected Err Notify Type	Defines the notification type for CCIX corrected errors: polled or using System Control Interrupt (SCI).	Polled	[Polled] [SCI]

The DRAM Post Package Repair (PPR) is a feature that allows a failed DRAM row to be replaced by another one, providing a recovery method for a failed portion of the memory. The PPR mechanism is handled by BIOS/UEFI FW.

The menu below allows the user to enable or disable this feature.

Menu Item	**Description**	**Default Config**	**Options**
DDR4 DRAM Hard Post Package Repair	Enable/Disable the DRAM Post Package Repair (PPR) feature.	Disabled	[Enabled] [Disabled]

The HEST (Hardware Error Status Table) is a structure used by the platform FW to transmit detailed information about hardware errors sources to the Operating System. The option below is used to enable or disable HEST for errors detected by the DMC (Deferred Machine Check) error source. A deferred error indicates that the data has been corrupted but it was not consumed by the system. So, the option below, when enabled, allows BIOS/UEFI to inform the OS about this type of error.

Menu Item	Description	Default Config	Options
HEST DMC Structure Support	Enable/Disable the HEST structure for DMC (Deferred Machine Check) errors.	Disabled	[Enabled] [Disabled]

3.3.2.3 SRIS mode debug

This option enables or disables the debugging of the SRIS (Separate Refclk with Independent SSC - Spread Spectrum Clocking) PCIe clock distribution architecture.

Menu Item	Description	Default Config	Options
SRIS mode debug	Enables/Disables the SRIS mode debug.	Auto	[Enabled] [Disabled]]Auto]

3.3.2.4 Skip interval

Configures the skip interval in the following order:

- SRNS (Gen2 and below); SRIS (Gen2 and below); SRNS (Gen3 and above); SRIS (Gen3 and above);

Menu Item	Description	Default Config	Options
Skip interval	Configures the skip interval for SRNS or SRIS.	1506 ; 144; 6050; 640	[1506 ; 144; 6050; 640] [1538 ; 154; 6068; 656] [1358 ; 128; 6032; 624] [1180 ; 112; 5996; 608]

3.3.2.5 LOWER_SKP_OS_GEN_SUPPORT

Configures the lower PCIe gen of the transmitter that will be handling skip (SKP) ordered sets.

Menu Item	Description	Default Config	Options
LOWER_SKP_OS_GEN_SUPPORT	Configures the lower PCIe gen of the transmitter that will be handling skip (SKP) ordered sets.	Disabled	[Disabled] [Gen1] [Gen2] [Gen3] [Gen4]

3.3.2.6 LOWER_SKP_OS_RCV_SUPPORT

Configures the lower PCIe gen of the receiver that will be handling skip (SKP) ordered sets.

Menu Item	Description	Default Config	Options
LOWER_SKP_OS_RCV_SUPPORT	Configures the lower PCIe gen of the receiver that will be handling skip (SKP) ordered sets.	Disabled	[Disabled] [Gen1] [Gen2] [Gen3] [Gen4]

3.3.2.7 SRIS autodetect

Enables or disables the autodetect function, in order to allow the system to automatically discover the PCIe clocking schema (SRNS, SRIS).

Menu Item	Description	Default Config	Options
SRIS autodetect	Enables/Disables the autodetect mode for SRIS PCIe clocking architecture.	Auto	[Enabled] [Disabled] [Auto]

3.3.2.8 SKP Interval Selection Mode

Configures how to define the SKP ordered set interval.

Menu Item	Description	Default Config	Options
SKP Interval Selection Mode	Configures how to define the SKP ordered set interval.	Dynamic SKP ordered set Interval Mode	[SKP ordered set Interval Lock Mode] [Dynamic SKP ordered set Interval Mode] [Far End nominal Empty Mode]

3.3.2.9 Autodetect Factor

Configures the SRIS autodetect factor.

Menu Item	Description	Default Config	Options
Autodetect Factor	Configures the SRIS autodetect factor.	1x	[1x] [0.95x] [0.9x] [0.85x]

3.3.3 AMD CBS

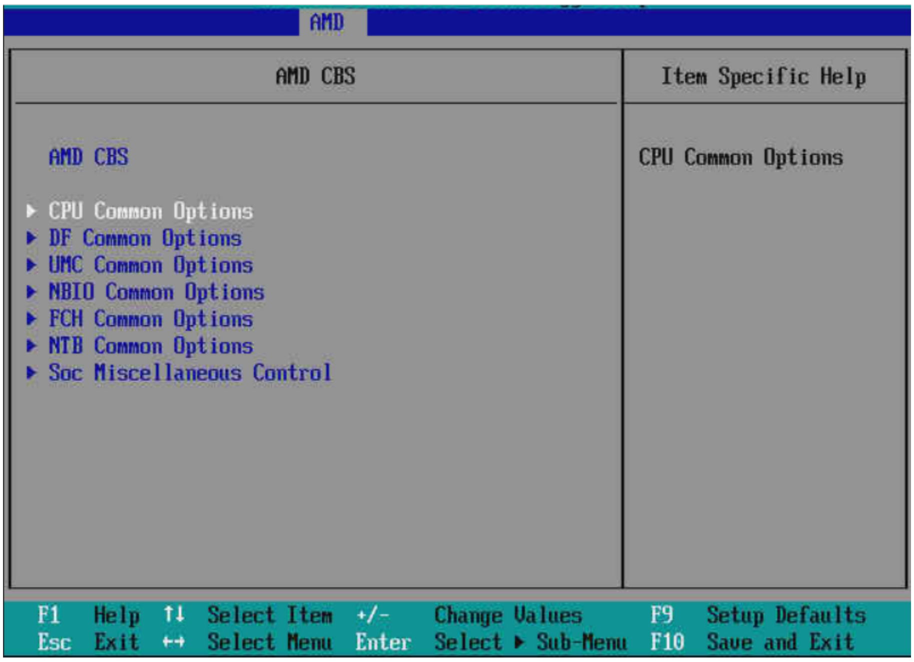

```
                              AMD
              AMD CBS                          Item Specific Help

   AMD CBS                                   CPU Common Options

 ▶ CPU Common Options
 ▶ DF Common Options
 ▶ UMC Common Options
 ▶ NBIO Common Options
 ▶ FCH Common Options
 ▶ NTB Common Options
 ▶ Soc Miscellaneous Control

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu  F10  Save and Exit
```

Figure 23: AMD CBS Menu

3.3.3.1 CPU Common Options

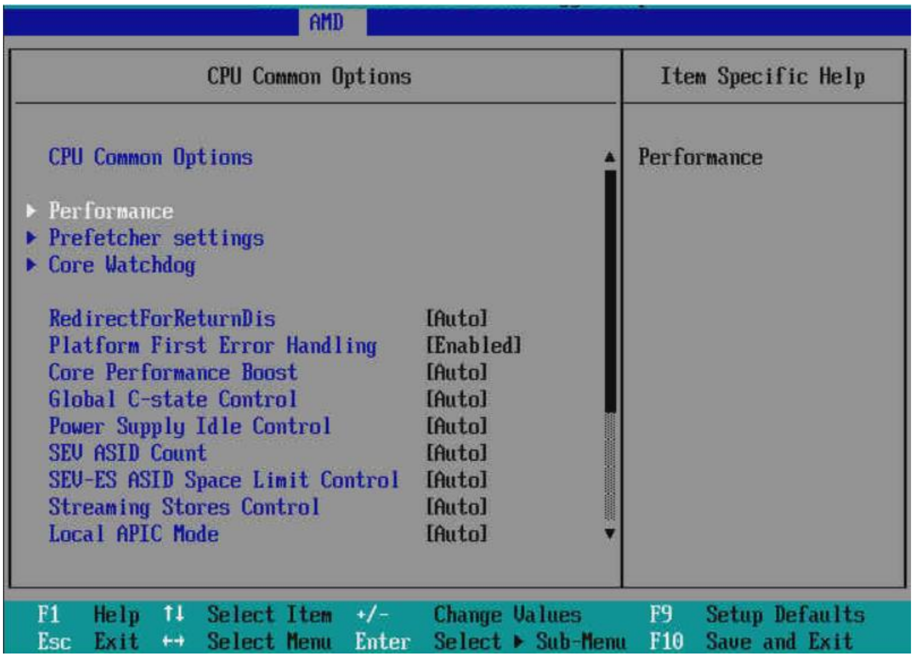

```
                              AMD
         CPU Common Options                    Item Specific Help

   CPU Common Options                        Performance

 ▶ Performance
 ▶ Prefetcher settings
 ▶ Core Watchdog

   RedirectForReturnDis          [Auto]
   Platform First Error Handling [Enabled]
   Core Performance Boost        [Auto]
   Global C-state Control        [Auto]
   Power Supply Idle Control     [Auto]
   SEV ASID Count                [Auto]
   SEV-ES ASID Space Limit Control [Auto]
   Streaming Stores Control      [Auto]
   Local APIC Mode               [Auto]

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu  F10  Save and Exit
```

Figure 24: CPU Common Options Menu

3.3.3.1.1 Performance

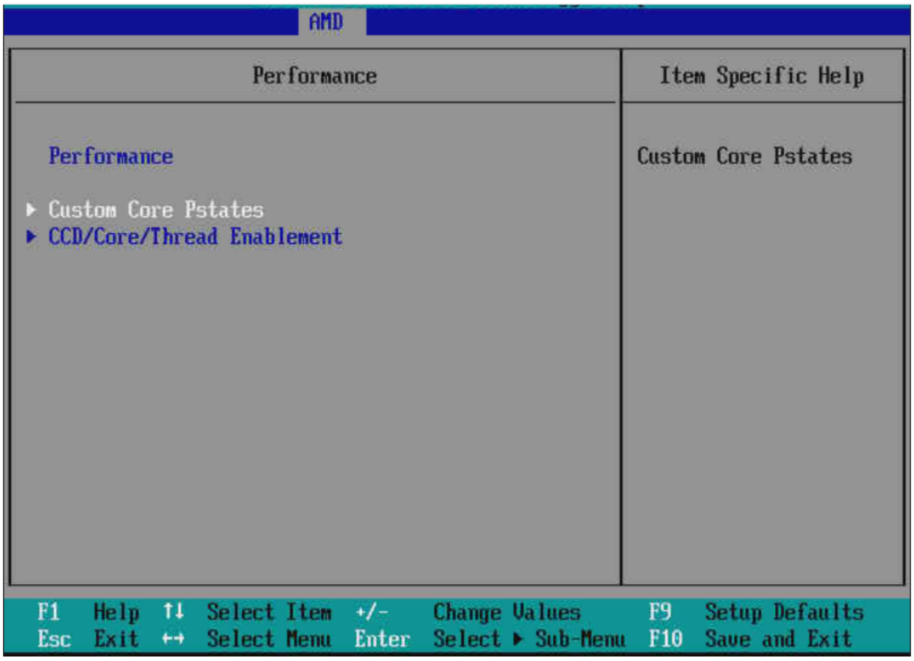

```
                         AMD
              Performance                     Item Specific Help

  Performance                               Custom Core Pstates

▸ Custom Core Pstates
▸ CCD/Core/Thread Enablement

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 25: Performance Menu

3.3.3.1.1.1 Custom Core Pstates

This option can be used to customize the Power States (Pstates) of the processor, defining which value of frequency and voltage they operate for which Pstate.

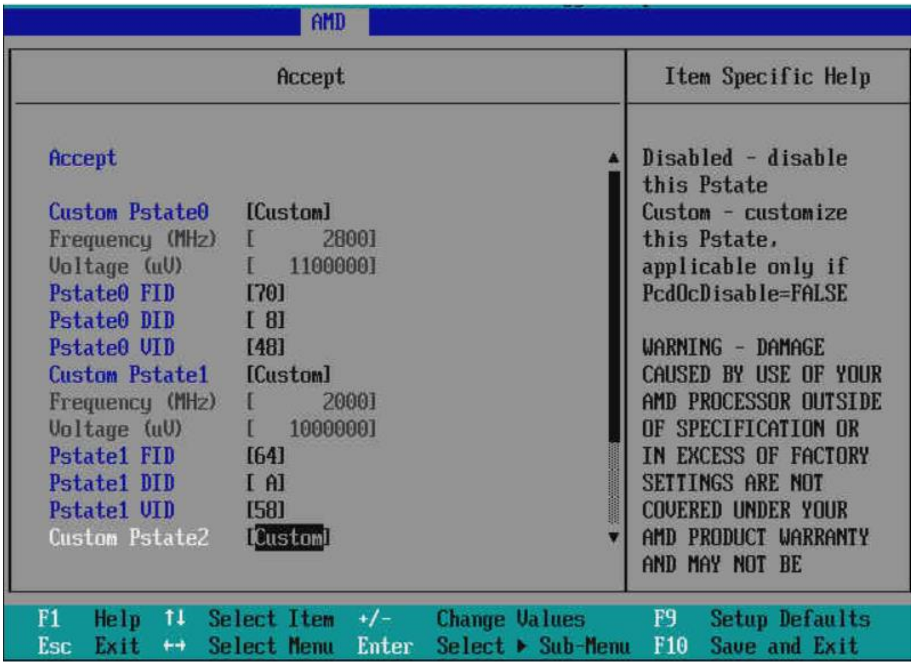

```
                         AMD
                Accept                         Item Specific Help

  Accept                                     Disabled - disable
                                             this Pstate
  Custom Pstate0      [Custom]               Custom - customize
  Frequency (MHz)     [      2800]           this Pstate,
  Voltage (uV)        [   1100000]           applicable only if
  Pstate0 FID         [70]                   PcdOcDisable=FALSE
  Pstate0 DID         [ 8]
  Pstate0 VID         [48]                   WARNING - DAMAGE
  Custom Pstate1      [Custom]               CAUSED BY USE OF YOUR
  Frequency (MHz)     [      2000]           AMD PROCESSOR OUTSIDE
  Voltage (uV)        [   1000000]           OF SPECIFICATION OR
  Pstate1 FID         [64]                   IN EXCESS OF FACTORY
  Pstate1 DID         [ A]                   SETTINGS ARE NOT
  Pstate1 VID         [58]                   COVERED UNDER YOUR
  Custom Pstate2      [Custom]               AMD PRODUCT WARRANTY
                                             AND MAY NOT BE

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 26: Custom Core Pstates Menu

CAUTION: It is not recommended to change the default values, once the processor can be damaged if the values are outside of the CPU specifications.

Menu Item	Description	Default Config	Options
Custom Pstate0	Defines if the CPU Pstate0 configuration is automatic or if a custom setting is used.	Auto	[Custom] [Auto] [Disabled]
Custom Pstate1	Defines if the CPU Pstate1 configuration is automatic or if a custom setting is used.	Auto	[Custom] [Auto] [Disabled]
Custom Pstate2	Defines if the CPU Pstate2 configuration is automatic or if a custom setting is used.	Auto	[Custom] [Auto] [Disabled]

When the Custom Pstate is active, the menu items below are shown.

Menu Item	Description	Default Config	Options
Pstate0 FID	Defines the value of FID (Frequency ID) to configure the CPU frequency for Pstate0. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	70	[<value in hex>]
Pstate0 DID	Defines the value of DID (Divisor ID) to configure the CPU frequency for Pstate0. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	8	[<value in hex>]
Pstate0 VID	Defines the value of VID (Voltage ID) to configure the CPU core voltage for Pstate0. The relationship between the hexadecimal values and the corresponding voltage levels are defined by AMD Serial VID Interface Specification.	48	[<value in hex>]
Pstate1 FID	Defines the value of FID (Frequency ID) to configure the CPU frequency for Pstate1. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	70	[<value in hex>]
Pstate1 DID	Defines the value of DID (Divisor ID) to configure the CPU frequency for Pstate1. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	8	[<value in hex>]
Pstate1 VID	Defines the value of VID (Voltage ID) to configure the CPU core voltage for Pstate1. The relationship between the hexadecimal values and the corresponding voltage levels are defined by AMD Serial VID Interface Specification.	48	[<value in hex>]
Pstate2 FID	Defines the value of FID (Frequency ID) to configure the CPU frequency for Pstate2. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	70	[<value in hex>]

Pstate2 DID	Defines the value of DID (Divisor ID) to configure the CPU frequency for Pstate2. The frequency value COF is defined by the formula: COF = 200MHz * FID / DID	8	[<value in hex>]
Pstate2 VID	Defines the value of VID (Voltage ID) to configure the CPU core voltage for Pstate2. The relationship between the hexadecimal values and the corresponding voltage levels are defined by AMD Serial VID Interface Specification.	48	[<value in hex>]

3.3.3.1.1.2 CCD/Core/Thread Enablement

The menu items below configure the number of CCDs (Core Cache Dies) and cores available in the AMD CPUs and also the availability of SMT (Symmetric Multithreading), which allows two threads for each core.

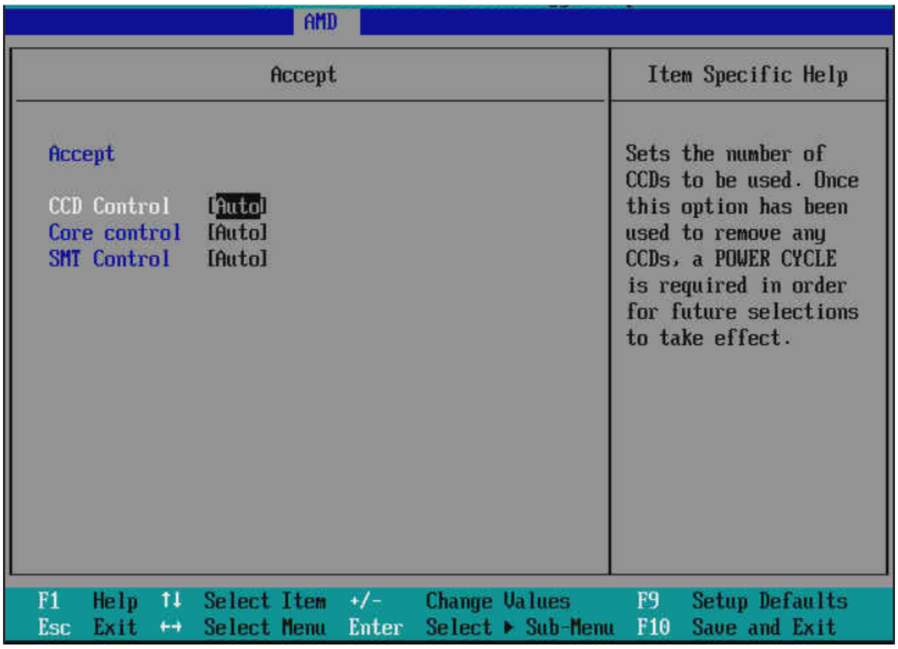

```
                              AMD
                Accept                       Item Specific Help

   Accept                                    Sets the number of
                                             CCDs to be used. Once
   CCD Control      [Auto]                   this option has been
   Core control     [Auto]                   used to remove any
   SMT Control      [Auto]                   CCDs, a POWER CYCLE
                                             is required in order
                                             for future selections
                                             to take effect.

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu    F10  Save and Exit
```

Figure 27: CCD/Core/Thread Enablement Menu

Some applications may achieve best performance by disabling the SMT, once the system reduces the interrupts/queue handling overhead. Adjusting the number of active cores may also be used to control the deployment cost for some hypervisors whose license is based on the number of virtual cores available. For additional information about SMT use cases, please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers”.

For additional information of the AMD CPUs CCDs, cores and additional CPU internal structure, please refer to the document (1) “Socket SP3 Platform NUMA Topology for AMD Family 17h Models 30h–3Fh”.

Menu Item	Description	Default Config	Options
CCD Control	Defines the number of CCDs available in the AMD CPU. The CPU requires a power cycle for changes to be applied. Depending on the CPU model, a limited number of CCDs may be available, so that some options may not be applicable. [Auto} option configures the maximum available number of CCDs. The AMD EPYC 7002 processors series can have up to 8 CCDs.	Auto	[2 CCDs] [3 CCDs] [4 CCDs] [6 CCDs]
Core Control	Defines the number of cores available for each CCD inside the AMD CPU. Therefore, the total number of cores of the CPU will be defined by the number of CCDs multiplied by the number of cores configured in this menu. The CPU requires a power cycle for changes to be applied. [Auto} option configures the maximum available number of cores. The AMD EPYC 7002 processors series can have up to 8 cores per CCD.	Auto	[TWO (1+1)] [FOUR (2+2)] [SIX (3+3)]
SMT Control	Enables/Disables the SMT (Symmetric Multithreading). When disabled, only one thread will be active per core inside the CPU. When enabled, two threads will be active per core inside the CPU. [Auto] option keeps SMT enabled by default.	Auto	[Auto] [Disable]

3.3.3.1.2 Prefetcher settings

The prefetcher is a mechanism that consists in allocating data in the cache memory before the application actually needs it. Depending on the workload running on the server CPU, enabling or disabling the cache prefetchers may improve the processing performance. Additional details about this feature can be found in the AMD document “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” (2).

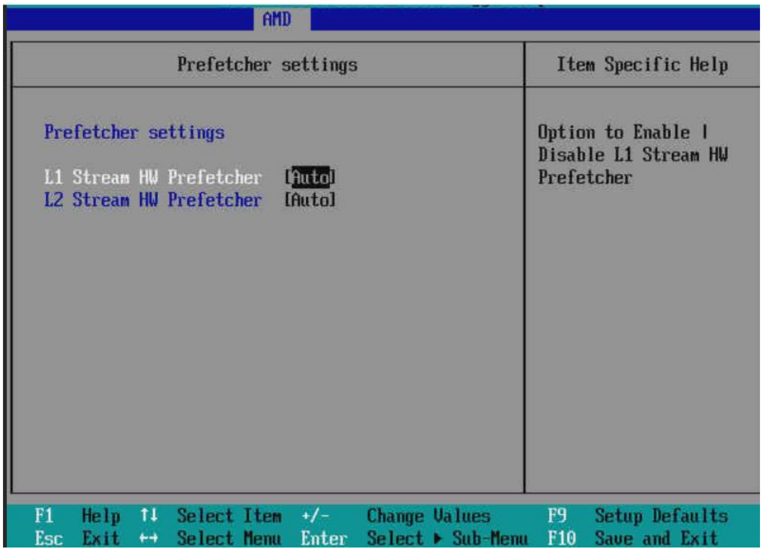

```
                              AMD
            Prefetcher settings                  Item Specific Help

   Prefetcher settings                         Option to Enable |
                                               Disable L1 Stream HW
   L1 Stream HW Prefetcher    [Auto]           Prefetcher
   L2 Stream HW Prefetcher    [Auto]

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 28: Prefetcher Settings Menu

Menu Item	Description	Default Config	Options
L1 Stream HW Prefetcher	Enables/Disables the Prefetcher feature for L1 cache. [Auto] option keeps the prefetcher enabled by default.	Auto	[Auto] [Enable] [Disable]
L2 Stream HW Prefetcher	Enables/Disables the Prefetcher feature for L2 cache. [Auto] option keeps the prefetcher enabled by default.	Auto	[Auto] [Enable] [Disable]

3.3.3.1.3 Core Watchdog

The Core Watchdog Timer (WDT) is a resource used to detect and treat abnormal CPU hang conditions. It is possible to configure a time interval, over which the CPU is checked to make sure it is running its instructions. If no instructions are executed within this time window, the system triggers a configurable kind of error.

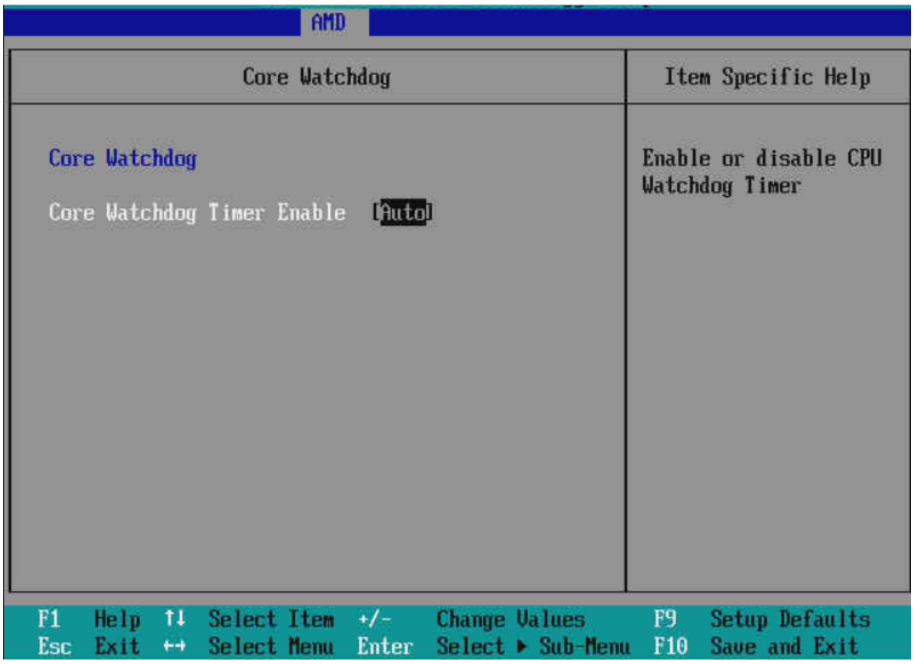

AMD

Core Watchdog

Item Specific Help

Core Watchdog

Core Watchdog Timer Enable [Auto]

Enable or disable CPU Watchdog Timer

Figure 29: Core Watchdog Menu

Menu Item	Description	Default Config	Options
Core Watchdog Timer Enable	Enables/Disables the Core Watchdog Timer functionality.	Auto	[Auto] [Enable] [Disable]

The menu items below are shown only when the Core Watchdog Timer is set to “Enable”.

Menu Item	Description	Default Config	Options
Core Watchdog Timer Interval	Defines the time interval of the WDT, over which the CPU will be checked. If the timer configured in this option expires without the CPU running any instruction, an error is triggered.	Auto	[Auto] [<select a time interval from the list>]
Core Watchdog Timer Severity	Defines the type of error to be generated when the Core Watchdog Timer expires due to CPU inactivity.	Auto	[Auto] [No Error] [Transparent] [Corrected] [Deferred] [Uncorrected] [Fatal]

3.3.3.1.4 RedirectForReturnDis

This option is a workaround used for AMD's Excavator (XV) cores used for Carrizo APUs. Therefore, this option is not expected to affect the DM-SV01 system and it is suggested to keep it with default settings.

Menu Item	Description	Default Config	Options
RedirectForReturnDis	Enables/Disables the RedirectForReturnDis workaround for specific AMD's APUs families.	Auto	[Enabled] [1] [0]

3.3.3.1.5 Platform First Error Handling

Menu Item	Description	Default Config	Options
Platform First Error Handling	When enabled, this option configures the CPU to report all errors to system FW (BIOS/UEFI) initially instead of reporting them directly to the Operating System. The BIOS/UEFI is then responsible to handle the errors and log them in the corresponding error registers.	Enabled	[Enabled] [Disabled] [Auto]

3.3.3.1.6 Core Performance Boost

Menu Item	Description	Default Config	Options
Core Performance Boost	Enables/Disables the Core Performance Boost. When enabled, the CPU maximizes performance when exposed to high workloads. If disabled, the CPU works on a power efficiency basis. [Auto] option keeps this function enabled by default.	Auto	[Auto] [Disabled]

3.3.3.1.7 Global C-state Control

Menu Item	Description	Default Config	Options
Global C-state Control	Enables/Disables the cores C-states. When this option is disabled, the CPU cores will not be able to enter low-power C-states.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.1.8 Power Supply Idle Control

This option is only used for legacy systems where power supplies were turned off when the motherboard was kept in idle state with low power being consumed. This is not applicable to the DM-SV01 power topology, so it's recommended to keep this option with default settings.

Menu Item	Description	Default Config	Options
Power Supply Idle Control	Configures the current limit for disabling the power supply when in idle state (not applicable to DM-SV01).	Auto	[Auto] [Low Current Idle] [Typical Current Idle]

3.3.3.1.9 SEV ASID Count

The ASID (Address Space Identifier) is an identifier that is associated with a particular guest virtual machine in a virtualization environment. The ASID is used by the SEV (Secure Encrypted Virtualization) technology to provide memory encryption for the guest VMs.

This option is used to configure the ASID count, once it affects the maximum physical address space available for the system. The options are defined as follows:

- 253 ASIDs: the physical address space of the system is limited to 16TB.
- 509 ASIDs: the physical address space of the system is limited to 8TB.

Menu Item	Description	Default Config	Options
SEV ASID Count	Configures the number of ASIDs (Address Space Identifiers) available in the system.	Auto	[Auto] [253 ASIDs] [509 ASIDs]

3.3.3.1.10 SEV-ES ASID Space Limit Control

This option is used to enable the manual configuration of the number of SEV-ES ASID guests available in the system.

When this option is set to “Manual”, the additional menu “SEV-ES ASID Space Limit” is shown. Please refer to section 3.3.3.1.11 SEV-ES ASID Space Limit for additional details about the SEV-ES feature and configuration.

Menu Item	Description	Default Config	Options
SEV-ES ASID Space Limit Control	Defines if the configuration of SEV-ES guest VMs is automatic or manual.	Auto	[Auto] [Manual]

3.3.3.1.11 SEV-ES ASID Space Limit

The SEV-ES (Secure Encrypted Virtualization - Encrypted State) is an advanced feature from SEV, which allows additional encryption of the state of the VM which has been subjected to a VMEXIT operation.

This option controls the number of virtual machines using the SEV-ES technology. All the VMs using SEV and with ASID count below the value set in this option will automatically switch to SEV-ES.

Note: if the number configured in this option is higher than the value of “SEV ASID Count” option (section 3.3.3.1.9 SEV ASID Count), then all VMs will be forced to use SEV-ES.

Menu Item	Description	Default Config	Options
SEV-ES ASID Space Limit Control	Configures the number of SEV-ES guests available in the system. Acceptable range: from 0 up to 510.	Auto	[<value>]

3.3.3.1.12 Streaming Stores Control

Streaming Store is a special feature of the CPU which allows some specific data to be written directly in the main memory of the system, without using the CPU cache mechanism. This option enables/disables this feature.

Menu Item	Description	Default Config	Options
Streaming Stores Control	Enables/Disables the Streaming Stores Function of the CPU.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.1.13 Local APIC Mode

This option configures the APIC (Advanced Programmable Interrupt Controller) feature, which allows the addressing of logical CPUs (threads). There are three options available for configuration, as follows:

- **xAPIC**: limits the number of logical CPUs to 255.
- **x2APIC**: limits the number of logical CPUs to 511. This option may not be compatible with some legacy OS.

- **Auto**: The system configures the APIC automatically, depending on the number of threads available in the system. The xAPIC option is configured if there are less than 256 threads, otherwise the x2APIC is used.

Menu Item	Description	Default Config	Options
Local APIC Mode	Configures the xAPIC or x2APIC setting for defining the limit of logical CPUs available.	Auto	[xAPIC] [x2APIC] [Auto]

3.3.3.1.14 ACPI _CST C1 Declaration

Menu Item	Description	Default Config	Options
ACPI _CST C1 Declaration	Controls BIOS/UEFI to declare or not the existence of power state C1 (CPU halt) to the Operating System.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.1.15 MCA error thresh enable

The MCA error thresholding is composed by a counter which is incremented when an error is detected in the MCA. When the counter value exceeds the configured thresholding, the system takes appropriate actions, depending on the error type (error logging, interrupt generation, etc.).

When this option is enabled, an additional menu is shown to configure the counter value. Please refer to section 3.3.3.1.16 MCA error thresh count for additional details.

Menu Item	Description	Default Config	Options
MCA error thresh enable	Enable or disable the MCA error thresholding.	Auto	[Auto] [True] [False]

3.3.3.1.16 MCA error thresh count

This option is used to configure the counter value for the MCA error thresholding and it is shown only if the MCA error thresh enable option is set to True (please refer to section 3.3.3.1.15 MCA error thresh enable).

The counter is configured as hexadecimal value, and the error count is defined as the difference between the value 0xFFF and the value configured in this option. For example, if the value configured in this option is 0xFF5, the threshold count value will be:

- 0xFFF - 0xFF5 = 0x00A (10 errors)

Menu Item	Description	Default Config	Options
MCA error thresh count	Configures the counter value for the MCA error thresholding.	Auto	[<value in hex>]

3.3.3.1.17 SMU and PSP Debug Mode

The SMU (System Management Unit) and PSP (Platform Security Processor) Debug Mode option, when enabled, prevents the system from generating a fatal error and resetting the system when specific uncorrected errors are detected by SMU FW or PSP FW.

Menu Item	Description	Default Config	Options
SMU and PSP Debug Mode	Enables/Disables the SMU and PSP Debug Mode.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.1.18 Xtrig7 Workaround

This option is used only for development and debug purposes, please keep it with default settings.

Menu Item	Description	Default Config	Options
Xtrig7 Workaround	Enables/Disables some workarounds used for debugging the CPU.	Auto	[Auto] [No Workaround] [Bronze Workaround] [Silver Workaround]

3.3.3.1.19 PPIN Opt-in

The PPIN (Protected Processor Inventory Number) is the serial number used to individually identify the CPU unit.

This option enables or disables the PPIN, making the processors’ serial number available or not for the Operating System.

Menu Item	Description	Default Config	Options
PPIN Opt-in	Enables/Disables the PPIN (Protected Processor Inventory Number).	Auto	[Enabled] [Disabled] [Auto]

3.3.3.1.20 RdRand

The RDRAND is a x86 CPU instruction used to access random values generated by the RNG (Random Number Generator) feature from the processor.

This option enables or disables the support for this CPU instruction.

Menu Item	Description	Default Config	Options
RdRand	Enables/Disables the RDRAND instruction support for generating random values.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2 DF (Data Fabric) Common Options

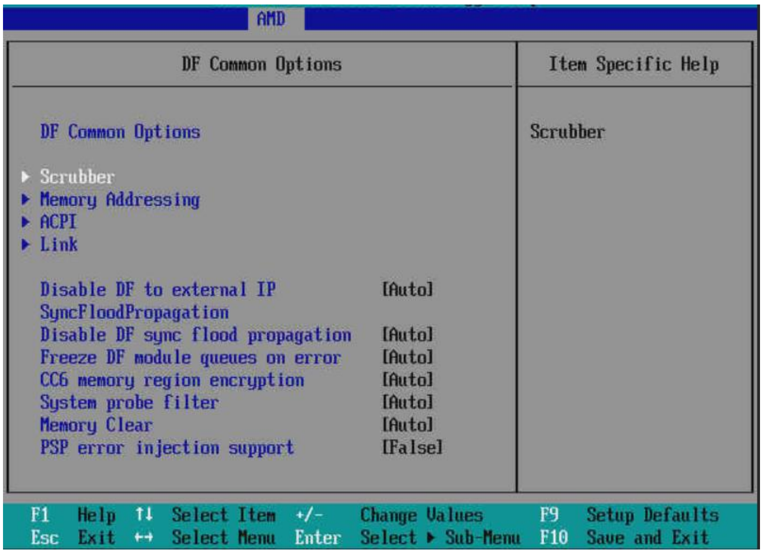

```
AMD
DF Common Options                                   Item Specific Help

DF Common Options                                   Scrubber

▶ Scrubber
▶ Memory Addressing
▶ ACPI
▶ Link

Disable DF to external IP             [Auto]
SyncFloodPropagation
Disable DF sync flood propagation     [Auto]
Freeze DF module queues on error      [Auto]
CC6 memory region encryption          [Auto]
System probe filter                   [Auto]
Memory Clear                          [Auto]
PSP error injection support           [False]

F1  Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
Esc Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu  F10  Save and Exit
```

Figure 30: DF Common Options Menu

3.3.3.2.1 Disable DF to external IP SyncFloodPropagation

The sync flood is a method of error reporting which halts the processor when an uncorrectable error is detected and sends a fatal error event message to the components of the internal data fabric (DF) of the CPU.

This option is used to enable or disable the propagation of sync flood to elements outside of CPU data fabric (DF), like UMC (Unified Memory Controller) and other modules, when an uncorrectable error is detected.

Menu Item	Description	Default Config	Options
Disable DF to external IP SyncFloodPropagation	Enables/Disables the propagation of sync flood to components outside of the data fabric (DF).	Auto	[Sync flood disabled] [Sync flood enabled] [Auto]

3.3.3.2.2 Disable DF sync flood propagation

The sync flood is a method of error reporting which halts the processor when an uncorrectable error is detected and sends a fatal error event message to the components of the internal data fabric (DF) of the CPU.

This option is used to enable or disable the propagation of sync flood to the data fabric (DF) components inside the CPU when an uncorrectable error is detected.

Menu Item	Description	Default Config	Options
Disable DF sync flood propagation	Enables/Disables the propagation of sync flood to the data fabric (DF) components.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2.3 Freeze DF module queues on error

When a fatal error is detected by the system, it’s expected that the internal data fabric (DF) of the CPU “freezes” the data queues in order to prevent the corrupted data to be sent out to other modules of the system. This option is used to enable or disable this behavior.

Menu Item	Description	Default Config	Options
Freeze DF module queues on error	Enables/Disables the action of freezing the data queues of internal data fabric (DF).	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2.4 CC6 memory region encryption

This option defines if the memory region used for entering/exiting the CC6 (Core C6 state) must be encrypted or not. The CC6 is a core power state where the power is gated off from the core when it is in idle for energy efficiency purposes.

Menu Item	Description	Default Config	Options
CC6 memory region encryption	Enables/Disables the encryption of the memory region used for CC6 state.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2.5 System probe filter

The probe filter creates a directory structure inside the core’s cache of the CPU to help limit the number of coherency probes and therefore improves efficiency for accessing local DRAM.

Menu Item	Description	Default Config	Options
System probe filter	Enables/Disables the probe filter.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2.6 Memory Clear

This option enables or disables the memory clear procedure after the DRAM training. The memory clear disabled state is only valid for non-ECC DIMMs. Once DM-SV01 system only uses ECC DIMMs, the option below must be always enabled.

Menu Item	Description	Default Config	Options
Memory Clear	Enables/Disables the memory clear procedure after the DRAM training.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.2.7 PSP error injection support

The PSP (Platform Security Processor) error injection is used only for development and debug purposes, so it’s recommended to keep this option with default settings.

Menu Item	Description	Default Config	Options
PSP error injection support	Enables/Disables the error injection support for PSP.	False	[True] [False]

3.3.3.2.8 Scrubber

The scrubber is a process responsible for reading each DRAM memory location looking for errors, correcting them when possible (using ECC), and writing back the corrected data to the corresponding location. If the error detected is classified as uncorrectable, then the scrubber writes back poisoned data to the memory location.

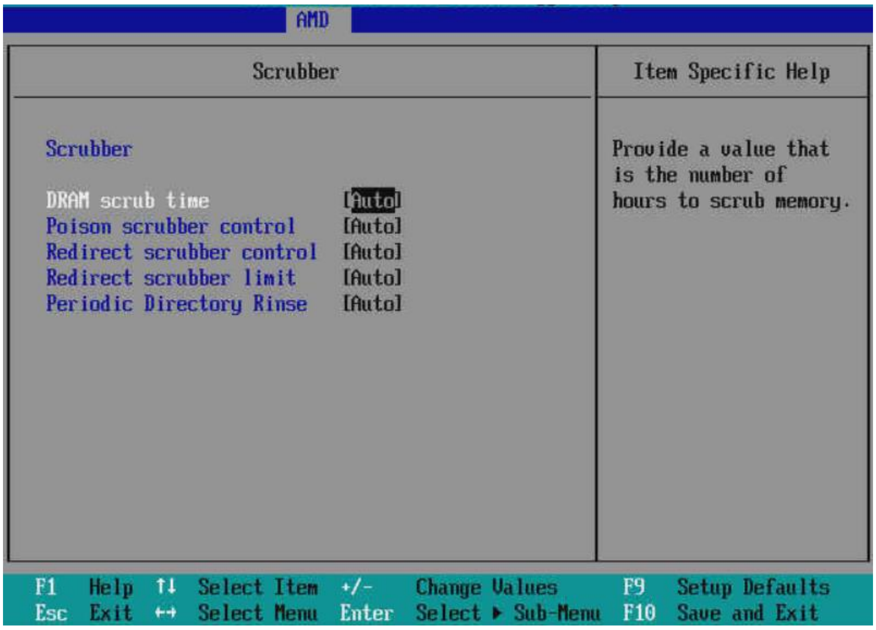

```
                              AMD
                  Scrubber                          Item Specific Help

   Scrubber                                       Provide a value that
                                                  is the number of
   DRAM scrub time              [Auto]            hours to scrub memory.
   Poison scrubber control      [Auto]
   Redirect scrubber control    [Auto]
   Redirect scrubber limit      [Auto]
   Periodic Directory Rinse     [Auto]

F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu    F10  Save and Exit
```

Figure 31: Scrubber Menu

The scrubber is done periodically in the entire DRAM populated in the system, and the user can configure some parameters of the scrubbing process using the settings described in the tables below.

Menu Item	Description	Default Config	Options
DRAM scrub time	Configures the time interval (in hours) for initiating each scrubber process.	Auto	[Disabled] [1 hour] [4 hours] [8 hours] [16 hours] [24 hours] [48 hours] [Auto]

The poison scrubber is a special type of scrubber, which runs whenever an uncorrectable ECC error is detected. Then, the scrubber writes poisoned data back to the memory and logs a deferred error. Once the data is marked as poison, it will not be consumed by processes or applications nor cause a fatal error. The poison data may cause only related processes to fail and must be terminated.

Menu Item	Description	Default Config	Options
Poison scrubber control	Enables/Disables the poison scrubber.	Auto	[Auto] [Enabled] [Disabled]

The redirect scrubber is another type of scrubber, which is started whenever a correctable error is detected in the system's DRAM memory, in order to correct the data and write it back to the memory location. It is also possible to configure the maximum number of redirect scrubbers that can be invoked.

Menu Item	Description	Default Config	Options
Redirect scrubber control	Enables/Disables the redirect scrubber.	Auto	[Auto] [Enabled] [Disabled]
Redirect scrubber limit	Configures the maximum number of redirect scrubbers that can be issued. It is possible to configure unlimited redirect scrubbers by using the “Infinite” setting.	Auto	[2] [4] [8] [Infinite] [Auto]
Periodic Directory Rinse	Enables/Disables the Periodic Directory Rinse function.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.2.9 Memory Addressing

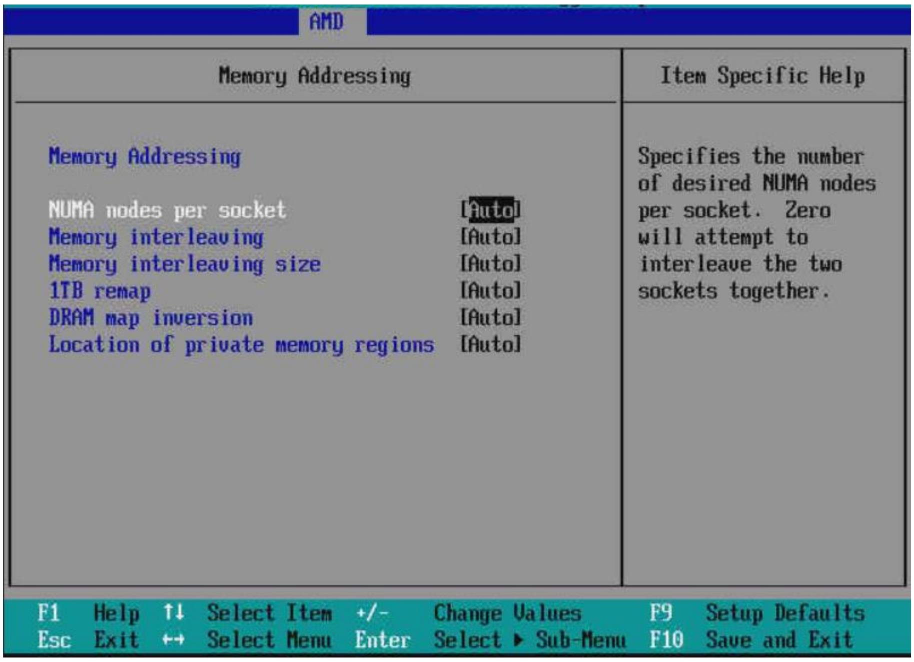
AMD
Memory Addressing
Item Specific Help

Memory Addressing

```
NUMA nodes per socket                 [Auto]
Memory interleaving                   [Auto]
Memory interleaving size              [Auto]
1TB remap                             [Auto]
DRAM map inversion                    [Auto]
Location of private memory regions    [Auto]
```

Specifies the number of desired NUMA nodes per socket. Zero will attempt to interleave the two sockets together.

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 32: Memory Addressing Menu

The option below configures the number of NUMA (Non-Uniform Memory Access) nodes for each CPU socket in the system. The options are as follows (where NPS stands for “Numa Per Socket”):

- NPS0: having zero NUMA nodes per socket means the system will try to interleave the two CPU sockets together, i.e., configuring a single NUMA node composed of both CPUs.
- NPS1: configures one NUMA node for each CPU socket.
- NPS2: configures two NUMA nodes for each CPU socket.
- NPS4: configures four NUMA nodes for each CPU socket.

Please refer to (1) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more information about AMD's CPUs architecture and NUMA nodes topology.

Refer also to the document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for details about the recommended settings for the NUMA node topology.

Menu Item	Description	Default Config	Options
NUMA nodes per socket	Configures the number of NUMA nodes for each CPU socket.	Auto	[Auto] [NPS0] [NPS1] [NPS2] [NPS4]

Memory interleaving is a configuration mode of the DRAM memory which consists in mapping the consecutive memory addresses in different memory banks. Using this approach, when a set of sequential data must be read, the read operations will not be performed repeatedly in the same bank, but instead will run once per bank in turn. Due to DRAM topology reasons, the interleaved access results in better overall performance for most use cases. The options below are used to enable or disable this feature and also to configure the number of bytes to be interleaved.

Menu Item	Description	Default Config	Options
Memory interleaving	Enables/Disables the memory interleaving feature.	Auto	[Auto] [Disabled]
Memory interleaving size	Configures the size of each interleaved memory section, in bytes.	Auto	[Auto] [256 Bytes] [512 Bytes] [1 KB] [2 KB]

When IOMMU is active, it’s necessary to mark a memory region which is mapped just below the 1TB boundary as reserved, making this memory area unavailable to the system.

The 1TB remap option below is used in systems with more than 1TB of installed DRAM to remap this memory region to higher addresses in order to make it available again. Depending on the system configuration (NUMA nodes and interleaving options), it might not be possible to perform this remap, so that the memory region marked as reserved will remain unused.

Menu Item	Description	Default Config	Options
1TB remap	Enables/Disables the 1TB remap on systems with more than 1TB DRAM memory installed.	Auto	[Auto] [Do not remap] [Attempt to remap]

The DRAM map inversion option is used to invert the memory addressing schema on the system, so that the highest memory channels will be assigned with the lowest logical DRAM addresses.

Menu Item	Description	Default Config	Options
DRAM map inversion	Enables/Disables the memory addressing assignment inversion.	Auto	[Auto] [Enabled] [Disabled]

The option below configures the allocation of the private memory regions, i.e., the regions reserved for the CPU modules and not available for the Operating System. This memory region can be configured as follows:

- Distributed: requires a system with 8 CCDs and 8 DRAM modules installed. Distributes the private memory region across all the DRAM modules so that each die can access its corresponding memory space in the DRAM channel directly attached to it. Therefore, it’s expected that this configuration provides the best performance.
- Consolidated: consolidates all private memory space at the beginning of the first DRAM module. This option is worse in terms of performance, but it is a better approach for the OS, once it has to consider only a single reserved region.

Menu Item	Description	Default Config	Options
Location of private memory regions	Configures if the private memory region for CPU modules is consolidated on a single DRAM module or distributed between all DRAM channels.	Auto	[Auto] [Distributed] [Consolidated]

3.3.3.2.10 ACPI

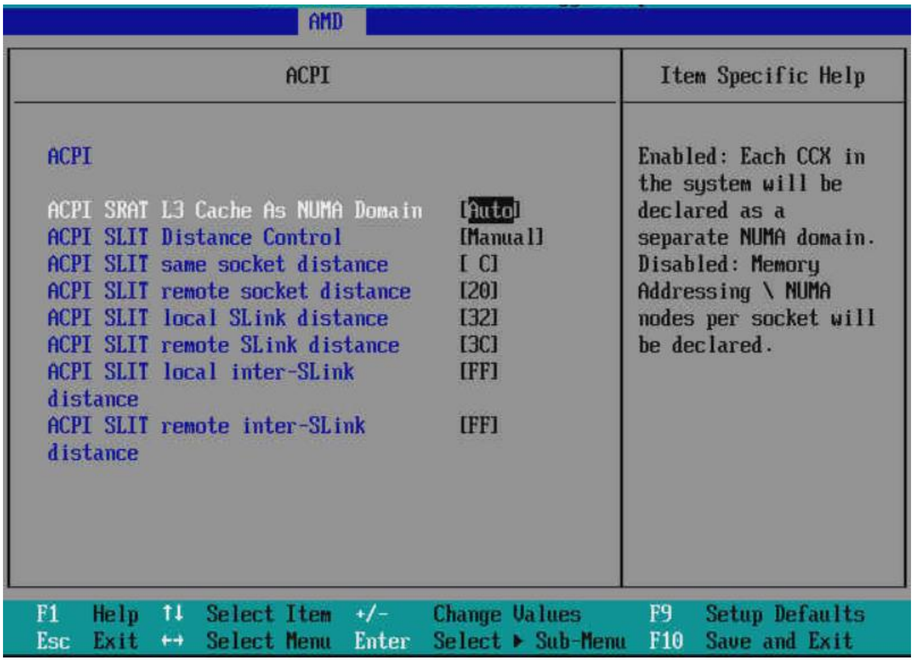

```
                              AMD
                 ACPI                          Item Specific Help

  ACPI                                         Enabled: Each CCX in
                                               the system will be
  ACPI SRAT L3 Cache As NUMA Domain  [Auto]    declared as a
  ACPI SLIT Distance Control         [Manual]  separate NUMA domain.
  ACPI SLIT same socket distance     [ C]      Disabled: Memory
  ACPI SLIT remote socket distance   [20]      Addressing \ NUMA
  ACPI SLIT local SLink distance     [32]      nodes per socket will
  ACPI SLIT remote SLink distance    [3C]      be declared.
  ACPI SLIT local inter-SLink        [FF]
  distance
  ACPI SLIT remote inter-SLink       [FF]
  distance

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 33: ACPI Menu

When the “ACPI (Advanced Configuration and Power Interface) SRAT (System Resource Affinity Table) L3 Cache As NUMA Domain” option is enabled, each CCX (Core Complex) of the CPUs is declared as a particular NUMA Domain, once each CCX has its own L3 cache. If this option is enabled, it will override the NUMA nodes configuration from section 3.3.3.2.9 Memory Addressing, and the number of NUMA nodes per socket (NPS) of the system will be equal to the number of CCX available in the CPUs (a maximum of 32 NPS can become available in a two socket system, depending on the CPU model installed on the system). For additional details regarding this configuration, please refer to AMD documents (1) Socket SP3 Platform NUMA Topology for AMD Family 17h Models 30h–3Fh and (2) Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Server.

Menu Item	Description	Default Config	Options
ACPI SRAT L3 Cache As NUMA Domain	When enabled, each CCX inside the CPU is declared as a separate NUMA domain. The resulting number of NUMA domains in the system will correspond to the total number of CCX modules available, which is dependent on the CPU model.	Auto	[Auto] [Enabled] [Disabled]

The SLIT (System Locality Distance Information Table) is a table containing numeric values which defines the relative “distances” for a CCD (Core Complex Die) to access another CCD in the system. The higher the number in the table, the higher will be the expected latency to access the corresponding node. The OSes may use this information to make decisions about memory allocation, once it can foresee the faster ways to access some elements of the CPUs. The option below enables some manual settings regarding the SLIT distance control.

Menu Item	Description	Default Config	Options
ACPI SLIT Distance Control	Configures the ACPI SLIT Distance control as automatic or manual. If set to manual, additional menu items are shown to configure SLIT parameters.	Auto	[Auto] [Manual]

When the ACPI SLIT Distance Control is set to “Manual”, the options below become available.

Menu Item	Description	Default Config	Options
ACPI SLIT same socket distance	Configures the distance between the die elements residing inside the CPU socket.	Auto	[<value in hexadecimal>]
ACPI SLIT remote socket distance	Configures the distance between the die elements residing in separate CPU sockets.	Auto	[<value in hexadecimal>]
ACPI SLIT local Slink distance	Configures the distance between the die elements and a S-Link residing inside the CPU socket.	Auto	[<value in hexadecimal>]
ACPI SLIT remote SLink distance	Configures the distance between the die elements and a S-Link residing in separate CPU sockets.	Auto	[<value in hexadecimal>]

ACPI SLIT local inter-SLink distance	Configures the distance between two S-Link domains residing inside the CPU socket.	Auto	[<value in hexadecimal>]
ACPI SLIT remote inter-SLink distance	Configures the distance between two S-Link domains residing in separate CPU sockets.	Auto	[<value in hexadecimal>]

The option below configures the SLIT value between one CPU and another in a two socket system, defining the distance as “near” or “far”. Configuring this distance as “near” may improve the performance of workloads spread into both CPU sockets. For additional details about the recommended configurations, please refer to AMD document (2) Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Server.

Menu Item	Description	Default Config	Options
ACPI SLIT remote relative distance	Configures the SLIT distance between the CPUs (in a two socket system) as near or far.	Auto	[Auto] [Near] [Far]

3.3.3.2.11 Link (GMI and xGMI)

The options below are used to control some settings of the GMI (Global Memory Interconnect) and the xGMI (External Global Memory Interconnect) buses of the CPU. The GMI and xGMI are high speed links used by the internal dies of the AMD’s CPU to communicate with each other. While the GMI is used to connect dies from the same CPU socket, the xGMI is used to interconnect dies from different CPU sockets.

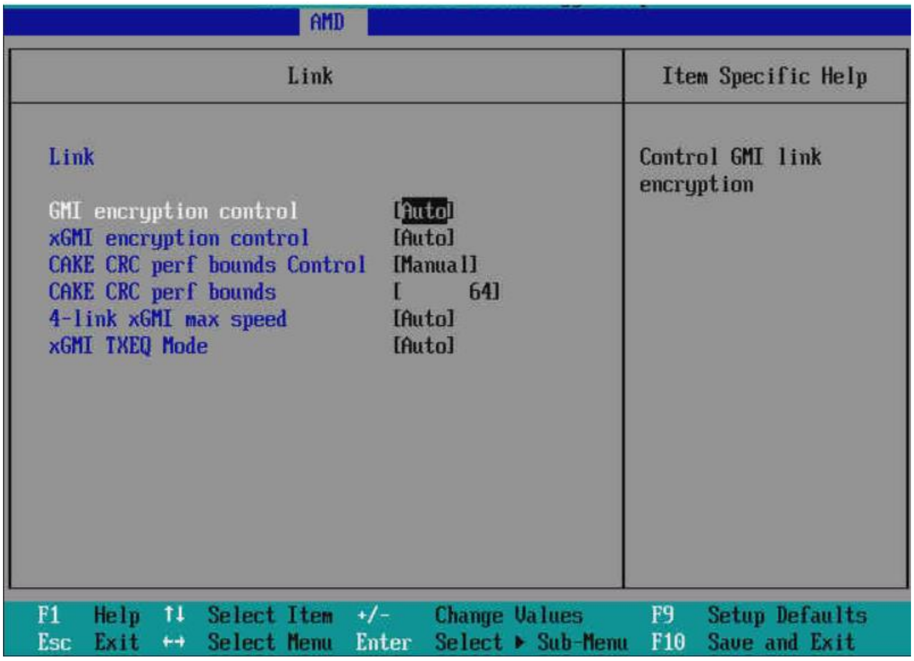

```
AMD
Link                                        Item Specific Help

Link                                        Control GMI link
                                            encryption
GMI encryption control      [Auto]
xGMI encryption control     [Auto]
CAKE CRC perf bounds Control [Manual]
CAKE CRC perf bounds        [      64]
4-link xGMI max speed       [Auto]
xGMI TXEQ Mode              [Auto]

F1   Help  ↑↓ Select Item  +/-   Change Values       F9   Setup Defaults
Esc  Exit  ↔  Select Menu  Enter Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 34: Link Menu

Menu Item	Description	Default Config	Options
GMI encryption control	Enables/Disables the encryption of the GMI bus.	Auto	[Auto] [Enabled] [Disabled]
xGMI encryption control	Enables/Disables the encryption of the xGMI bus.	Auto	[Auto] [Enabled] [Disabled]
CAKE CRC perf bounds Control	When set to “Manual”, allows the user to manually define the acceptable performance loss of the GMI/xGMI buses.	Auto	[Auto] [Manual]

The option below is shown only when “CAKE CRC perf bounds Control” is set to “Manual” and it configures an acceptable limit for CRC errors in the CAKE (Coherent AMD socKet Extender) module of the CPU, which is the element that handles the communication of the GMI/xGMI buses. When the rate of CRC errors on a GMI/xGMI bus crosses the limit defined in the option below, then a CRC error is logged by the system.

Menu Item	Description	Default Config	Options
CAKE CRC perf bounds	Configure the acceptable performance loss for the GMI/xGMI buses. The percentage of allowable losses is defined as: 0.00001% x [value from this option] As an example, configuring a value of 1,000,000 in this menu, results in 10% allowed loss.	Auto	[<value in decimals>]
4-link xGMI max speed	Configures the speed of the xGMI link. The [Auto] option configures the xGMI at 16 Gbps, which is the highest speed supported by the DM-SV01 server.	Auto	[Auto] [10.667Gbps] [13Gpbs] [16Gpbs]
xGMI TXEQ Mode	Configures the equalization model used for training the xGMI link. It’s highly suggested to keep this option with [Auto] setting.	Auto	[Auto] [TXEQ_Disabled] [TXEQ_Lane] [TXEQ_Link] [TXEQ_RX_Vet]

3.3.3.3 UMC (Unified Memory Controller) Common Options

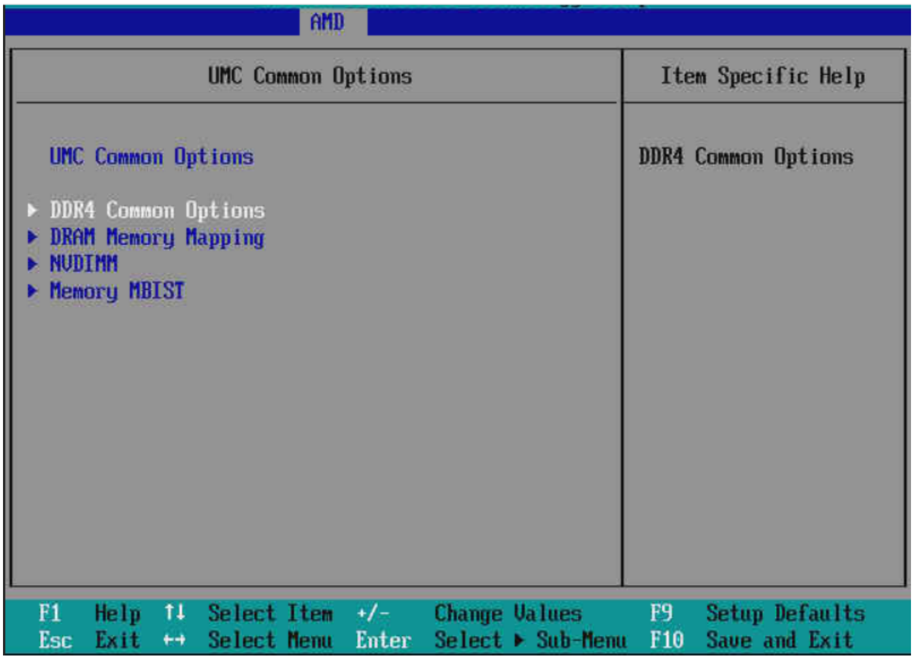

```
AMD
UMC Common Options                          Item Specific Help

  UMC Common Options                        DDR4 Common Options

▸ DDR4 Common Options
▸ DRAM Memory Mapping
▸ NVDIMM
▸ Memory MBIST

F1   Help  ↑↓  Select Item  +/-    Change Values     F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 35: UMC Common Options Menu

3.3.3.3.1 DDR4 Common Options

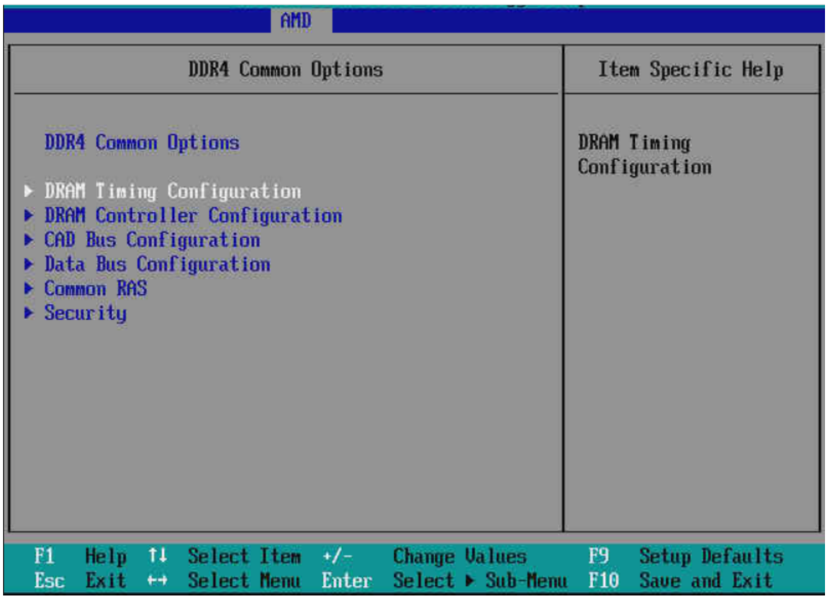

```
AMD
DDR4 Common Options                         Item Specific Help

  DDR4 Common Options                       DRAM Timing
                                            Configuration
▸ DRAM Timing Configuration
▸ DRAM Controller Configuration
▸ CAD Bus Configuration
▸ Data Bus Configuration
▸ Common RAS
▸ Security

F1   Help  ↑↓  Select Item  +/-    Change Values     F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 36: DDR4 Common Options Menu

3.3.3.3.1.1 DRAM Timing Configuration

Warning! “Operating your AMD processor outside of specification or in excess of factory settings, including but not limited to overclocking, may damage or shorten the life of your processor or other system components, create system instabilities (e.g., data loss and corrupted images) and in extreme cases may result in total system failure. AMD does not

provide support or service for issues or damages related to use of an AMD processor outside of processor specifications or in excess of factory settings."

When the user enters the DRAM Timing Configuration menu, the warning above is shown and the user is prompted to Accept or Decline. When the "Accept" option is selected, the options below are shown.

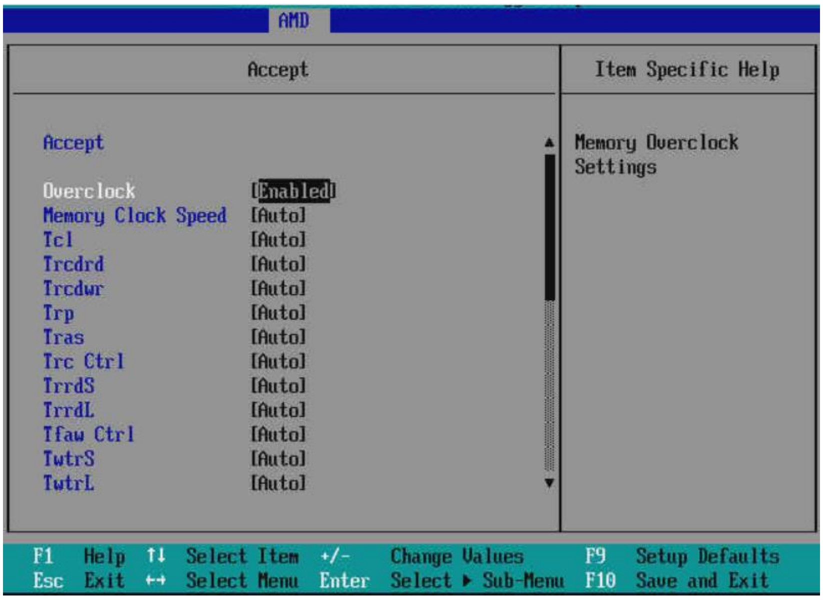

```
                               AMD
                  Accept                        Item Specific Help

 Accept                                         Memory Overclock
                                                Settings
 Overclock            [Enabled]
 Memory Clock Speed   [Auto]
 Tcl                  [Auto]
 Trcdrd               [Auto]
 Trcdwr               [Auto]
 Trp                  [Auto]
 Tras                 [Auto]
 Trc Ctrl             [Auto]
 TrrdS                [Auto]
 TrrdL                [Auto]
 Tfaw Ctrl            [Auto]
 TwtrS                [Auto]
 TwtrL                [Auto]

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 37: DRAM Timing Configuration Menu

Menu Item	Description	Default Config	Options
Overclock	Enables/Disables the overclocking options.	Auto	[Auto] [Enabled]

When the "Overclock" option is set to "Enabled", all the timing options are displayed.

When the "Memory Clock Speed" option is set to [Auto], which is the default setting, the memory is configured to run at the maximum speed available. From the DM-SV01 point of view, the maximum speed capability is 1600MHz. From the DIMMs perspective, on the other hand, the system adjusts the memory clock to match the DIMM with the lower speed available at the system. As an example, if the server is populated with 15 DIMMs rated at 1600MHz and 1 DIMM rated at 1467MHz, the server will run all memory channels at 1467MHz to match the lower speed DIMM.

Although the 1600MHz configuration provides the best performance in terms of throughput, this is not always the best setting for the system, due to latency issues. The infinity fabric (IF) of the AMD CPU runs its internal clock at 1467MHz, so configuring the memory speed to 1467MHz in order to match the IF clock provides the best performance in terms of latency for the memory access.

Additionally, reducing the clock speed also helps the CPU memory controller to save power. This allows the CPU to allocate the remaining power for other processing tasks, eventually improving the performance of the workloads running at the server.

For additional details about this configuration, please refer to the AMD document (2) "Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers".

Menu Item	Description	Default Config	Options
Memory Clock Speed	Configures the memory clock frequency, in MHz.	Auto	[Auto] [667MHz] [800MHz] [933MHz] [1067MHz] [1200MHz] [1333MHz] [1467MHz] [1600MHz]
Tcl	Configures the CAS latency. The CAS latency is the minimum time in clock cycles between a CAS assertion for a read cycle and data returning.	Auto	[Auto] [value in hex from 0x08 up to 0x21 clock cycles]
Trcdrd	Configures the delay between RAS Active and CAS read. The Trcdrd delay is the minimum time in clock cycles between an activate command and a read command, in the same bank.	Auto	[Auto] [value in hex from 0x08 up to 0x1B clock cycles]
Trcdwr	Configures the delay between RAS Active and CAS write. The Trcdwr delay is the minimum time in clock cycles between an activate command and a write command, in the same bank.	Auto	[Auto] [value in hex from 0x08 up to 0x1B clock cycles]
Trp	Configures the row precharge delay time. The row precharge delay is the minimum time in clock cycles between a precharge command and an activate or auto refresh command, in the same bank.	Auto	[Auto] [value in hex from 0x08 up to 0x1B clock cycles]
Tras	Configures the Active to Precharge delay time. The Tras is the minimum time in clock cycles between an activate command and a precharge command, in the same bank.	Auto	[Auto] [value in hex from 0x15 up to 0x3A clock cycles]

When “Trc Ctrl” is configured as “Manual”, the “Trc” option is shown.

Menu Item	Description	Default Config	Options
Trc Ctrl	Configures the Trc delay control as automatic or manual.	Auto	[Auto] [Manual]
Trc	Configures the Trc delay time. The Trc is the minimum time in clock cycles between an activate command and another activate or an auto refresh command, in the same bank.	0x39	[Auto] [<value in hex from 0x1D up to 0x87 clock cycles>]

Menu Item	Description	Default Config	Options
TrrdS	Configures the TrrdS delay time. The TrrdS is the minimum time in clock cycles between an activate command and another activate command, in different chip select banks.	Auto	[Auto] [value in hex from 0x04 up to 0x0C clock cycles]
TrrdL	Configures the TrrdL delay time. The TrrdL is the minimum time in clock cycles between an activate command and another activate command, in the same bank.	Auto	[Auto] [value in hex from 0x04 up to 0x0C clock cycles]

When “Tfaw Ctrl” is configured as “Manual”, the option “Tfaw” is shown.

Menu Item	Description	Default Config	Options
Tfaw Ctrl	Configures the Tfaw delay control as automatic or manual.	Auto	[Auto] [Manual]
Tfaw	Configures the four activate window time. The Tfaw is the window time in clock cycles during which a maximum of four banks can be activated.	0x1A	[Auto] [<value in hex from 0x06 up to 0x36 clock cycles>]

Menu Item	Description	Default Config	Options
TwtrS	Configures the TwtrS (write to read) delay time. The TwtrS is the minimum time in clock cycles between a write operation and a read operation, in different chip select banks.	Auto	[Auto] [value in hex from 0x02 up to 0x0E clock cycles]
TwtrL	Configures the TwtrL (write to read) delay time. The TwtrL is the minimum time in clock cycles between a write operation and a read operation, in the same bank.	Auto	[Auto] [value in hex from 0x02 up to 0x0E clock cycles]

When “Twr Ctrl” is configured as “Manual”, the option “Twr” is shown.

Menu Item	Description	Default Config	Options
Twr Ctrl	Configures the Twr (write recovery time) control as automatic or manual.	Auto	[Auto] [Manual]
Twr	Configures the write recovery time. The Twr is the minimum time in clock cycles between a write operation and a chip select bank precharge.	0x12	[Auto] [<value in hex from 0x0A up to 0x51 clock cycles>]

When “Trcpage Ctrl” is configured as “Manual”, the option “Trcpage” is shown.

Menu Item	Description	Default Config	Options
Trcpage Ctrl	Configures the Trcpage control as automatic or manual.	Auto	[Auto] [Manual]
Trcpage	Configures the Trcpage time. The Trcpage is the minimum time in clock cycles within a refresh window between an activate command and another activate command, in the same bank.	0x00	[Auto] [<value in hex from 0x00 up to 0x3FF clock cycles>]

When “TrdrdScl Ctrl” is configured as “Manual”, the option “TrdrdScL” is shown.

Menu Item	Description	Default Config	Options
TrdrdScL Ctrl	Configures the TrdrdScL (CAS to CAS delay time) control as automatic or manual.	Auto	[Auto] [Manual]
TrdrdScL	Configures the TrdrdScL (CAS to CAS delay time). The TrdrdScL is the minimum time in clock cycles between a CAS deassertion for a read-burst operation and a subsequent CAS assertion for a read-burst operation, in the same bank.	0x03	[Auto] [<value in hex from 0x01 up to 0x0F clock cycles>]

When “TwrwrScl Ctrl” is configured as “Manual”, the option “TwrwrScL” is shown.

Menu Item	Description	Default Config	Options
TwrwrScL Ctrl	Configures the TwrwrScL (CAS to CAS delay time) control as automatic or manual.	Auto	[Auto] [Manual]
TwrwrScL	Configures the TwrwrScL (CAS to CAS delay time). The TwrwrScL is the minimum time in clock cycles between a CAS deassertion for a write-burst operation and a subsequent CAS assertion for a write-burst operation, in the same bank.	0x03	[Auto] [<value in hex from 0x01 up to 0x0F clock cycles>]

When “Trfc Ctrl” is configured as “Manual”, the option below is shown.

Menu Item	Description	Default Config	Options
Trfc Ctrl	Configures the Trfc (Refresh Recovery delay time) control as automatic or manual.	Auto	[Auto] [Manual]
Trfc	Configures the Trfc (Refresh Recovery delay time).	0x138	[Auto] [<value in hex from 0x3C up to 0x3DE clock cycles>]

When “Trfc2 Ctrl” is configured as “Manual”, the option “Trfc2” is shown.

Menu Item	Description	Default Config	Options
Trfc2 Ctrl	Configures the Trfc2 (Refresh Recovery delay time 2x) control as automatic or manual.	Auto	[Auto] [Manual]
Trfc2	Configures the Trfc2 (Refresh Recovery delay time 2x) when using the refresh Fine Granularity as 2x mode (refresh at 2x rate with one half the number of bits as 1x standard mode).	0xC0	[Auto] [<value in hex from 0x3C up to 0x3DE clock cycles>]

When “Trfc4 Ctrl” is configured as “Manual”, the option “Trfc4” is shown.

Menu Item	Description	Default Config	Options
Trfc4 Ctrl	Configures the Trfc4 (Refresh Recovery delay time 4x) control as automatic or manual.	Auto	[Auto] [Manual]
Trfc4	Configures the Trfc4 (Refresh Recovery delay time 4x) when using the refresh Fine Granularity as 4x mode (refresh at 4x rate with one fourth the number of bits as 1x standard mode).	0x84	[Auto] [<value in hex from 0x3C up to 0x3DE clock cycles>]

Menu Item	Description	Default Config	Options
Tcwl	Configures the Tcwl (CAS write latency). The Tcwl is the minimum time in clock cycles between an internal write command and the first write data in the memory.	Auto	[Auto] [value in hex from 0x09 up to 0x14 clock cycles]
Trtp	Configures the Trtp (Read CAS to precharge) delay time. The Trtp is the minimum time in clock cycles between a page read and the corresponding page closing.	Auto	[value in hex from 0x05 up to 0x0E clock cycles]
Tcke	Configures the Tcke (clock enable pulse) time. The Tcke is the minimum time in clock cycles for the CKE (Clock Enable) signal high and low pulses.	Auto	[value in hex from 0x01 up to 0x1F clock cycles]
Trdwr	Configures the Trdwr (read to write) delay time. The Trdwr is the minimum time in clock cycles between a read operation and a write operation.	Auto	[value in hex from 0x01 up to 0x1F clock cycles]
Twrrd	Configures the Twrrdr (write to read) delay time. The Twrrd is the minimum time in clock cycles between a write operation and a read operation.	Auto	[value in hex from 0x01 up to 0x1F clock cycles]
TwrwrSc	Configures the TwrwrSc (write to write) delay time. The TwrwrSc is the minimum time in clock cycles between a write operation and another write operation, in the same bank.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]

TwrwrSd	Configures the TwrwrSd (write to write) delay time. The TwrwrSd is the minimum time in clock cycles between a write operation and another write operation, in the same DIMM.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]
TwrwrDd	Configures the TwrwrDd (write to write) delay time. The TwrwrDd is the minimum time in clock cycles between a write operation and another write operation, in a different DIMM.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]
TrdrdSc	Configures the TrdrdSc (read to read) delay time. The TrdrdSc is the minimum time in clock cycles between a read operation and another read operation, in the same bank.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]
TrdrdSd	Configures the TrdrdSd (read to read) delay time. The TrdrdSd is the minimum time in clock cycles between a read operation and another read operation, in the same DIMM.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]
TrdrdDd	Configures the TrdrdDd (read to read) delay time. The TrdrdDd is the minimum time in clock cycles between a read operation and another read operation, in a different DIMM.	Auto	[value in hex from 0x01 up to 0x0F clock cycles]
ProcODT	Configures the strength (impedance) of the processor ODT (On Die Termination).	Auto	[Auto] [High Impedance] [480 ohm] [240 ohm] [160 ohm] [120 ohm] [96 ohm] [80 ohm] [68.6 ohm] [60 ohm] [53.3 ohm] [48 ohm] [43.6 ohm] [40 ohm] [36.9 ohm] [34.3 ohm] [32 ohm] [30 ohm] [28.2 ohm]

3.3.3.3.1.2 DRAM Controller Configuration

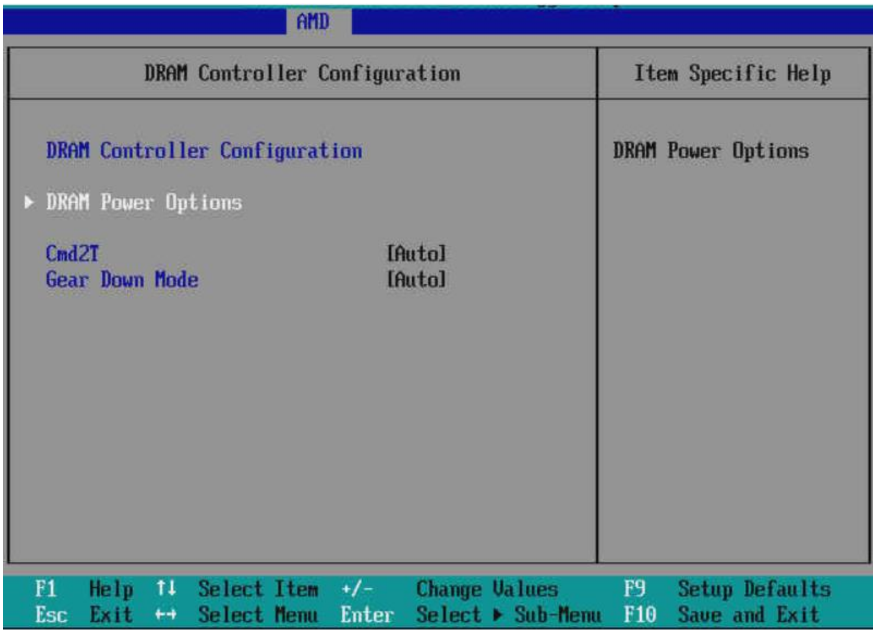

```
                              AMD
  DRAM Controller Configuration                   Item Specific Help

  DRAM Controller Configuration                   DRAM Power Options

► DRAM Power Options

  Cmd2T                               [Auto]
  Gear Down Mode                      [Auto]

F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ► Sub-Menu    F10  Save and Exit
```

Figure 38: DRAM Controller Configuration Menu

DRAM Power Options:

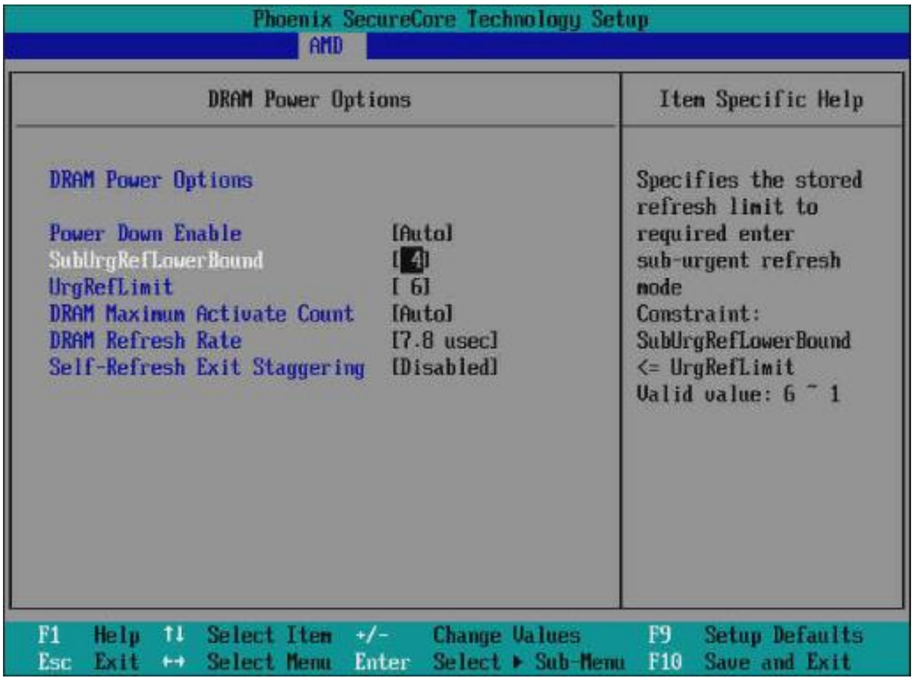

```
                 Phoenix SecureCore Technology Setup
                              AMD
  DRAM Power Options                              Item Specific Help

  DRAM Power Options                              Specifies the stored
                                                  refresh limit to
  Power Down Enable                 [Auto]        required enter
  SubUrgRefLowerBound               [4]           sub-urgent refresh
  UrgRefLimit                       [ 6]          mode
  DRAM Maximum Activate Count       [Auto]        Constraint:
  DRAM Refresh Rate                 [7.8 usec]    SubUrgRefLowerBound
  Self-Refresh Exit Staggering      [Disabled]    <= UrgRefLimit
                                                  Valid value: 6 ~ 1

F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ► Sub-Menu    F10  Save and Exit
```

Figure 39: DRAM Power Options Menu

Menu Item	Description	Default Config	Options
Power Down Enable	Enables/Disables the DRAM power down mode. When enabled, the DRAM will enter power down state when not being used, in order to save energy.	Auto	[Auto] [Enabled] [Disabled]

SubUrgRefLower Bound	Configures the stored refresh limit to enter in the sub-urgent refresh mode. The valid range is from 1 up to 6.	4	[<value in decimal>]
UrgRefLimit	Configures the stored refresh limit to enter in the urgent refresh mode. The valid range is from 1 up to 6.	6	[<value in decimal>]
DRAM Maximum Activate Count	This option overrides DIMM SPD Byte 7 [3:0], configuring the maximum activate count (MAC). The MAC is the maximum number of times a row can be accessed before the adjacent rows require a refresh. The default option is Auto, which keeps the SPD setting.	Auto	[Untested MAC] [700 K] [600 K] [500 K] [400 K] [300 K] [200 K] [Unlimited MAC] [Auto]
DRAM Refresh Rate	Configures the DRAM Refresh period, in microseconds.	7.8 usec	[7.8 usec] [3.9 usec]
Self-Refresh Exit Staggering	Configures a parameter “n” to apply staggering in the self-refresh operations of the DRAMs. This parameter can be “3” or “4” and it is used in the formula below: Tcksrx = Tcksrx + (Trfc/n * (UMC_Number % 4)) When the “n” parameter is configured, the DRAM must wait for Tcksrx before exit the self-refresh mode.	Disabled	[Trfc / 3] [Trfc / 4]

The option below configures the DDR command rate, i.e. it defines if the chip select can be asserted in a single clock (1T) or if two clocks (2T) are needed. In general, 1T setting provides better performance, while 2T setting can be used for compatibility reasons when necessary.

Menu Item	**Description**	**Default Config**	**Options**
Cmd2T	Configures the DDR command rate between 1T or 2T.	Auto	[Auto] [1T] [2T]

The Gear Down Mode allows the DDR memory to use half of its clock frequency by using every other rising edge of the clock for the command/address and control signals latching. This option may be used as a fallback when signal integrity is needed.

Menu Item	**Description**	**Default Config**	**Options**
Gear Down Mode	Enables/Disables the DRAM gear down mode.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.1.3 CAD Bus Configuration

The option below can be used to enable the manual configuration of DDR CAD (command/address) timing parameters.

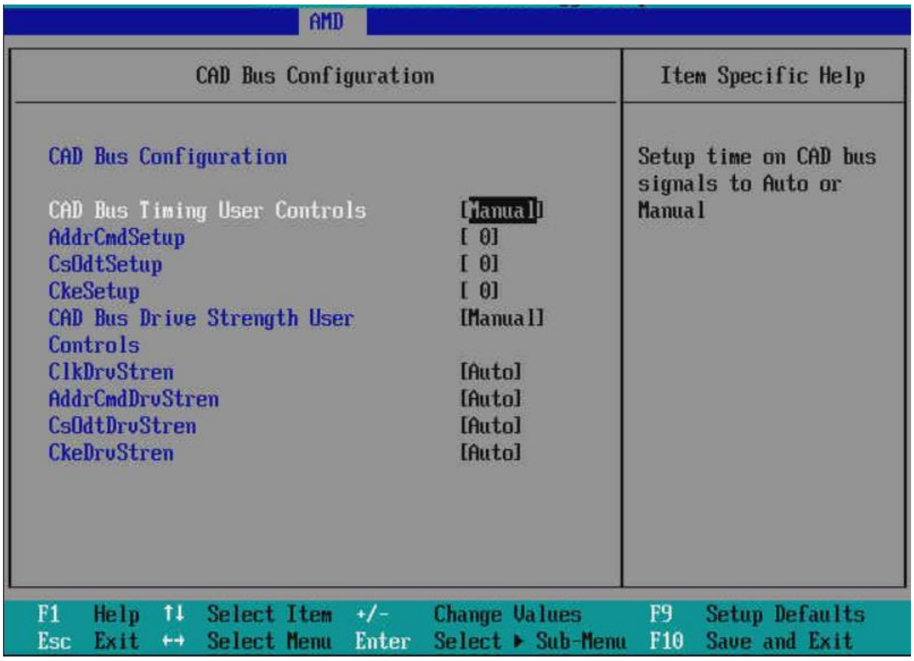

```
                                AMD
          CAD Bus Configuration                  Item Specific Help

  CAD Bus Configuration                          Setup time on CAD bus
                                                 signals to Auto or
  CAD Bus Timing User Controls       [Manual]    Manual
  AddrCmdSetup                       [ 0]
  CsOdtSetup                         [ 0]
  CkeSetup                           [ 0]
  CAD Bus Drive Strength User        [Manual]
  Controls
  ClkDrvStren                        [Auto]
  AddrCmdDrvStren                    [Auto]
  CsOdtDrvStren                      [Auto]
  CkeDrvStren                        [Auto]

 F1   Help  ↑↓  Select Item  +/-    Change Values     F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu F10  Save and Exit
```

Figure 40: CAD Bus Configuration Menu

It’s highly recommended to keep this option in “Auto”, which is the setting already tested and validated for use in the server.

Menu Item	Description	Default Config	Options
CAD Bus Timing User Controls	Configures the CAD (Command/ADdress) timing parameters as automatic or manual.	Auto	[Auto] [Manual]

When the CAD Bus Timing User Controls option is set to manual, the options below are displayed.

Menu Item	Description	Default Config	Options
AddrCmdSetup	Configures the setup time for DDR memory Address and Command signals with respect to the memory clock. Valid input range: 0x00 up to 0x3F.	0	[<value in hex>]
CsOdtSetup	Configures the setup time for DDR memory CS (chip select) and ODT (On Die Termination) signals with respect to the memory clock. Valid input range: 0x00 up to 0x3F.	0	[<value in hex>]
CkeSetup	Configures the setup time for DDR memory CKE (Clock Enable) signal with respect to the memory clock. Valid input range: 0x00 up to 0x3F.	0	[<value in hex>]

The option below can be used to enable the manual configuration of DDR CAD (command/address) bus drive strength parameters.

It's highly recommended to keep this option in "Auto", which is the setting already tested and validated for use in the server.

Menu Item	Description	Default Config	Options
CAD Bus Drive Strength User Controls	Configures the CAD (Command/ADdress) bus drive strength parameters as automatic or manual.	Auto	[Auto] [Manual]

When the CAD Bus Drive Strength User Controls option is set to manual, the options below are displayed.

Menu Item	Description	Default Config	Options
ClkDrvStren	Configures the drive strength (impedance) of the DDR clock signals in ohms.	Auto	[Auto] [120.0 Ohm] [60.0 Ohm] [40.0 Ohm] [30.0 Ohm] [24.0 Ohm] [20.0 Ohm]
AddrCmdDrvStren	Configures the drive strength (impedance) of the DDR address and command signals in ohms.	Auto	[Auto] [120.0 Ohm] [60.0 Ohm] [40.0 Ohm] [30.0 Ohm] [24.0 Ohm] [20.0 Ohm]
CsOdtDrvStren	Configures the drive strength (impedance) of the DDR CS (chip select) and ODT (On Die Termination) signals in ohms.	Auto	[Auto] [120.0 Ohm] [60.0 Ohm] [40.0 Ohm] [30.0 Ohm] [24.0 Ohm] [20.0 Ohm]
CkeDrvStren	Configures the drive strength (impedance) of the DDR CKE (Clock Enable) signal in ohms.	Auto	[Auto] [120.0 Ohm] [60.0 Ohm] [40.0 Ohm] [30.0 Ohm] [24.0 Ohm] [20.0 Ohm]

3.3.3.3.1.4 Data Bus Configuration

The option below can be used to enable the manual configuration of DDR data bus parameters.

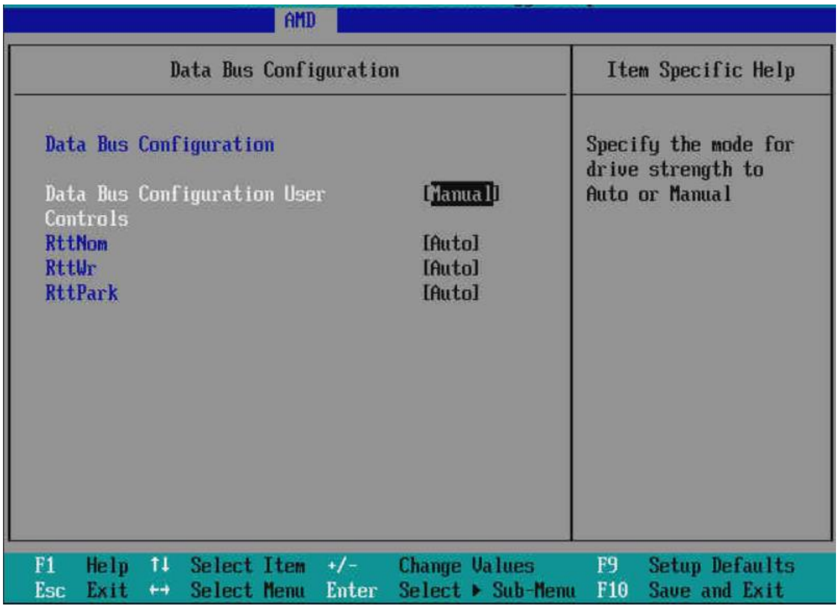

```
                                   AMD
        Data Bus Configuration                    Item Specific Help

  Data Bus Configuration                        Specify the mode for
                                                drive strength to
  Data Bus Configuration User       [Manual]    Auto or Manual
  Controls
  RttNom                            [Auto]
  RttWr                             [Auto]
  RttPark                           [Auto]

F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ► Sub-Menu    F10  Save and Exit
```

Figure 41: Data Bus Configuration Menu

It's highly recommended to keep this option in "Auto", which is the setting already tested and validated for use in the server.

Menu Item	Description	Default Config	Options
Data Bus Configuration User Controls	Configures the Data bus parameters as automatic or manual.	Auto	[Auto] [Manual]

When the Data Bus Configuration User Controls option is set to manual, the options below are displayed.

Menu Item	Description	Default Config	Options
RttNom	Configures the nominal termination resistance of the DRAM.	Auto	[Rtt_Nom Disable] [RZQ/4] [RZQ/2] [RZQ/6] [RZQ/1] [RZQ/5] [RZQ/3] [RZQ/7] [Auto]
RttWr	Configures the termination resistance of the DRAM used for write operations.	Auto	[Dynamic ODT Off] [RZQ/2] [RZQ/1] [Hi-Z] [RZQ/3] [Auto]

RttPark	Configures the park termination resistance of the DRAM used when ODT signal is low.	Auto	[Rtt_PARK Disable] [RZQ/4] [RZQ/2] [RZQ/6] [RZQ/1] [RZQ/5] [RZQ/3] [RZQ/7] [Auto]

3.3.3.3.1.5 Common RAS

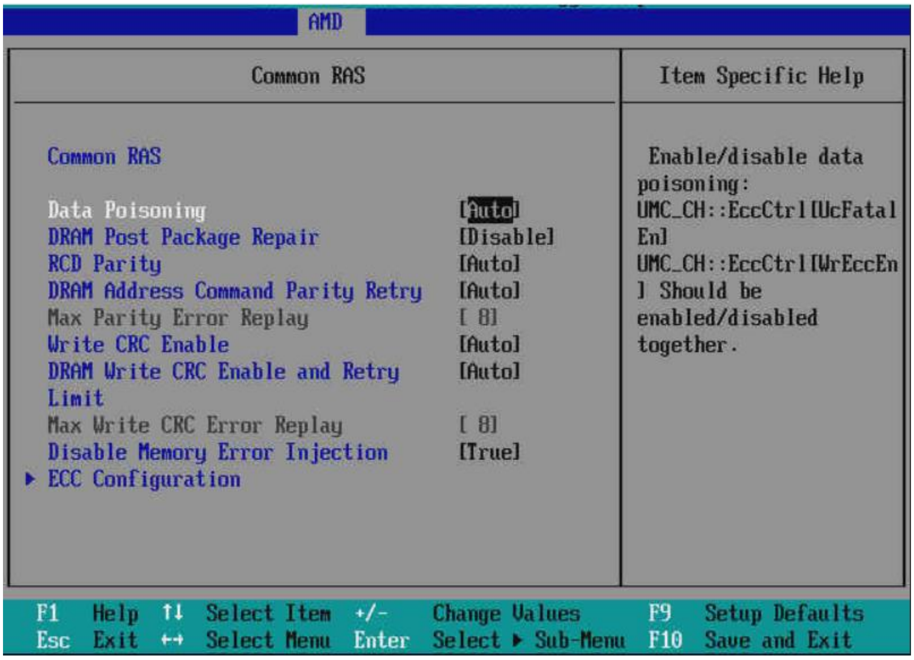

```
                                AMD
           Common RAS                          Item Specific Help

  Common RAS                                  Enable/disable data
                                             poisoning:
  Data Poisoning                  [Auto]     UMC_CH::EccCtrl[UcFatal
  DRAM Post Package Repair        [Disable]  En]
  RCD Parity                      [Auto]     UMC_CH::EccCtrl[WrEccEn
  DRAM Address Command Parity Retry [Auto]   ] Should be
  Max Parity Error Replay         [ 8]       enabled/disabled
  Write CRC Enable                [Auto]     together.
  DRAM Write CRC Enable and Retry [Auto]
  Limit
  Max Write CRC Error Replay      [ 8]
  Disable Memory Error Injection  [True]
▸ ECC Configuration

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 42: Common RAS Menu

The system marks data as poisoned when an uncorrectable error is detected. Once the data is marked as poisoned, other elements in the system will not consume it, avoiding errors and data corruption. The option below is used to enable or disable the data poisoning of DRAM data.

Menu Item	**Description**	**Default Config**	**Options**
Data Poisoning	Enables/Disables the poisoning of data containing an uncorrectable error.	Auto	[Auto] [Enabled] [Disabled]

The DRAM Post Package Repair (PPR) is a feature that allows a failed DRAM row to be replaced by another one, providing a recovery method for a failed portion of the memory. The PPR mechanism is handled by BIOS/UEFI FW.

The menu below allows the user to enable or disable this feature.

Menu Item	Description	Default Config	Options
DRAM Post Package Repair	Enables/Disables the DRAM PPR (Post Package Repair) feature.	Disable	[Enable] [Disable]

The parity is a hardware mechanism used to detect single bit errors. The options below configure the parity for the DRAM modules in the system.

Menu Item	Description	Default Config	Options
RCD Parity	Enables/disables the parity checking for DRAM RCD (Register Clock Driver) memory controller.	Auto	[Auto] [Enabled] [Disabled]
DRAM Address Command Parity Retry	Enables/disables the parity retries for DRAM Address and Command.	Auto	[Auto] [Enabled] [Disabled]

When the “DRAM Address Command Parity Retry” option is enabled, the setting below is shown. The Max Parity Error Replay option configures the maximum number of retries that must be attempted when a parity error is detected.

Menu Item	Description	Default Config	Options
Max Parity Error Replay	Configures the maximum number of retries that must be attempted when a parity error is detected. The value is in hex and varies from 0x00 up to 0x3F.	0x08	[<value in hex>]

The CRC is a hardware mechanism used to detect errors, including burst errors. The options below configure the CRC for the DRAM modules in the system.

Menu Item	Description	Default Config	Options
Write CRC Enable	Enables/disables the CRC (Cyclic Redundancy Check) feature for DRAM data.	Auto	[Auto] [Enabled] [Disabled]
DRAM Write CRC Enable and Retry Limit	Enables/disables the CRC retries for DRAM data.	Auto	[Auto] [Enabled] [Disabled]

When the “DRAM Write CRC Enable and Retry Limit” option is enabled, the setting below is shown. The “Max Write CRC Error Replay” option configures the maximum number of retries that must be attempted when a CRC error is detected.

Menu Item	Description	Default Config	Options
Max Write CRC Error Replay	Configures the maximum number of retries that must be attempted when a CRC error is detected. The value is in hex and varies from 0x00 up to 0x3F.	0x08	[<value in hex>]

The memory error injection option shown below is used only for development and debug purposes, so it's recommended to keep this option with default settings.

Menu Item	Description	Default Config	Options
Disable Memory Error Injection	Configures the error injection disable for DDR memories as true or false.	True	[True] [False]

ECC Configuration:

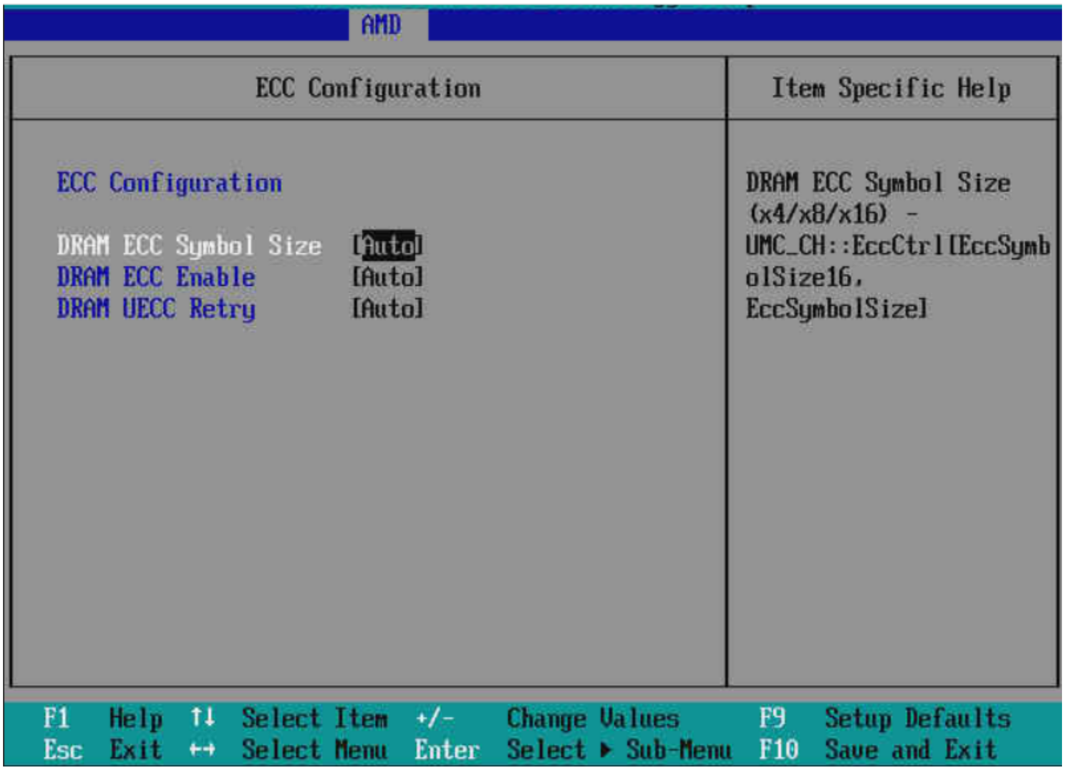

```
                               AMD
            ECC Configuration                    Item Specific Help

  ECC Configuration                            DRAM ECC Symbol Size
                                               (x4/x8/x16) -
  DRAM ECC Symbol Size     [Auto]              UMC_CH::EccCtrl[EccSymb
  DRAM ECC Enable          [Auto]              olSize16,
  DRAM UECC Retry          [Auto]              EccSymbolSize]

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 43: ECC Configuration Menu

Menu Item	Description	Default Config	Options
DRAM ECC Symbol Size	Configures the ECC (Error Correction Code) symbol size as x4, x8 or x16. ● x4: ECC uses 36 symbols, each containing 4 bits, resulting in a 144-bit ECC word with 128 data bits and 16 check bits. ● x8: ECC uses 18 symbols, each containing 8 bits, resulting in a 144-bit ECC word with 128 data bits and 16 check bits ● x16: ECC uses 18 symbols, each containing 16 bits, resulting in a 288-bit ECC word with 256 data bits and 32 check bits	Auto	[Auto] [x4] [x8] [x16]
DRAM ECC Enable	Enables/Disables the ECC (Error Correction Code) feature.	Auto	[Auto] [Enabled] [Disabled]
DRAM UECC Retry	Enables/disables the Uncorrectable ECC error retries for DRAM data. The number of retries for the UECC is the same configured in the "DRAM Write CRC Enable and Retry Limit" option.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.1.6 Security

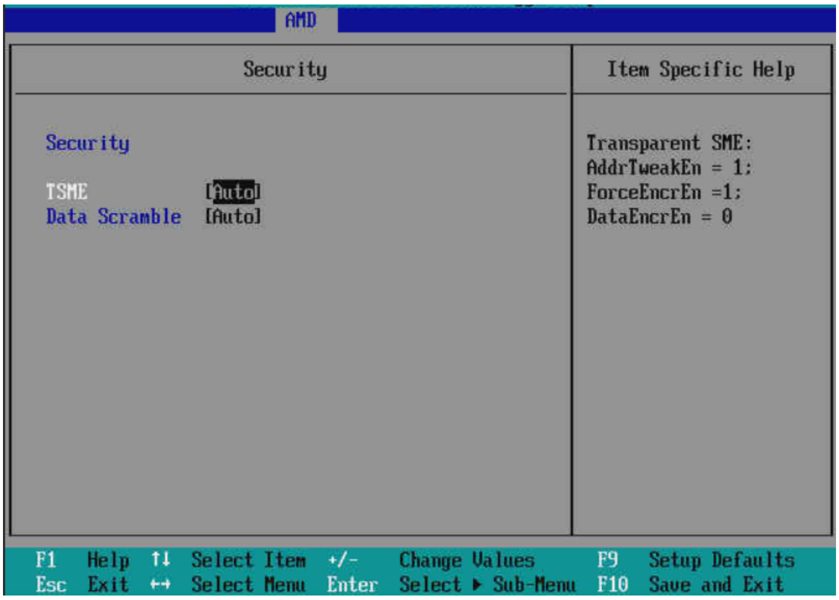
AMD
Security
Item Specific Help
Security

TSME [Auto]
Data Scramble [Auto]

Transparent SME:
AddrTweakEn = 1;
ForceEncrEn =1;
DataEncrEn = 0

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 44: Security Menu

The option below is used to enable or disable the TSME (Transparent Secure Memory Encryption) feature. TSME is a simpler memory encryption mechanism that does not require OS intervention.

Additional details regarding memory encryption and the TSME feature can be found in the AMD document (6) "AMD MEMORY ENCRYPTION".

Menu Item	Description	Default Config	Options
TSME	Enables/Disables the TSME (Transparent Secure Memory Encryption) feature.	Auto	[Auto] [Enabled] [Disabled]

The option below enables or disables the data scrambling of the DDR memories. The data scrambling is a process of randomizing the data bits, providing an additional security layer for the system memory.

Menu Item	Description	Default Config	Options
Data Scramble	Enables/Disables the Data Scramble feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2 DRAM Memory Mapping

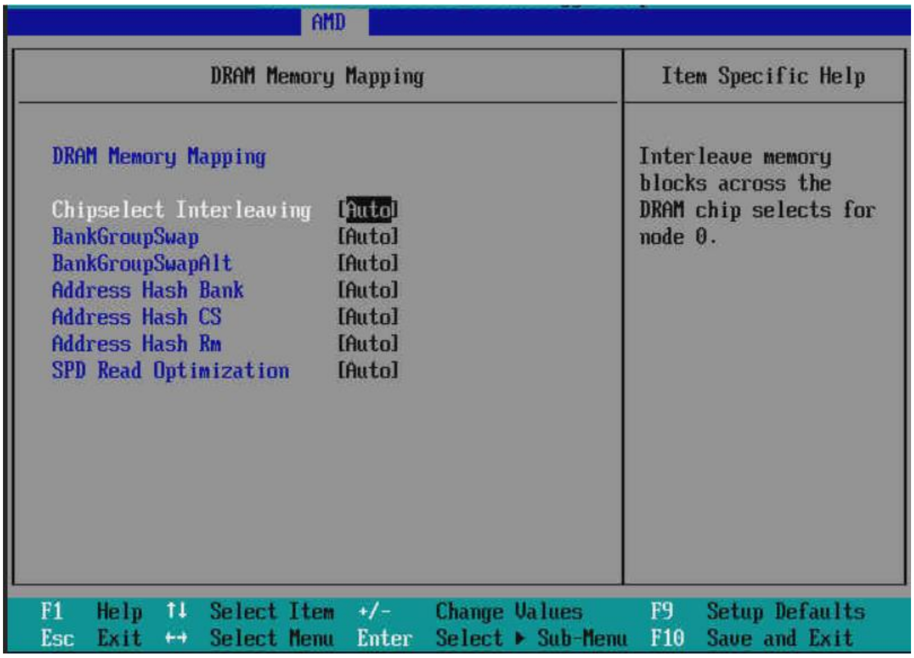

```
DRAM Memory Mapping                          Item Specific Help

DRAM Memory Mapping                          Interleave memory
                                             blocks across the
Chipselect Interleaving      [Auto]          DRAM chip selects for
BankGroupSwap                [Auto]          node 0.
BankGroupSwapAlt             [Auto]
Address Hash Bank            [Auto]
Address Hash CS              [Auto]
Address Hash Rm              [Auto]
SPD Read Optimization        [Auto]
```

AMD
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 45: DRAM Memory Mapping Menu

3.3.3.3.2.1 Chipset Interleaving

This option is used to disable the DDR4 memory interleaving across the chip selects or to keep it in automatic configuration. For details regarding the memory interleaving functionality, please refer to section 3.3.3.2.9 Memory Addressing.

Menu Item	Description	Default Config	Options
Chipset Interleaving	Configures the memory interleaving across the chip selects as automatic or manual.	Auto	[Auto] [Disabled]

3.3.3.3.2.2 BankGroupSwap

The BGS (Bank Group Swap) is a feature from AMD which changes the way applications get assigned to physical addresses in the DDR memory, in order to improve performance for some workloads.

The option below allows the user to enable or disable this feature.

Menu Item	Description	Default Config	Options
BankGroupSwap	Enables/Disables the BGS (Bank Group Swap) feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2.3 BankGroupSwapAlt

The Bank Group Swap Alt option has preference against the Bank Group Swap (3.3.3.3.2.2 BankGroupSwap) and both cannot be used at the same time. Therefore, if the “BankGrouSwapAlt” option is enabled, the “BankGroupSwap” feature is disabled automatically.

Menu Item	Description	Default Config	Options
BankGroupSwap Alt	Enables/Disables the BGS (Bank Group Swap) Alt feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2.4 Address Hash Bank

The option below is used to enable or disable the hashing function for DDR memory banks.

Menu Item	Description	Default Config	Options
Address Hash Bank	Enables/Disables the hashing of DDR memory banks.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2.5 Address Hash CS

The option below is used to enable or disable the hashing function for DDR memory chip selects (CS).

Menu Item	Description	Default Config	Options
Address Hash CS	Enables/Disables the hashing of DDR memory chip selects (CS).	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2.6 Address Hash Rm

The option below is used to enable or disable the hashing function for DDR memory RMs (Rank Multiplication).

Rank Multiplication (RM) is the ratio between the number of physical ranks of the DIMM and the logical ranks in the controller.

Menu Item	Description	Default Config	Options
Address Hash Rm	Enables/Disables the hashing of DDR memory RMs (Rank Multiplication).	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.2.7 SPD Read Optimization

The SPD (Serial Presence Detect) is a standard way of providing inventory information regarding the DDR memory modules to the system. The SPD information is recorded in the DDR module and the system can access this data to identify the DDR module characteristics.

The SPD Read Optimization is a feature that allows the system to improve performance by skipping the reading of some unnecessary fields of the SPD data. The user can enable or disable this optimization by means of the option below.

Menu Item	Description	Default Config	Options
SPD Read Optimization	Enables/Disables the read optimization of the DDR modules SDP (Serial Presence Detect) data.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.3 NVDIMM

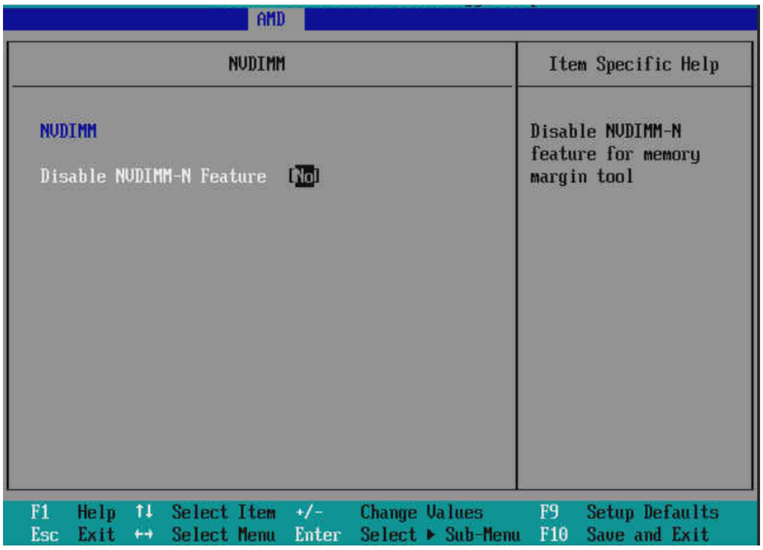

```
                          AMD
            NVDIMM                        Item Specific Help

  NVDIMM                                  Disable NVDIMM-N
                                          feature for memory
  Disable NVDIMM-N Feature   [No]         margin tool

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 46: NVDIMM Menu

3.3.3.3.3.1 Disable NVDIMM-N Feature

This option allows the user to disable the NVDIMM-N (Non-Volatile Dual In-Line Memory Module) feature for debug purposes.

Menu Item	Description	Default Config	Options
Disable NVDIMM-N Feature	Configures the NVDIMM-N feature disable (yes=disable NVDIMM-n / no=do not disable NVDIMM-N).	No	[No] [Yes]

3.3.3.3.4 Memory MBIST

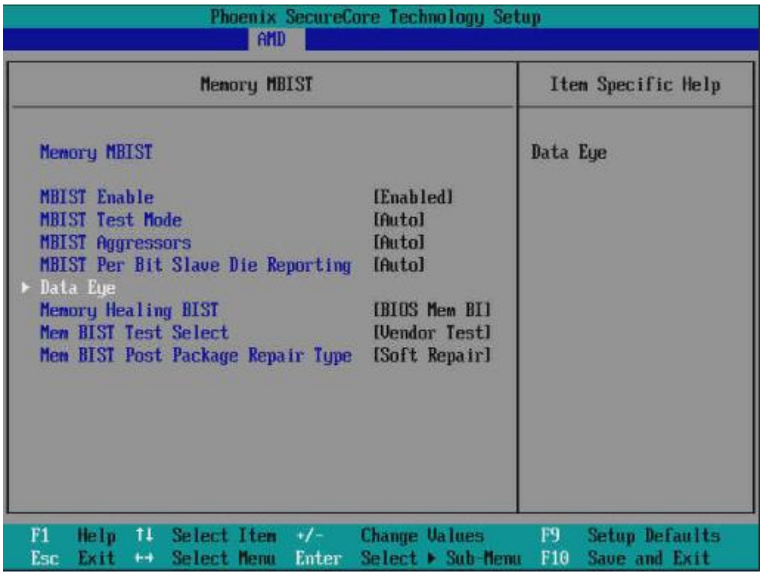

```
                   Phoenix SecureCore Technology Setup
                          AMD
            Memory MBIST                  Item Specific Help

  Memory MBIST                            Data Eye

  MBIST Enable                    [Enabled]
  MBIST Test Mode                 [Auto]
  MBIST Aggressors                [Auto]
  MBIST Per Bit Slave Die Reporting [Auto]
▸ Data Eye
  Memory Healing BIST             [BIOS Mem BI]
  Mem BIST Test Select            [Vendor Test]
  Mem BIST Post Package Repair Type [Soft Repair]

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 47: Memory MBIST Menu

3.3.3.3.4.1 MBIST Enable

The MBIST (Memory Built In Self Test) is a tool that can be used to perform a set of diagnosys in the DDR memories plugged in the system.

The option below is used to enable or disable the test mode.

Menu Item	Description	Default Config	Options
MBIST Enable	Enables/Disables the MBIST (Memory Built In Self Test) test mode.	Disabled	[Enabled] [Disabled]

3.3.3.3.4.2 MBIST Test Mode

This option is used to select the test mode for the MBIST. There are the following options available:

- Interface Mode: used for testing if the electrical interface is functional.
- Data Eye Mode: used for checking the data eye or “margin”.
- Both: uses both “Interface Mode” and “Data Eye Mode”.

Menu Item	Description	Default Config	Options
MBIST Test Mode	Select the MBIST test mode between “Interface”, “Data eye” or “Both”.	Auto	[Interface Mode] [Data Eye Mode] [Both] [Auto]

3.3.3.3.4.3 MBIST Aggressors

The option below is used to enable or disable the aggressor during the MBIST test. The aggressor is responsible for maximizing the possibility of crosstalk by performing a high rate of memory read/write operations.

Menu Item	Description	Default Config	Options
MBIST Aggressors	Enables/Disables the Aggressor during the MBIST memory test.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.4.4 MBIST Per Bit Slave Die Reporting

This option enables or disables the reporting of the slave per bit results to the console. If disabled, only the master results are reported to the console.

Menu Item	Description	Default Config	Options
MBIST Per Bit Slave Die Reporting	Enables/Disables the reporting of per bit results of the slave.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.3.4.5 Memory Healing BIST

This option enables or disables a DRAM test and healing mechanism. The options for running the test are as follows:

- **BIOS Mem BIST:** tests all the memory after training. When failing memory is found, it can be repaired using the mechanism defined in the option explained in section “3.3.3.3.4.7 Mem BIST Post Package Repair Type”. The expected test duration is about 3 minutes per 16GB of installed memory.
- **Self-Healing Mem BIST:** runs the self-healing test as defined in JEDEC DRAM standard. Whenever a failing memory is found, a hard repair is performed to fix it. The expected test duration is about 10 seconds per memory rank per channel.
- **BIOS and Self-Healing Mem BIST:** performs the “BIOS Mem BIST” test first and then the Self-Healing Mem BIST test.

Menu Item	Description	Default Config	Options
Memory Healing BIST	Enables/Disables the full memory BIST test. **Caution:** enabling this test will significantly increase the boot time.	Disabled	[Disabled] [BIOS Mem BIST] [Self-Healing Mem BIST] [BIOS and Self-Healing Mem BIST]

3.3.3.3.4.6 Mem BIST Test Select

This option configures the Self-healing BIOS memory tests. There are three options available:

- **Vendor Tests Enabled:** enables the memory module manufacturer tests to be run in the corresponding manufacturer’s DIMMs.
- **Vendor Tests Disabled:** disables all vendor tests.
- **All Tests - All Vendors:** enables the memory module manufacturer tests to be run on all DIMMs from all manufacturers.

Menu Item	Description	Default Config	Options
Mem BIST Test Select	Configures which tests should be performed when running memory healing BIST tests.	Vendor Tests Enabled	[Vendor Tests Enabled] [Vendor Tests Disabled] [All Tests - All Vendors]

3.3.3.3.4.7 Mem BIST Post Package Repair Type

Configures the type of memory repair to be used when a failed memory is detected by the test described in section “3.3.3.3.4.5 Memory Healing BIST”. There are two types of repair available:

- **Soft Repair (recommended):** performs a temporary repair which can be reverted if needed. When the system is rebooted, the repair is lost and must be applied again.
- **Hard Repair:** performs a permanent repair, modifying the DIMM in an irreversible way.

The user may also configure “no repairs” for debug purposes. If this option is chosen, then the DRAM test will be performed, but the memory will not be repaired in case of failure.

Menu Item	Description	Default Config	Options
Mem BIST Post Package Repair Type	Selects which type of repair will be performed when failed DRAM is detected by means of the BIOS Memory BIST test: soft repair, hard repair or do not attempt to repair.	Soft Repair	[Soft Repair] [Hard Repair] [No Repairs - Test only]

3.3.3.3.4.8 Data Eye

The options below are used to configure some parameters used in the MBIST Data Eye test.

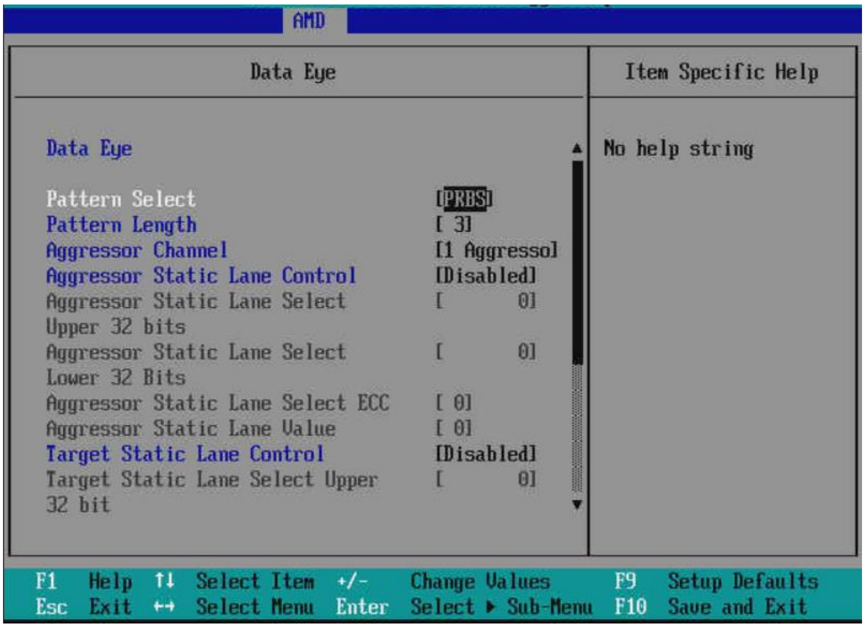

```
                              AMD
              Data Eye                          Item Specific Help

  Data Eye                                      No help string

  Pattern Select                   [PRBS]
  Pattern Length                   [ 3]
  Aggressor Channel                [1 Aggressor]
  Aggressor Static Lane Control    [Disabled]
  Aggressor Static Lane Select     [       0]
  Upper 32 bits
  Aggressor Static Lane Select     [       0]
  Lower 32 Bits
  Aggressor Static Lane Select ECC [ 0]
  Aggressor Static Lane Value      [ 0]
  Target Static Lane Control       [Disabled]
  Target Static Lane Select Upper  [       0]
  32 bit

F1   Help  ↑↓  Select Item  +/-    Change Values     F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu F10  Save and Exit
```

Figure 48: Data Eye Menu

Menu Item	Description	Default Config	Options
Pattern Select	Selects the bit pattern used for the data to be sent during the test: ● PRBS (Pseudo Random Bit Sequence) ● SSO (Simultaneous Switching Output) ● Both PRBS and SSO	PRBS	[PRBS] [SSO] [Both]
Pattern Length	Configures the pattern length from 3 up to 12.	3	[<value from 3

	As an example, a pattern length configured to 7 will generate a sequence of (2^7 - 1) bits, which is equivalent to 127 bits.		up to 12>]
Aggressor Channel	Configures the number of channels that can be set as aggressors. ● 1 Aggressor Channel: 1 channel is tested and 1 channel is set as aggressor. ● 3 Aggressor Channels: 1 channel is tested and 3 channels are set as aggressors. ● 7 Aggressor Channels: 1 channel is tested and 7 channels are set as aggressors.	1 Aggressor Channel	[Disabled] [1 Aggressor Channel] [3 Aggressor Channels] [7 Aggressor Channels]
Aggressor Static Lane Control	Enables/Disables the aggressor static lane control, which allows the user to manually specify the channels which are set as aggressors.	Disabled	[Enabled] [Disabled]
Aggressor Static Lane Select Upper 32 bits	When “Aggressor Static Lane Control” option is enabled, this option can be used to configure the upper 32 DQ lanes of the DDR memory individually as aggressors or not. Each bit represents a single DQ lane. When the bit is set to “1”, the lane is used as an aggressor.	0	[<value in hex>]
Aggressor Static Lane Select Lower 32 bits	When “Aggressor Static Lane Control” option is enabled, this option can be used to configure the lower 32 DQ lanes of the DDR memory individually as aggressors or not. Each bit represents a single DQ lane. When the bit is set to “1”, the lane is used as an aggressor.	0	[<value in hex>]
Aggressor Static Lane Select ECC	When “Aggressor Static Lane Control” option is enabled, this option can be used to configure the 8 ECC DQ lanes of the DDR memory individually as aggressors or not. Each bit represents a single ECC DQ lane. When the bit is set to “1”, the lane is used as an aggressor.	0	[<value in hex>]
Aggressor Static Lane Value	Configures the value that must be set for Aggressor Static Lanes and can be set to “0” or “1”.	0	[<value in hex>]
Target Static Lane Control	Enables/Disables the aggressor static lane control, which allows the user to manually specify the channels which are set as targets to be tested.	Disabled	[Enabled] [Disabled]
Target Static Lane Select Upper 32 bits	When “Target Static Lane Control” option is enabled, this option can be used to configure the upper 32 DQ lanes of the DDR memory individually as targets to be tested or not. Each bit represents a single DQ lane. When the bit is set to “1”, the lane is used as a target.	0	[<value in hex>]
Target Static Lane Select Lower 32 bits	When “Target Static Lane Control” option is enabled, this option can be used to configure the lower 32 DQ lanes of the DDR memory individually as target to be tested or not. Each bit represents a single DQ lane. When the bit is set to “1”, the lane is used as a target.	0	[<value in hex>]
Target Static Lane Select ECC	When “Target Static Lane Control” option is enabled, this option can be used to configure the 8 ECC DQ lanes of the DDR memory individually as targets to be tested or not. Each bit represents a single ECC DQ lane. When the bit is set to “1”, the lane is used as a target.	0	[<value in hex>]
Target Static Lane Value	Configures the value that must be set for Target Static Lanes and can be set to “0” or “1”.	0	[<value in hex>]
Data Eye Type	Configures which results are captured for the data eye:	Worst	[1D Voltage

	● 1D Voltage Sweep: captures the voltage margin (vertical eye opening). ● 1D Timing Sweep: captures the timing margin (horizontal eye opening). ● 2D Full Data Eye: captures both voltage and timing margins. ● Worst Case Margin Only: captures the worst voltage and timing margin from the measurements.	Case	Sweep] [1D Timing Sweep] [2D Full Data Eye] [Worst Case Margin Only]
Worst Case Margin Granularity	When “Data Eye Type” is set to “Worst Case Margin Only”, this option is used to determine the granularity of the data capture: per chip select or per nibble.	Per Chip Select	[Per Chip Select] [Per Nibble]
Read Voltage Sweep Step Size	Configures the step size for reading data on voltage margin test.	2	[1] [2] [4]
Read Timing Sweep Step Size	Configures the step size for reading data on timing margin test.	1	[1] [2] [4]
Write Voltage Sweep Step Size	Configures the step size for writing data on voltage margin test.	2	[1] [2] [4]
Write Timing Sweep Step Size	Configures the step size for writing data on timing margin test.	1	[1] [2] [4]

3.3.3.4 NBIO (North Bridge IO) Common Options

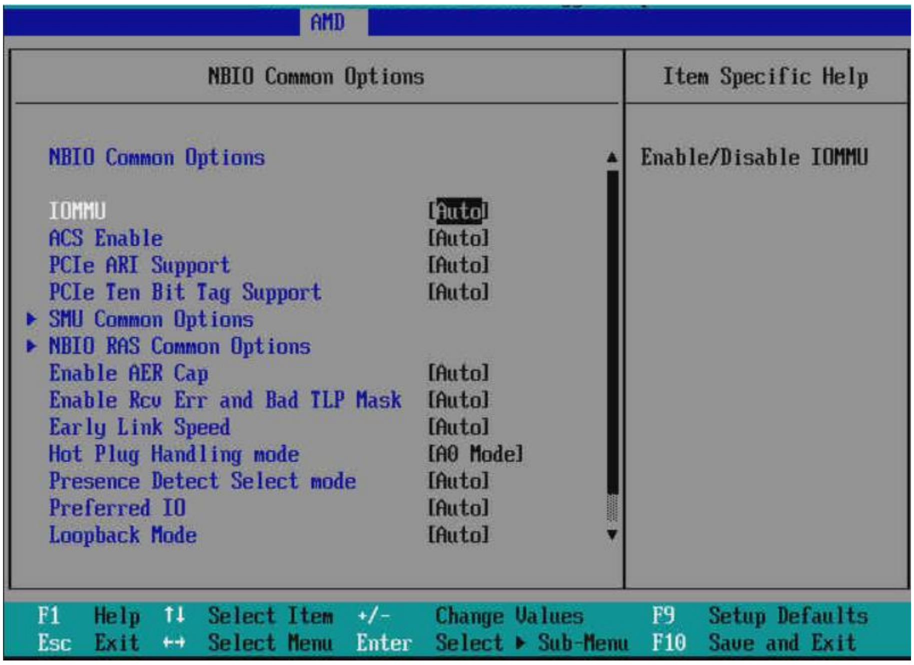
AMD
NBIO Common Options
Item Specific Help
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

```
NBIO Common Options

IOMMU                              [Auto]
ACS Enable                         [Auto]
PCIe ARI Support                   [Auto]
PCIe Ten Bit Tag Support           [Auto]
▸ SMU Common Options
▸ NBIO RAS Common Options
Enable AER Cap                     [Auto]
Enable Rcv Err and Bad TLP Mask    [Auto]
Early Link Speed                   [Auto]
Hot Plug Handling mode             [A0 Mode]
Presence Detect Select mode        [Auto]
Preferred IO                       [Auto]
Loopback Mode                      [Auto]
```

Enable/Disable IOMMU

Figure 49: NBIO Common Options Menu

3.3.3.4.1 IOMMU

The IOMMU (Input-Output Memory Management Unit) is a technology from AMD which is responsible for handling the access to the main memory of the system by external devices. The IOMMU allows virtualization environments to work, once it works as the translation layer between the host physical memory addressing and the guest VMs memory mapping.

For additional Details of IOMMU, please refer to AMD document (3) AMD I/O Virtualization Technology (IOMMU) Specification.

Menu Item	Description	Default Config	Options
IOMMU	Enables/Disables the IOMMU technology.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.2 ACS Enable

The ACS (Access Control Services) is a mechanism that controls PCIe communication in order to make sure that any data transaction is routed to the correct root complex (RC) or endpoint (EP). Therefore, ACS prevents a PCIe device from accidentally or deliberately accessing an unauthorized endpoint, thus protecting the communication against data corruption or malicious access.

The ACS functionality can be used only if AER capability is also enabled, as shown in section 3.3.3.4.5 Enable AER Cap.

Menu Item	Description	Default Config	Options
ACS Enable	Enables/Disables ACS (Access Control Services) for PCIe devices.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.3 PCIe ARI Support

Please refer to the section 3.3.1.5 ARI Support for details regarding the ARI functionality.

Menu Item	Description	Default Config	Options
PCIe ARI Support	Enables/Disables the ARI (Alternative Routing ID) functionality for PCIe devices.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.4 PCIe Ten Bit Tag Support

The “Ten Bit Tag Support” is a feature used by PCIe gen4 devices to improve performance by increasing the “tag” field length of the PCIe header from 8 to 10 bits, thus expanding the number of outstanding non-posted requests available. However, some devices do not support this feature, which may cause issues for the system to initialize. If this is the case, the Ten Bit Tag Support can be disabled using the option below.

Menu Item	Description	Default Config	Options
PCIe Ten Bit Tag Support	Enables/Disables the Ten Bit Tag Support for PCIe devices.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.5 Enable AER Cap

The AER (Advanced Error Reporting) is a feature that allows PCIe devices to report errors with enhanced debug information to the system, allowing better error recovery processes to take place.

Menu Item	Description	Default Config	Options
Enable AER Cap	Enable/Disable the Advanced Error Reporting (AER) feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.6 Enable Rcv Err and bad TLP Mask

This option may be used to mask Receiver Errors (physical layer error) or Bad TLP Errors (correctable error in the PCIe TLP packet) in the PCIe buses, so that the system will not be informed of such error types.

Menu Item	Description	Default Config	Options
Enable Rcv Err and bad TPL Mask	Enables/Disables the masking of Receiver Error and Bad TLP Error for the PCIe links.	Auto	[Auto] [Enable] [Disabled]

3.3.3.4.7 Early Link Speed

The Early Link Speed is a feature that allows the PCIe bus between the CPU and the BMC to train and initialize early in the boot process, before other PCIe buses. This is particularly useful for some system functionalities, once the BMC functions are important to manage and control the system.

This option allows the user to configure which PCIe speed will be used for this early PCIe link establishment between CPU and BMC. It’s highly recommended to keep this option in “Auto”, which is the setting already tested and validated for use in the server.

Menu Item	Description	Default Config	Options
Early Link Speed	Configures the early speed of the PCIe bus between CPU and BMC.	Auto	[Auto] [Gen1] [Gen2]

3.3.3.4.8 Hot Plug Handling Mode

This option configures the hotplug mode for the E1.S NVMe SSDs in the DM-SV01 server. Currently, only the “A0 Mode” option is available, which is the hotplug mode tested and validated for use in the server.

Menu Item	Description	Default Config	Options
Hot Plug Handling Mode	Configures the hotplug mode of operation for the E1.S NVMe SSDs.	A0 Mode	[A0 Mode]

3.3.3.4.9 Presence Detect Select mode

The option below configures the method used by the system to detect the presence of the E1.S NVMe SSDs. It is possible to configure a logic OR or a logic AND between the inband presence detection (using the PCIe bus) and the sideband presence detection (using a specific hardware presence signal). It’s highly recommended to keep this option in “Auto”, which is the setting already tested and validated for use in the server.

Menu Item	Description	Default Config	Options
Presence Detect Select mode	Configures the presence detect mode of operation for the E1.S NVMe SSDs.	Auto	[Auto] [OR] [AND]

3.3.3.4.10 Preferred IO

The following option allows the user to enable the manual configuration of the Preferred IO feature. The Preferred IO is used to improve the performance of a PCIe endpoint, by giving it priority in the I/O transactions.

Please refer to the AMD document (4) “AMD SP3 Family 17h Models 30h–3Fh Preferred IO Usage Guide” for details about the Preferred IO feature and a detailed explanation of how to correctly configure this feature.

Menu Item	Description	Default Config	Options
Preferred IO	Configures the Preferred IO feature to automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.11 Preferred IO Bus

This option requires the Preferred IO option (3.3.3.4.10 Preferred IO) set to manual in order to be shown.

Please refer to the AMD document (4) “AMD SP3 Family 17h Models 30h–3Fh Preferred IO Usage Guide” for details about the Preferred IO feature and a detailed explanation of how to correctly configure this feature.

Menu Item	Description	Default Config	Options
Preferred IO Bus	Configures the PCIe bus number to make use of the Preferred IO feature.	Auto	[<number in decimal>]

3.3.3.4.12 Enhanced Preferred IO Mode

This option requires the Preferred IO option (3.3.3.4.10 Preferred IO) set to manual in order to be shown.

Please refer to the AMD document (4) “AMD SP3 Family 17h Models 30h–3Fh Preferred IO Usage Guide” for details about the Preferred IO feature and a detailed explanation of how to correctly configure this feature. The Enhanced Preferred IO improves the performance of the Preferred IO feature by additionally keeping the PCIe device operating with maximum clock frequencies.

Menu Item	Description	Default Config	Options
Enhanced Preferred IO Mode	Enables/Disables the Enhanced Preferred IO Mode.	Auto	[Disable] [Enable] [Auto]

3.3.3.4.13 Loopback Mode

The Loopback Mode option configures the PCIe lanes receive lanes to directly forward the data to the respective transmit lanes.

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
Loopback Mode	Enables/Disables the Loopback Mode for the PCIe buses.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.14 CV test

This option enables or disables the PCIe test mode for use with the PCIECV tool. This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
CV test	Enables/Disables the test mode for the PCIECV tool.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.15 CAC Weight Adjustment

This option allows configuring the CAC weight algorithm mode of operation. It may be necessary to enable the CAC weight adjustment for some very specific workloads per AMD advice.

Caution: It is recommended to keep this option with default settings, unless the customer is explicitly advised by AMD to change it.

Menu Item	Description	Default Config	Options
CAC Weight Adjustment	Enables/Disables the CAC weight algorithm adjustment.	Auto	[Auto] [Enable] [Disable]

3.3.3.4.16 EDC Control Throttle

This option is used to enable or disable the EDC (Electrical Design Current) shutdown protection (refer to section “3.3.3.4.21.24 EDC Current Tracking” for details regarding EDC). It can be used to avoid EDC shutdown when using some specific high workloads. AMD recommends keeping this option with the default setting.

Menu Item	Description	Default Config	Options
EDC Control Throttle	Enables/Disables the throttling of the EDC (Electrical Design Current).	Auto	[Auto] [Enable] [Disable]

3.3.3.4.17 SRIS

This option is used to enable or disable the SRIS (Separate Refclk Independent SSC Architecture) feature. For details regarding SRIS, please refer to the section “3.3.2.3 SRIS mode debug”.

Menu Item	Description	Default Config	Options
SRIS	Enables/Disables the SRIS feature.	Auto	[Auto] [Enable] [Disable]

3.3.3.4.18 Compliance Loopback

This option is used to enable the loopback mode of the PCIe buses for compliance tests. When the compliance loopback mode is enabled, all hotplug features are disabled.

Menu Item	Description	Default Config	Options
Compliance Loopback	Enables/Disables the loopback mode for PCIe compliance testing.	Auto	[Auto] [Enable] [Disable]

3.3.3.4.19 Multi Upstream Auto Speed Change

This option is used for enabling or disabling the speed change requests from PCIe endpoint devices. By default, Gen2 and Gen3 devices have this option enabled and Gen1 devices have this option disabled.

Menu Item	Description	Default Config	Options
Multi Upstream Auto Speed Change	Enables/Disables the automatic upstream speed change requests from PCIe endpoint devices.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.20 Multi Auto Speed Change On Last Rate

This option configures the PCIe link speed on training as defined below:

- **Enabled:** uses the last data rate advertised.
- **Disabled:** uses the highest data rate ever advertised.

Menu Item	Description	Default Config	Options
Multi Auto Speed Change On Last Rate	Enables/Disables the usage of the last data rate advertised for the PCIe link speed.	Auto	[Auto] [Enable] [Disable]

3.3.3.4.21 SMU (System Management Unit) Common Options

The SMU (System Management Unit) is a subsystem responsible for handling features related to power consumption and temperature monitoring inside the CPU.

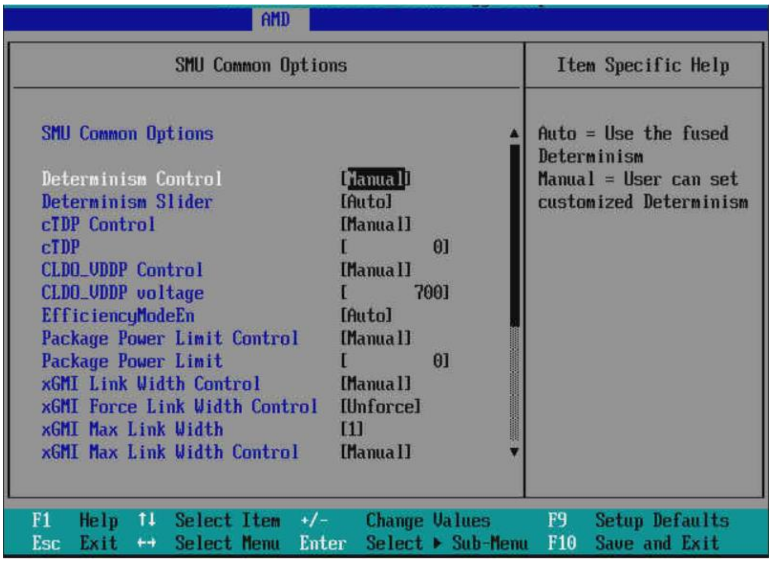

```
                              AMD
                 SMU Common Options                 Item Specific Help

   SMU Common Options                               Auto = Use the fused
                                                    Determinism
   Determinism Control            [Manual]          Manual = User can set
   Determinism Slider             [Auto]            customized Determinism
   cTDP Control                   [Manual]
   cTDP                           [        0]
   CLDO_VDDP Control              [Manual]
   CLDO_VDDP voltage              [      700]
   EfficiencyModeEn               [Auto]
   Package Power Limit Control    [Manual]
   Package Power Limit            [        0]
   xGMI Link Width Control        [Manual]
   xGMI Force Link Width Control  [Unforce]
   xGMI Max Link Width            [1]
   xGMI Max Link Width Control    [Manual]

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 50: SMU Common Options Menu

3.3.3.4.21.1 Determinism Control

This option allows the user to configure the Determinism as automatic or manual. Please refer to section 3.3.3.4.21.2 Determinism Slider for details about the manual determinism configuration.

Menu Item	Description	Default Config	Options
Determinism Control	Configures the Determinism as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.2 Determinism Slider

When Determinism Control (3.3.3.4.21.1 Determinism Control) is configured as manual, the option below is shown and allows the user to choose which Determinism setting will be applied: Performance or Power.

- Power: each CPU can reach a high performance level individually in a set of identically configured CPUs.
- Performance: CPUs performance levels are equalized in a set of identically configured CPUs.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the determinism configuration.

Menu Item	Description	Default Config	Options
Determinism Slider	Selects which Determinism Slider will be used by the CPUs: Power or Performance.	Auto	[Auto] [Power] [Performance]

3.3.3.4.21.3 cTDP Control

The cTDP (Configurable Thermal Design Power) can be set to automatic or manual by means of the option below. Please refer to section 3.3.3.4.21.4 cTDP for details about the manual cTDP configuration.

Menu Item	Description	Default Config	Options
cTDP Control	Configures the cTDP as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.4 cTDP

When cTDP Control (3.3.3.4.21.3 cTDP Control) is configured as manual, the option below is shown. The cTDP (Configurable Thermal Design Power) is the maximum power that can be dissipated by the CPU before it starts throttling (throttling is the process of reducing the CPU frequency in order to limit power). Each CPU from EPYC 7002 family has its own maximum TDP value, so the maximum cTDP value allowed in this option may be different according to the CPU model.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the cTDP configuration.

Menu Item	Description	Default Config	Options
cTDP	Configures the cTDP value in W (watts).	Auto	[<value in watts>]

3.3.3.4.21.5 CLDO_VDDP Control

The cLDO VDDP (adjustable voltage for DRAM memories) can be set to automatic or manual by means of the option below. Please refer to section 3.3.3.4.21.6 CLDO_VDDP voltage for details about the manual cTDP configuration.

Menu Item	Description	Default Config	Options
CLDO_VDDP Control	Configures the CLDO_VDDP as automatic or manual.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.6 CLDO_VDDP voltage

When CLDO_VDDP Control (3.3.3.4.21.5 CLDO_VDDP Control) is configured as manual, the option below is shown. The cLDO_VDDP (Configurable LDO VDDP voltage) setting allows the user to configure the VDDP voltage, which powers the DRAM physical layer interface. The value is configured in millivolts (mv) and the acceptable range varies from 700mv up to 1100mV and a cold reboot is needed for the changes to take effect.

Menu Item	Description	Default Config	Options
CLDO_VDDP voltage	Configures the voltage level of CLDO_VPPD in millivolts (mV). Acceptable range = 700mV up to 1100mV. Note: a cold reset is needed in order for the changes to take effect.	Auto	[<value in mV>]

3.3.3.4.21.7 EfficiencyModeEn

The EfficiencyModeEn option configures the CCLK DPM (Core Clock Dynamic Power Management), which allows the CPU to reduce the core clocks in order to improve the power efficiency. There are two options available:

- Auto: CCLK DPM optimized for performance.
- Enabled: CCLK DPM optimized for power efficiency.

Menu Item	Description	Default Config	Options
EfficiencyModeEn	Configures the CCLK DPM (Core Clock Dynamic Power Management) as auto (performance based) or enabled (power efficiency based).	Auto	[Auto] [Enabled]

3.3.3.4.21.8 Package Power Limit Control

The PPL (Package Power Limit) can be set to automatic or manual by means of the option below. Please refer to section 3.3.3.4.21.9 Package Power Limit for details about the manual PPL configuration.

Menu Item	Description	Default Config	Options
Package Power Limit Control	Configures the PPL (Package Power Limit) as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.9 Package Power Limit

When Package Power Limit Control (3.3.3.4.21.8 Package Power Limit Control) is configured as manual, the option below is shown. The PPL (Package Power Limit) is the maximum power that can be dissipated by the CPU. The CPU controls boost to keep power within the specified limit.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the PPL configuration.

Menu Item	Description	Default Config	Options
Package Power Limit	Configures the PPL (Package Power Limit) value in W (watts).	Auto	[<value in watts>]

3.3.3.4.21.10 xGMI Link Width Control

The xGMI Link Width Control can be set to automatic or manual by means of the option below. A brief description of the xGMI bus can be found at 3.3.3.2.11 Link (GMI and xGMI).

Please refer to sections 3.3.3.4.21.11 xGMI Force Link Width Control and 3.3.3.4.21.13 xGMI Max Link Width Control for details about the manual xGMI Link Width configuration.

Menu Item	Description	Default Config	Options
xGMI Link Width Control	Configures the xGMI link width control as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.11 xGMI Force Link Width Control

When xGMI Link WIdth Control (3.3.3.4.21.10 xGMI Link Width Control) is configured as manual, the option below is shown. The xGMI Force Link Width Control can be set to Forced or Unforced by means of the option below. Please refer to sections 3.3.3.5.11 xGMI Force Link Width Control and 3.3.3.5.12 xGMI Max Link Width Control for details about the manual xGMI Link Width configuration.

Menu Item	Description	Default Config	Options
xGMI Force Link Width Control	Configures the xGMI link width as forced or unforced.	Auto	[Unforce] [Force]

3.3.3.4.21.12 xGMI Force Link Width

When xGMI Force Link WIdth Control (3.3.3.4.21.11 xGMI Force Link Width Control) is configured as forced, the option below is shown. There are three possible settings:

- 0: force the xGMI link width to x2 (2 data lanes).
- 1: force the xGMI link width to x4 (4 data lanes).
- 2: force the xGMI link width to x16 (16 data lanes - maximum supported).

Menu Item	Description	Default Config	Options
xGMI Force Link Width	Force the xGMI link width as x2 (option 0), x4 (option 1) or x16 (option 2).	Auto	[0] [1] [2]

3.3.3.4.21.13 xGMI Max Link Width Control

When xGMI Link Width Control (3.3.3.4.21.10 xGMI Link Width Control) is configured as manual, the option below is shown. The xGMI Maximum Link Width Control can be set to automatic or manual by means of the option below. A brief description of the xGMI bus can be found at 3.3.3.2.11 Link.

Menu Item	Description	Default Config	Options
xGMI Max Link Width Control	Configures the xGMI maximum link width control as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.14 xGMI Max Link Width

When xGMI Max Link Width Control (3.3.3.4.21.13 xGMI Max Link Width Control) is configured as manual, the option below is shown. There are two possible settings:

- 0: configure the xGMI maximum link width to x8 (8 data lanes).
- 1: configure the xGMI maximum link width to x16 (16 data lanes).

Menu Item	Description	Default Config	Options
xGMI Max Link Width	Configures the xGMI maximum link width as x8 (option 0) or x16 (option 1).	Auto	[0] [1]

3.3.3.4.21.15 APBDIS

The APBDIS (Algorithm Performance Boost Disable) is used to configure the behavior of the Power States (P-states) switching by the CPU. There are two possible settings:

- 0: P-states are dynamically switched according to the CPU usage.
- 1: P-states remain fixed.

When APBDIS is set to 1 (APB is disabled), it is possible to configure the value of the fixed P-state of the CPU by means of the option “Fixed SOC Pstate”, as shown in section 3.3.3.4.21.16 Fixed SOC Pstate.

Menu Item	Description	Default Config	Options
APBDIS	Enables (option 0) or disables (option 1) the APB (Algorithm Performance Boost) feature of the CPU.	Auto	[Auto] [0] [1]

3.3.3.4.21.16 Fixed SOC Pstate

When APBDIS (3.3.3.4.21.15 APBDIS) is set to 1, the option below is shown. There are four possible settings:

- P0: configure the P0 P-state (maximum power).

- P1: configure the P1 P-state.
- P2: configure the P2 P-state.
- P3: configure the P3 P-state (minimum power).

Menu Item	Description	Default Config	Options
Fixed SOC Pstate	Configures the fixed Power state (P-state) of the CPU when APBDIS is set to 1.	P0	[P0] [P1] [P2] [P3]

3.3.3.4.21.17 DF Cstates

The AMD CPU can switch its internal Data Fabric to low power states when it is operating in idle. The DF Cstates (Data Fabric C-states) option allows the user to enable or disable this feature.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the DF Cstates configuration.

Menu Item	Description	Default Config	Options
DF Cstates	Enables/Disables the Data Fabric (DF) to switch to low power states when the CPU enters C states.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.18 CPPC

The CPPC (Collaborative Processor Performance Control) is a feature used by the system to exchange performance related data between the hardware and the operating system, allowing the OS to take control of power/efficiency balance operations. This option can be used to enable or disable this feature.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the CPPC feature.

Menu Item	Description	Default Config	Options
CPPC	Enables/Disables the CPPC (Collaborative Processor Performance Control) feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.19 HSMP Support

The HSMP (Host System Management Port) is an interface which allows the OS to access some system management functions. This option can be used to enable or disable this feature.

Please refer to the AMD document (5) “Processor Programming Reference (PPR) for Family 19h Model 01h, Revision B1 Processors” for more details regarding the CPPC feature.

Menu Item	**Description**	**Default Config**	**Options**
HSMP Support	Enables/Disables the HSMP (Host System Management Port) feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.20 Diagnostic Mode

Menu Item	**Description**	**Default Config**	**Options**
Diagnostic Mode	Enables/Disables the Diagnostic Mode.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.21 DLWM Support

The DLWM (Dynamic Link Width Management) feature works at the xGMI bus by reducing the bus width from x16 to x8 in periods of low data traffic. This option allows the user to enable or disable this feature. A brief description of the xGMI bus can be found at 3.3.3.2.11 Link.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the DLWM feature.

Menu Item	**Description**	**Default Config**	**Options**
DLWM Support	Enables/Disables the DLWM (Dynamic Link Width Management) feature.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.21.22 BoostFmaxEn

The BoostFmaxEn control can be set to automatic or manual by means of the option below.

Please refer to section 3.3.3.4.21.23 BoostFmax for details about the manual Boost Frequency configuration.

Menu Item	**Description**	**Default Config**	**Options**
BoostFmaxEn	Configures the Maximum Boost Frequency control as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.21.23 BoostFmax

When BoostFmaxEn (3.3.3.4.21.22 BoostFmaxEn) is configured as manual, the option below is shown. The BoostFmax option allows the user to configure the maximum boost frequency that the CPU core can reach. If the user enters a value higher than the maximum allowed boost frequency, the frequency will be limited and will not reach the configured value.

Please refer to the AMD document (2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers” for more details regarding the Maximum Boost Frequency feature.

Menu Item	Description	Default Config	Options
BoostFmax	Configures the Maximum Boost Frequency value in MHz (Mega Hertz).	Auto	[<value in MHz>]

3.3.3.4.21.24 EDC Current Tracking

EDC stands for “Electrical Design Current” and it is defined as the maximum peak current value that the CPU core can reach. This option allows the user to enable a tracking function for this CPU peak current value in order to generate an MCE (Machine Check Exception) when the threshold defined in “3.3.3.4.21.25 EDC Tracking Current Threshold” has been exceeded.

Menu Item	Description	Default Config	Options
EDC Current Tracking	Enables/Disables the EDC (Electrical Design Current) tracking.	Disable	[Disable] [Enable]

3.3.3.4.21.25 EDC Tracking Current Threshold

When EDC Current Tracking (3.3.3.4.21.24 EDC Current Tracking) is configured to “enabled”, the option below is shown. This option configures the EDC (Electrical Design Current) value that triggers a correctable MCE when exceeded.

Menu Item	Description	Default Config	Options
EDC Tracking Current Threshold	Configures the EDC (Electrical Design Current) value in “amperes” that triggers a MCE when exceeded.	N/A	[<value in amperes>]

3.3.3.4.21.26 EDC Tracking Report Interval

When EDC Current Tracking (3.3.3.4.15.24 EDC Current Tracking) is configured to “enabled”, the option below is shown. This option configures how many times the EDC (Electrical Design Current) value must exceed the limit defined in “3.3.3.4.21.25 EDC Tracking Current Threshold” for generating a correctable MCE.

Menu Item	Description	Default Config	Options
EDC Tracking Report Interval	Configures the number of times the EDC (Electrical Design Current) must exceed the threshold for logging a correctable MCE.	N/A	[<value in decimals>]

3.3.3.4.22 NBIO RAS Common Options

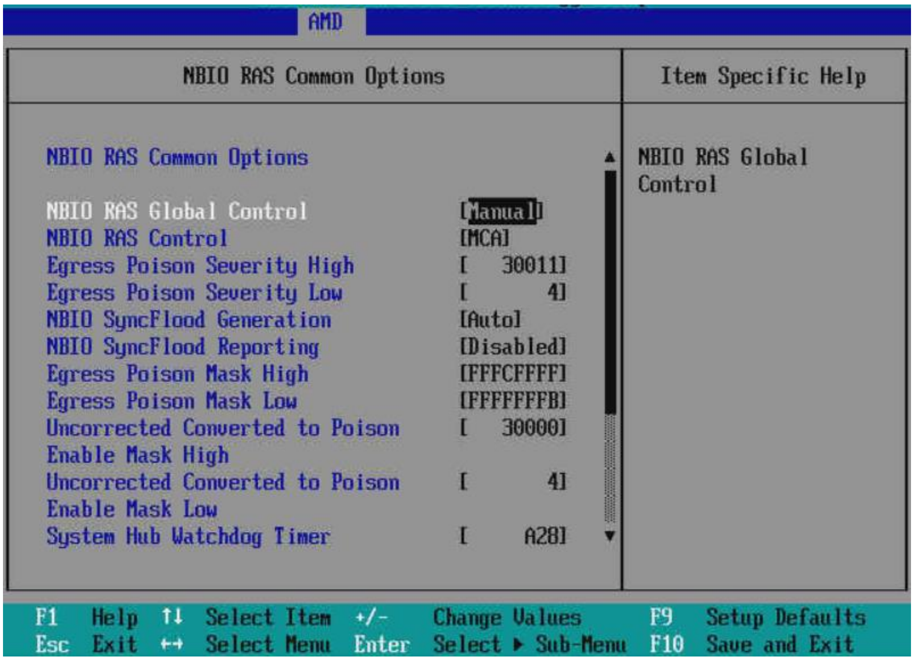

```
                                AMD
            NBIO RAS Common Options                    Item Specific Help

   NBIO RAS Common Options                           NBIO RAS Global
                                                     Control
   NBIO RAS Global Control            [Manual]
   NBIO RAS Control                   [MCA]
   Egress Poison Severity High        [   30011]
   Egress Poison Severity Low         [       4]
   NBIO SyncFlood Generation          [Auto]
   NBIO SyncFlood Reporting           [Disabled]
   Egress Poison Mask High            [FFFCFFFF]
   Egress Poison Mask Low             [FFFFFFFB]
   Uncorrected Converted to Poison    [   30000]
   Enable Mask High
   Uncorrected Converted to Poison    [       4]
   Enable Mask Low
   System Hub Watchdog Timer          [    A28]

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 51: NBIO RAS Common Options Menu

3.3.3.4.22.1 NBIO RAS Global Control

The NBIO RAS (North Bridge I/O - Reliability, Availability, Serviceability) Global Control can be set to automatic or manual by means of the option below.

Please refer to section 3.3.3.4.22.2 NBIO RAS Control for details about the manual NBIO RAS configuration.

Menu Item	Description	Default Config	Options
NBIO RAS Global Control	Configures the NBIO RAS Global Control as automatic or manual.	Auto	[Auto] [Manual]

3.3.3.4.22.2 NBIO RAS Control

When NBIO RAS Global Control (3.3.3.4.22.1 NBIO RAS Global Control) is configured as manual, the option below is shown. The NBIO RAS Control allows the user to select which mechanism will be used for error logging and reporting in the NBIO (North Bridge I/O unit). There are three options available:

- Disabled: no mechanism for error reporting/recovery is active.
- MCA: uses the Machine Check Architecture for error reporting/logging.
- Legacy: uses Legacy functions for error reporting/logging.

Menu Item	Description	Default Config	Options
NBIO RAS Control	Configures the NBIO RAS as disabled, legacy or MCA.	Auto	[MCA] [Legacy] [Disabled]

3.3.3.4.22.3 Egress Poison Severity High

The system marks data as poisoned when an uncorrectable error is detected. Once the data is marked as poisoned, other elements in the system will not consume it, avoiding errors and data corruption. This option is used to configure the severity level of poisoned data for each egress port of the IOHC (I/O Hub Control). Each bit set in this register affects a corresponding IOHC egress port. When the bit is set to binary 1, the severity is set to HIGH. When the bit is set to binary 0, the severity is set to LOW.

The IOHC mapping is out of the scope of this document once this option is used for debug purposes only.

Menu Item	Description	Default Config	Options
Egress Poison Severity High	Configures the severity level (low or high) for each egress port of the IOHC structure of the CPU.	0x30011	[<value in hex>]

3.3.3.4.22.4 Egress Poison Severity Low

This option is the complementary register of Egress Poison Severity High, extending the IOHC ports mapping. Please refer to section 3.3.3.4.22.3 Egress Poison Severity High for a description of the register functionality.

Menu Item	Description	Default Config	Options
Egress Poison Severity Low	Configures the severity level (low or high) for each egress port of the IOHC structure of the CPU.	0x4	[<value in hex>]

3.3.3.4.22.5 NBIO SyncFlood Generation

This option is used to enable or disable the generation of Sync Flood (error messages report mechanism) to the system when the NBIO detects an error. Please refer to section 3.3.3.2.1 Disable

DF to external IP SyncFloodPropagation for a brief description of Sync Flood.

Menu Item	Description	Default Config	Options
NBIO SyncFlood Generation	Enables/Disables the generation of SyncFlood by NBIO when an uncorrectable error is detected.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.22.6 NBIO SyncFlood Reporting

This option is used to enable or disable the reporting of Sync Flood (error messages report mechanism) to the APML (Advanced Platform Management Link) when the NBIO detects an error.

Please refer to section 3.3.3.2.1 Disable DF to external IP SyncFloodPropagation for a brief description of Sync Flood.

Menu Item	Description	Default Config	Options
NBIO SyncFlood Reporting	Enables/Disables the reporting of SyncFlood by NBIO when an uncorrectable error is detected.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.22.7 Egress Poison Mask High

The system marks data as poisoned when an uncorrectable error is detected. Once the data is marked as poisoned, other elements in the system will not consume it, avoiding errors and data corruption. This option is used to mask the poisoned data for each egress port of the IOHC (I/O Hub Control). Each bit set in this register affects a corresponding IOHC egress port. When the bit is set to binary 1, the corresponding poison error is masked so that the system will not take any actions regarding the error.

The IOHC mapping is out of the scope of this document once this option is used for debug purposes only.

Menu Item	Description	Default Config	Options
Egress Poison Mask High	Configures the mask of poison errors for each egress port of the IOHC structure of the CPU.	0xFFFC FFFF	[<value in hex>]

3.3.3.4.22.8 Egress Poison Mask Low

This option is the complementary register of Egress Poison Severity High, extending the IOHC ports mapping. Please refer to section 3.3.3.4.22.7 Egress Poison Mask High for a description of the register functionality.

Menu Item	Description	Default Config	Options
Egress Poison Mask Low	Configures the mask of poison errors for each egress port of the IOHC structure of the CPU.	0xFFFF FFFB	[<value in hex>]

3.3.3.4.22.9 Uncorrected Converted to Poison Enable Mask High

The system marks data as poisoned when an uncorrectable error is detected. Once the data is marked as poisoned, other elements in the system will not consume it, avoiding errors and data corruption. This option is used to mask the poisoned data for each egress port of the IOHC (I/O Hub Control). Each bit set in this register affects a corresponding IOHC egress port. When the bit is set to binary 1, the corresponding uncorrected converted to poison (UCP) error is masked so that the system will not take any actions regarding the error.

The IOHC mapping is out of the scope of this document once this option is used for debug purposes only.

Menu Item	Description	Default Config	Options
Uncorrected Converted to Poison Enable Mask High	Configures the mask of uncorrected converted to poison (UCP) errors for each egress port of the IOHC structure of the CPU.	0x30000	[<value in hex>]

3.3.3.4.22.10 Uncorrected Converted to Poison Enable Mask Low

Menu Item	Description	Default Config	Options
Uncorrected Converted to Poison Enable Mask Low	Configures the mask of uncorrected converted to poison (UCP) errors for each egress port of the IOHC structure of the CPU.	0x4	[<value in hex>]

3.3.3.4.22.11 System Hub Watchdog Timer

This option is used to configure the System Hub (SYSHUB) watchdog timer. The value of the timer is configured in hexadecimal and milliseconds. The default value is 0xA20 which corresponds to 2592ms. When the timer expires without the system refreshing it, an error is generated.

Menu Item	Description	Default Config	Options
System Hub Watchdog Timer	Configures the value of the watchdog timer.	0xA20	[<value in hex>]

3.3.3.4.22.12 SLINK Read Response OK

This option is used to convert a read response error detected in a S-Link to an “okay” response and it is used for debug purposes only.

Menu Item	Description	Default Config	Options
SLINK Read Response OK	Enables/Disables the conversion of a read response error detected in a S-Link to an “okay” response.	Disabled	[Enabled] [Disabled]

3.3.3.4.22.13 SLINK Read Response Error Handling

This option specifies the behavior of the system when a S-Link write response error is detected. There are three possible options:

- Enabled: write response errors are converted to “okay”.
- Trigger MCOMMIT Error: write response errors trigger a MCOMMIT instruction to return an error.
- Log Errors in MCA: the error is logged in the MCA and a fatal error is generated.

Menu Item	Description	Default Config	Options
SLINK Read Response Error Handling	Configures the system behavior when a write response error is detected in a S-Link.	Log Errors in MCA	[Enabled] [Trigger MCOMMIT Error] [Log Errors in MCA]

3.3.3.4.22.14 Log Poison Data from SLINK

This option defines if a poison data propagated to a S-Link will generate a deferred error to be logged or not.

Menu Item	Description	Default Config	Options
Log Poison Data from SLINK	Enables/Disables the generation of a deferred error to be logged when a poison data is propagated to a S-Link.	Disabled	[Enabled] [Disabled]

3.3.3.4.22.15 PCIe Aer Reporting Mechanism

The PCIe Aer Reporting Mechanism option defines the method of reporting AER (Advanced Error Reporting) errors from PCIe buses. There are three options available:

- Firmware First: the Firmware (BIOS/UEFI) is responsible for the error reporting.
- OS First: the OS handles the error through the generation of a system control interrupt (SCI).
- MCA: the server reports the error through MCA (Machine Check Architecture) registers.

Menu Item	Description	Default Config	Options
PCIe Aer Reporting Mechanism	Configures the mechanism used by the system to report AER errors from PCIe devices.	Auto	[Auto] [Firmware First] [OS First] [MCA]

3.3.3.4.22.16 Edpc Control

The Edpc (Enhanced Downstream Port Containment) is a feature from PCIe devices used for error handling and recovery, which allows disabling a single PCIe link when an error is detected. When Edpc is enabled, the system uses the eDPC interrupts mechanism to report fatal or non-fatal errors directly to the operating system, instead of using the AER (Advanced Error Reporting) feature. The option below allows the user to enable or disable the Edpc feature.

Menu Item	Description	Default Config	Options
Edpc Control	Enables/Disables the Edpc (Enhanced Downstream Port Containment) feature.	Disabled	[Auto] [Enabled] [Disabled]

3.3.3.4.22.17 NBIO Poison Consumption

The NBIO Poison Consumption defines the system behavior regarding poison data. When NBIO Poison Consumption is enabled, the system issues a fatal error when poison data is detected on a PCIe device. When the option is disabled, the system propagates the poison data to the PCIe devices.

Menu Item	Description	Default Config	Options
NBIO Poison Consumption	Enables/Disables the consumption of poison data by PCIe devices by means of the NBIO.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.4.22.18 Sync Flood on PCIe Fatal Error

The option below defines if the system must issue a sync flood (fatal error) in the event of a PCIe uncorrectable error that cannot be poisoned.

Menu Item	Description	Default Config	Options
Sync Flood on PCIe Fatal Error	Enables/Disables the sync flood (fatal error) generation in case of a PCIe uncorrectable error that cannot be poisoned.	Auto	[Auto] [True] [False]

3.3.3.5 FCH (Fusion Controller Hub) Common Options

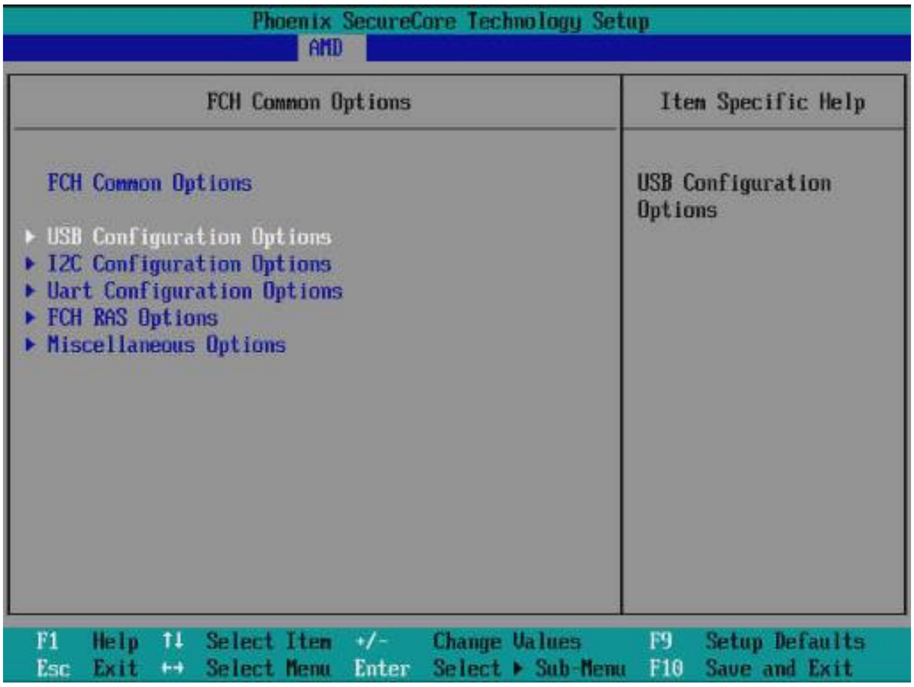
Phoenix SecureCore Technology Setup
AMD
FCH Common Options
Item Specific Help
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

FCH Common Options

▸ USB Configuration Options
▸ I2C Configuration Options
▸ Uart Configuration Options
▸ FCH RAS Options
▸ Miscellaneous Options

USB Configuration Options

Figure 52: FCH Common Options Menu

3.3.3.5.1 USB Configuration Options

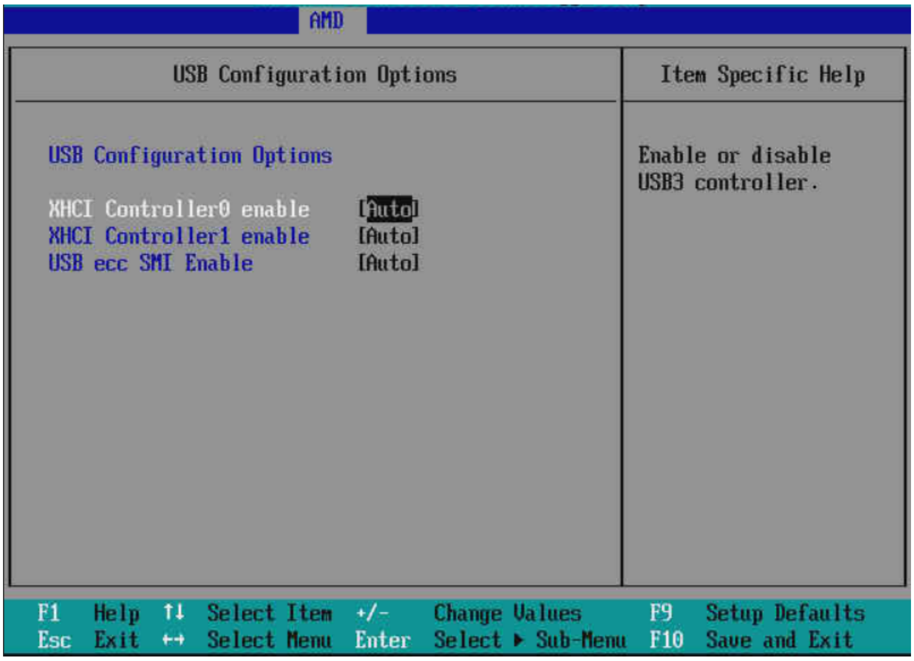
AMD
USB Configuration Options
Item Specific Help
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

USB Configuration Options

XHCI Controller0 enable [Auto]
XHCI Controller1 enable [Auto]
USB ecc SMI Enable [Auto]

Enable or disable USB3 controller.

Figure 53: USB Configuration Options Menu

3.3.3.5.1.1 XHCI Controller0 enable

This option is used to enable or disable the Extensible Host Controller Interface (xHCI) of the front panel USB ports. If the option is set to disabled, both USB ports become unavailable.

Menu Item	Description	Default Config	Options
XHCI Controller0 enable	Enable/Disable the xHCI controller of the USB port 0. [Auto] option keeps the USB port enabled by default.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.5.1.2 XHCI Controller1 enable

This option is used to enable or disable the Extensible Host Controller Interface (xHCI) responsible for the BMC KVM window. If the option is set to disabled, the BMC KVM becomes unavailable.

Menu Item	Description	Default Config	Options
XHCI Controller1 enable	Enable/Disable the xHCI controller of the BMC KVM windows. [Auto] option keeps the USB port enabled by default.	Auto	[Enabled] [Disabled] [Auto]

3.3.3.5.1.3 USB ecc SMI Enable

This option is used to enable or disable the generation of a System Management Interrupt (SMI) when ECC () errors are detected in one of the USB interfaces.

Menu Item	Description	Default Config	Options
USB ecc SMI Enable	Enable/Disable the generation of a System Management Interrupt (SMI) when ECC () errors are detected in one of the USB interfaces.	Auto	[Enable] [Off] [Auto]

3.3.3.5.2 I2C Configuration Options

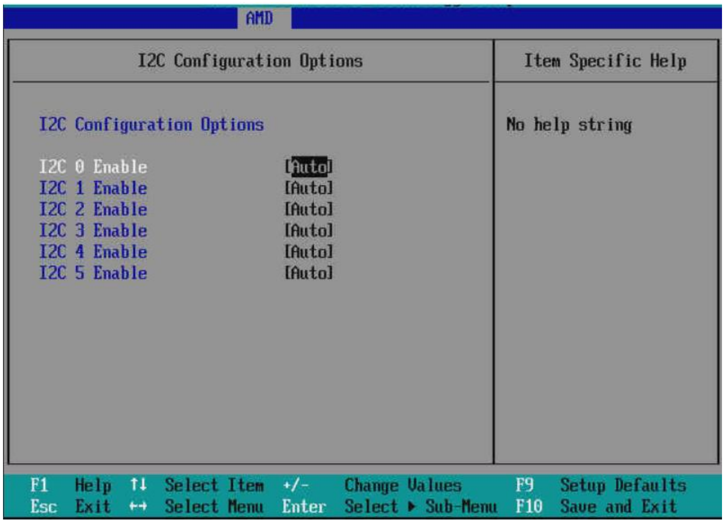

```
                          AMD
        I2C Configuration Options           Item Specific Help

  I2C Configuration Options                 No help string

  I2C 0 Enable              [Auto]
  I2C 1 Enable              [Auto]
  I2C 2 Enable              [Auto]
  I2C 3 Enable              [Auto]
  I2C 4 Enable              [Auto]
  I2C 5 Enable              [Auto]

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 54: USB Configuration Options

3.3.3.5.2.1 I2C (0-5) Enable

This option is for debug purposes only, please keep it with default settings.

Menu Item	Description	Default Config	Options
I2C (0-5) Enable	Enable/Disable the I2C buses.	Auto	[Enabled] [Disabled]

3.3.3.5.3 Uart Configuration Options

This is a read only option, used to check the settings of the UART interfaces. The image below shows an example of the screen showing the summary UART settings.

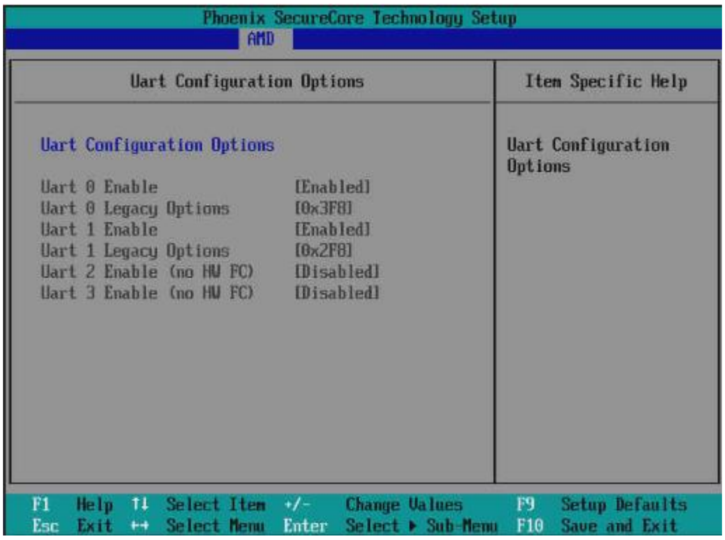

```
                   Phoenix SecureCore Technology Setup
                          AMD
        Uart Configuration Options          Item Specific Help

  Uart Configuration Options                Uart Configuration
                                            Options
  Uart 0 Enable               [Enabled]
  Uart 0 Legacy Options       [0x3F8]
  Uart 1 Enable               [Enabled]
  Uart 1 Legacy Options       [0x2F8]
  Uart 2 Enable (no HW FC)    [Disabled]
  Uart 3 Enable (no HW FC)    [Disabled]

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 55: Uart Configuration Options menu

Menu Item	Description	Default Config	Options
Uart Configuration Options	Read Only menu.	N/A	N/A

3.3.3.5.4 FCH RAS Options

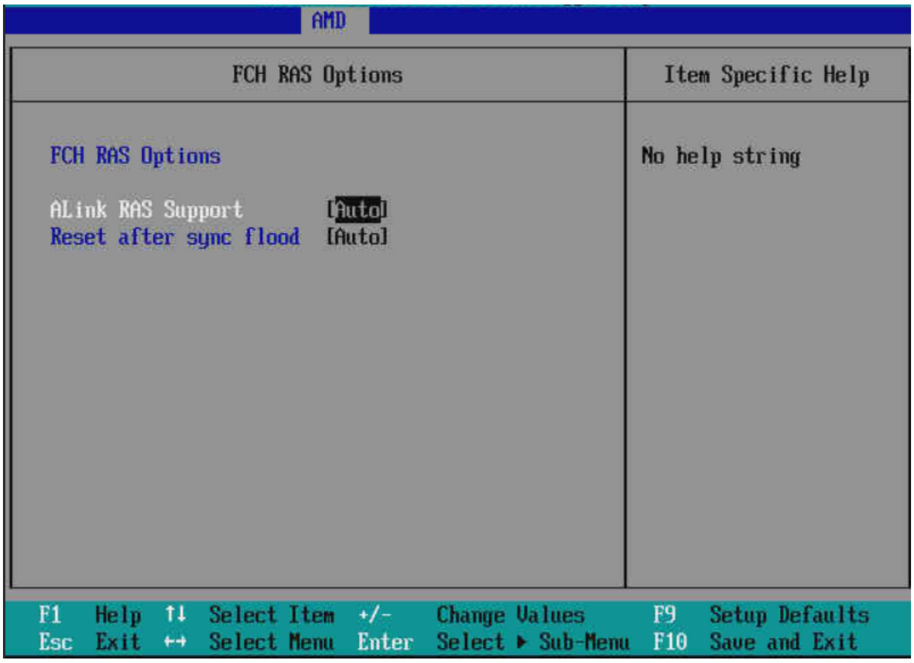

AMD

FCH RAS Options

Item Specific Help

FCH RAS Options

ALink RAS Support [Auto]
Reset after sync flood [Auto]

No help string

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 56: FCH RAS Options Menu

3.3.3.5.4.1 ALink RAS Support

This option, when enabled, provides support for logging parity errors from legacy A-link buses of the PCIe devices.

Menu Item	Description	Default Config	Options
Alink RAS Support	Enable/Disable PCIe A-link RAS for detecting parity errors.	[Auto]	[Enabled] [Disabled]

3.3.3.5.4.2 Reset after sync flood

Defines if the system must be reset when a Fatal Error (Sync Flood) event is detected by the CPU FCH.

Menu Item	Description	Default Config	Options
Reset after sync flood	Enable/Disable the reset of the system after a Fatal Error (Sync Flood) event is detected by the CPU FCH.	[Auto]	[Enabled] [Disabled]

3.3.3.5.5 Miscellaneous Options

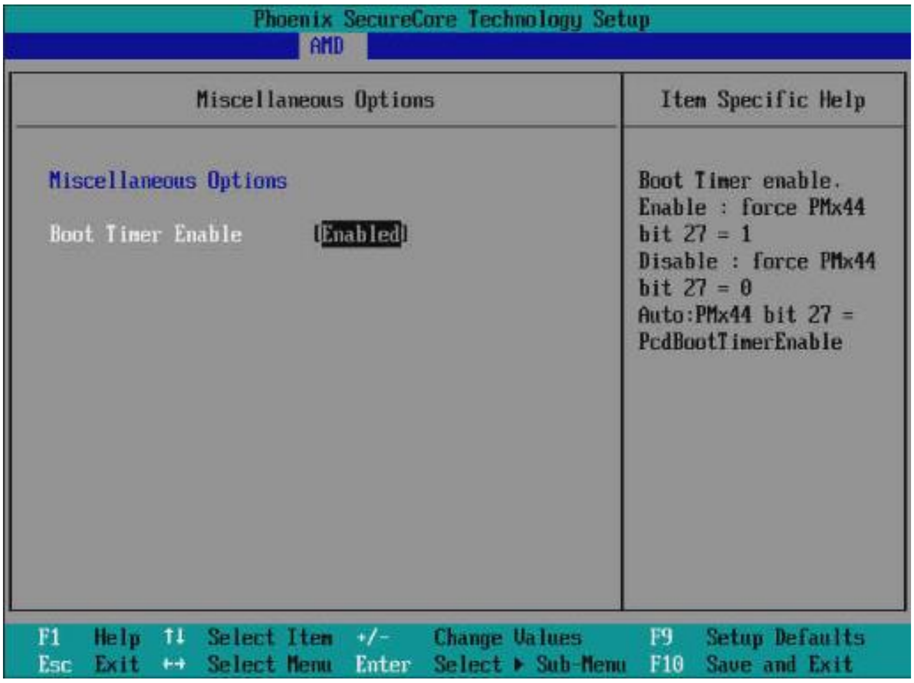

Miscellaneous Options

Boot Timer Enable [Enabled]

Boot Timer enable.
Enable : force PMx44 bit 27 = 1
Disable : force PMx44 bit 27 = 0
Auto:PMx44 bit 27 = PcdBootTimerEnable

Figure 57: Miscellaneous Options menu

3.3.3.5.5.1 Boot Timer Enable

The Boot Timer Enable option allows the user to enable or disable the timer of the FCH (Fusion Controller Hub) module. The boot timer is used to avoid that the BIOS/UEFI firmware hangs or stops responding. If the FW does not disable the timer before 1.17 seconds, then the system will start a warm reset for recovering the system.

Menu Item	Description	Default Config	Options
Boot Timer Enable	Enable/Disable the Boot Timer Enable feature.	Enabled	[Auto] [Enabled] [Disabled]

3.3.3.6 NTB (Non-Transparent Bridge) Common Options

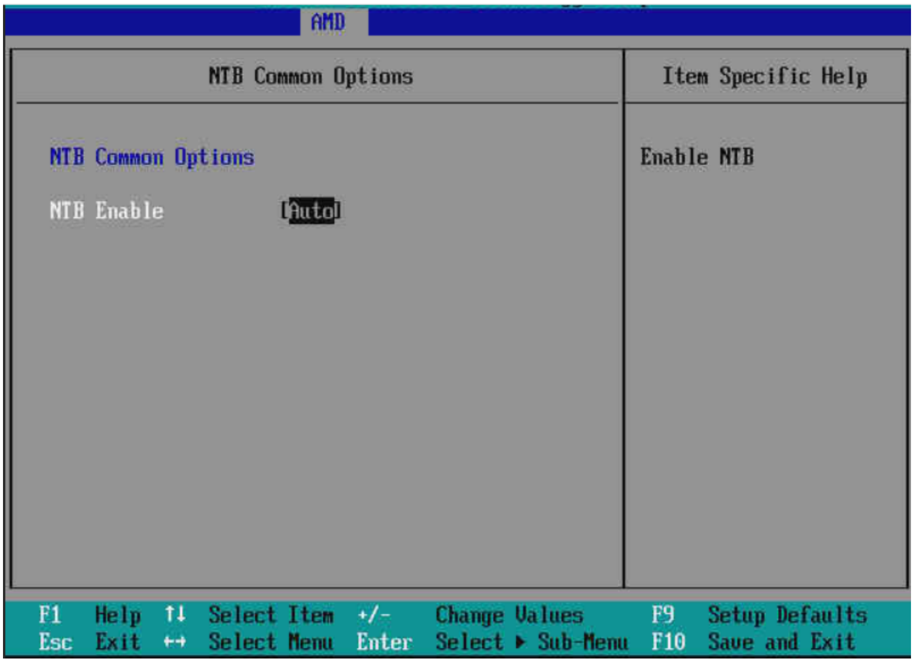

```
AMD
NTB Common Options                         Item Specific Help

NTB Common Options                         Enable NTB

NTB Enable            [Auto]

F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 58: NTB Common Options Menu

3.3.3.6.1 NTB Enable

Standard PCIe implementations are composed of a single RC (root complex) which can communicate with one or more EPs (endpoints).

The NTB (Non-transparent Bridge) is used to overcome this limitation, allowing the PCIe communication between buses from different domains. The NTB is a bridge because it is capable of forwarding the PCIe communication out of the EP domain, interconnecting RCs from different switches. Furthermore, it is “non-transparent” because the PCIe RC is not capable of “seeing” the other RCs interconnected, it just “sees” the EP directly connected.

This option is used to enable the NTB feature for the DM-SV01 PCIe buses. When it is enabled, a complementary set of options are displayed to allow additional configurations for the NTB.

Note: The DM-SV01 does not support NTB. Although some BIOS releases may have this option available, the user must keep NTB menus with default settings.

Menu Item	Description	Default Config	Options
NTB Enable	Enable/Disable the NTB feature of the PCIe buses.	Auto	[Auto] [Enabled]

3.3.3.6.2 NTB Location

Configures the location (CPU socket/die) of the PCIe which will use the NTB function.

Menu Item	Description	Default Config	Options
NTB Location	Selects the location (CPU socket/die) of the PCIe which will use the NTB feature.	Auto	[Auto] [Socket0-Die0] [Socket0-Die1] [Socket0-Die2] [Socket0-Die3] [Socket1-Die0] [Socket1-Die1] [Socket1-Die2] [Socket1-Die3]

3.3.3.6.3 NTB active on PCIeCore

Selects which PCIe core is used for NTB.

Menu Item	Description	Default Config	Options
NTB active on PCIeCore	Selects which PCIe core is used for NTB.	Auto	[Auto] [Core0] [Core1]

3.3.3.6.4 NTB Mode

This option configures the NTB Mode of operation. The following options are available:

- NTB Disabled: disables the NTB.
- NTB Primary: configures the NTB as primary. The primary side of the NTB is connected to one of the PCIe root-complex ports.
- NTB Secondary: configures the NTB as secondary. The secondary side of the NTB is the side working as an endpoint device.
- NTB Random: randomly chooses the NTB configuration.

Menu Item	Description	Default Config	Options
NTB Mode	Configures the NTB mode of operation.	Auto	[NTB Disabled] [NTB Primary] [NTB Secondary] [NTB Random] [Auto]

3.3.3.6.5 Link Speed

Defines the PCIe link speed to run in the NTB connection:

- PCIe Gen1: 2.5GT/s
- PCIe Gen2: 5GT/s

- PCIe Gen3: 8GT/s
- PCIe Gen4: 16GT/s

If the “Max Speed” option is used, then the link speed follows the configuration from 3.2.2.3 Max link speed.

Menu Item	Description	Default Config	Options
Link Speed	Defines the PCIe link speed to run in the NTB connection.	Auto	[Max Speed] [Gen1] [Gen2] [Gen3] [Auto] [Gen4]

3.3.3.7 Soc Miscellaneous Control

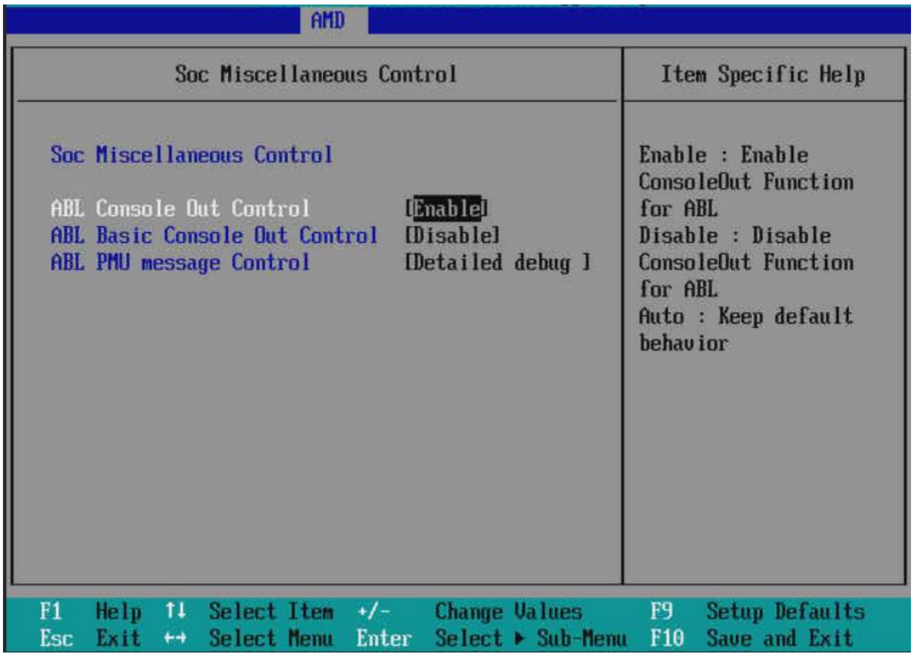

```
                              AMD
        Soc Miscellaneous Control              Item Specific Help

  Soc Miscellaneous Control                    Enable : Enable
                                               ConsoleOut Function
  ABL Console Out Control       [Enable]       for ABL
  ABL Basic Console Out Control [Disable]      Disable : Disable
  ABL PMU message Control       [Detailed debug ]  ConsoleOut Function
                                               for ABL
                                               Auto : Keep default
                                               behavior

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 59: Soc Miscellaneous Control Menu

3.3.3.7.1 ABL Console Out Control

This option is used to enable the ABL (Agesa Boot Loader) logs to be output on console. The logs may be useful for debugging purposes. Please note that enabling the console output logs is expected to increase the boot time of the system.

For details about configuring the console, please refer to section 3.1.4.7 Console Redirection.

When the ABL Console Out Control is enabled, it releases access to the configuration of two additional log controls: “ABL Basic Console Out Control” and “ABL PMU message Control”.

Menu Item	Description	Default Config	Options
ABL Console Out Control	Enable/Disable the console function for outputting ABL logs. [Auto] option keeps the logs disabled by default.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.7.2 ABL Basic Console Out Control

In order to allow the configuration of the “ABL Basic Console Out Control”, the user must previously enable the “ABL Console Out Control” option, as explained in section 3.3.3.7 Soc Miscellaneous Control.

The “ABL Basic Console Out Control” option, when enabled, is used to limit the amount of log information displayed in the console, showing only some basic data and thus reducing the boot time. If this option is kept disabled and the “ABL Console Out Control” is enabled, then the complete ABL logs will be sent to the console.

Menu Item	Description	Default Config	Options
ABL Basic Console Out Control	When enabled, it restricts the console output to display only basic ABL log information. When disabled, it allows the console output to display full ABL log information. [Auto] option keeps the logs disabled by default.	Auto	[Auto] [Enabled] [Disabled]

3.3.3.7.3 ABL PMU message Control

This option controls the log message display level for the PMU (Phy Micro-controller Unit) module, which is the FW responsible for training the DDR memories during the system boot process. There are four possible configurations:

- **Detailed debug message**: show all available messages related to the PMU.
- **Coarse debug message**: show less messages than the detailed debug message option.
- **Stage completion**: display only messages indicating that a PMU related training stage is complete.
- **Firmware completion message only**: display a message only when the whole training procedure by PMU is finished.

Menu Item	Description	Default Config	Options
ABL PMU message Control	Configures the amount of PMU related log messages displayed in the console. [Auto] setting uses the “Stage completion” option by default.	Auto	[Detailed debug message] [Coarse debug message] [Stage completion] [Firmware completion message only] [Auto]

3.3.4 AMD Mem Configuration Status

This menu provides a way to view the summary information about the DDR memory configuration status. When accessing the menu, the user can see a list of DDR memory configuration items and the current configured value for each of them.

The AMD Mem Configuration Status menu has three levels, as described below:

- System level: shows information about the DDR configurations applied to the whole system, i.e., the settings that affect all DDR memory devices connected to the system.
- Socket level: shows information about the DDR configurations applied to the selected CPU socket, i.e., the settings that affect all DDR memory devices connected to the selected CPU socket.
- Channel level: shows information about the DDR configurations applied to the selected DDR channel, i.e., the settings that affect only the DDR memory device connected to the specified channel of the selected socket.

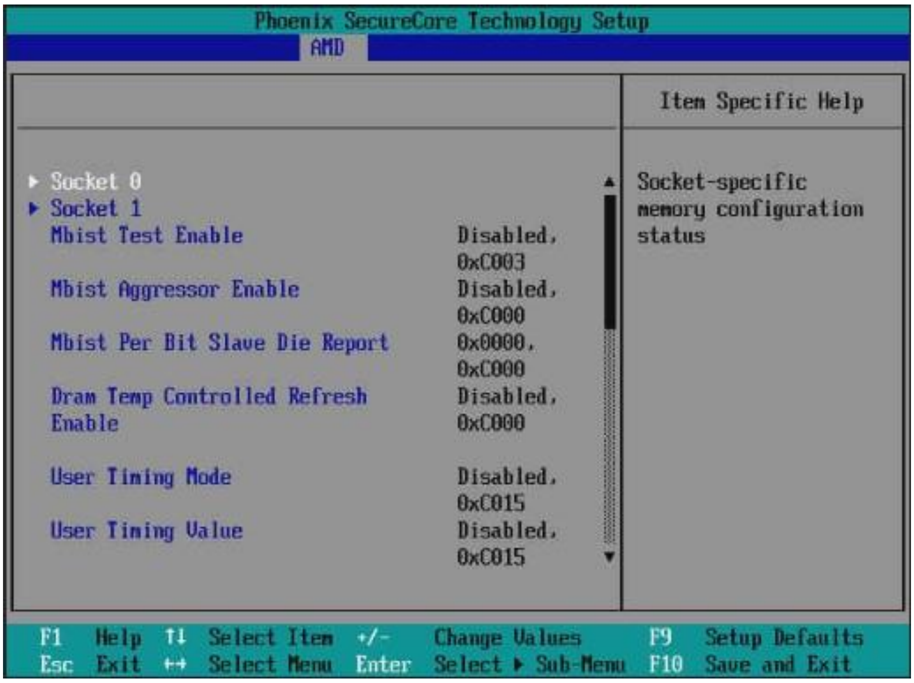

```
                  Phoenix SecureCore Technology Setup
                                  AMD
                                                    Item Specific Help

 ▸ Socket 0                                       Socket-specific
 ▸ Socket 1                                       memory configuration
   Mbist Test Enable              Disabled,       status
                                  0xC003
   Mbist Aggressor Enable         Disabled,
                                  0xC000
   Mbist Per Bit Slave Die Report 0x0000,
                                  0xC000
   Dram Temp Controlled Refresh   Disabled,
   Enable                         0xC000

   User Timing Mode               Disabled,
                                  0xC015
   User Timing Value              Disabled,
                                  0xC015

 F1   Help  ↑↓ Select Item  +/-   Change Values       F9   Setup Defaults
 Esc  Exit  ↔  Select Menu  Enter Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 60: AMD Mem Configuration Status menu - System level

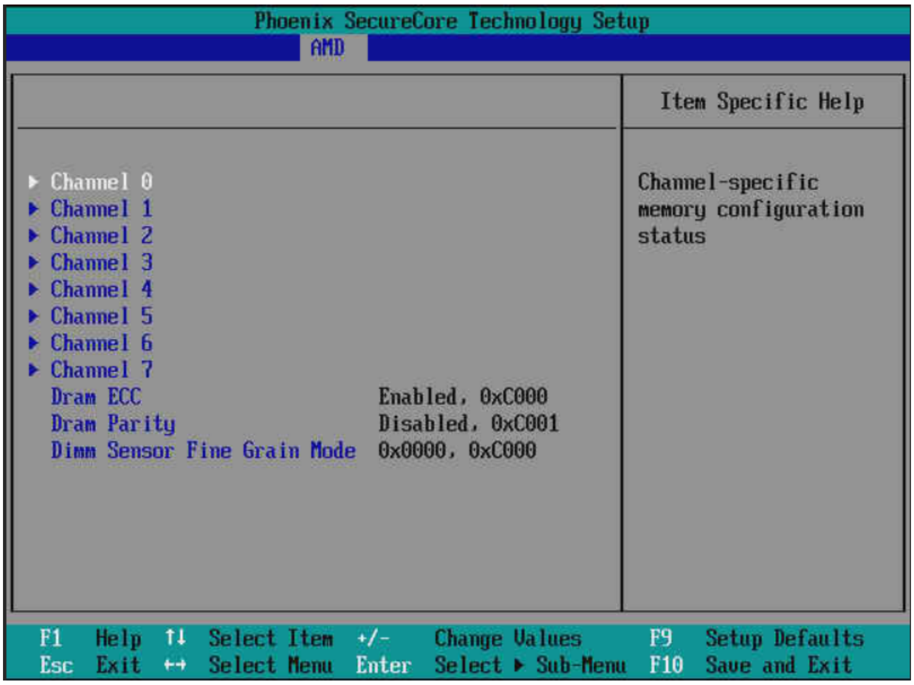

```
                 Phoenix SecureCore Technology Setup
                                AMD

                                                 Item Specific Help

 ▸ Channel 0                                    Channel-specific
 ▸ Channel 1                                    memory configuration
 ▸ Channel 2                                    status
 ▸ Channel 3
 ▸ Channel 4
 ▸ Channel 5
 ▸ Channel 6
 ▸ Channel 7
   Dram ECC                     Enabled, 0xC000
   Dram Parity                  Disabled, 0xC001
   Dimm Sensor Fine Grain Mode  0x0000, 0xC000

 F1   Help  ↑↓ Select Item  +/-   Change Values      F9   Setup Defaults
 Esc  Exit  ↔  Select Menu  Enter Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 61: AMD Mem Configuration Status menu - Socket level

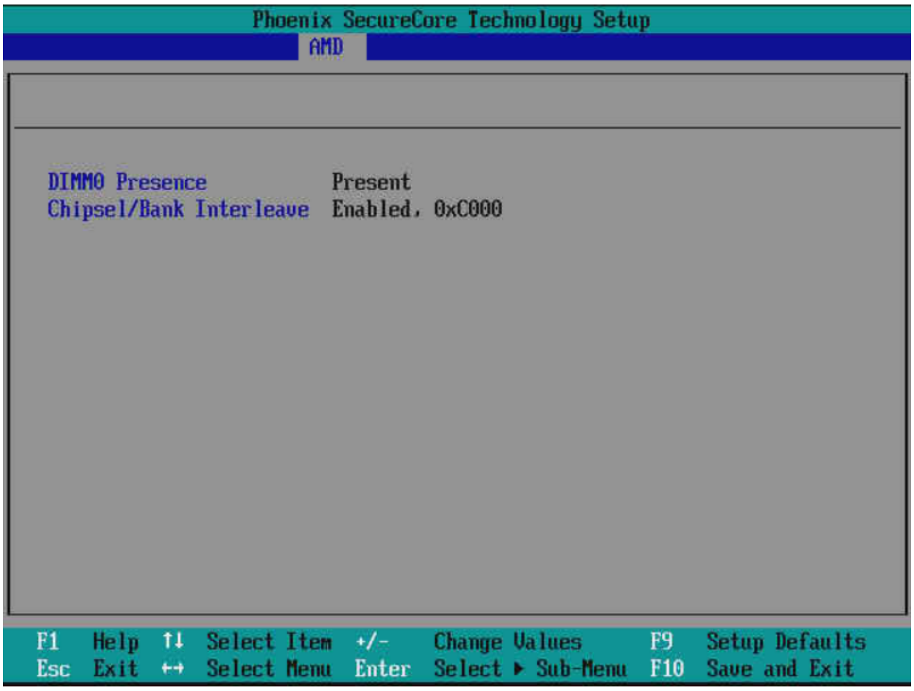

```
                 Phoenix SecureCore Technology Setup
                                AMD

   DIMM0 Presence            Present
   Chipsel/Bank Interleave   Enabled, 0xC000

 F1   Help  ↑↓ Select Item  +/-   Change Values      F9   Setup Defaults
 Esc  Exit  ↔  Select Menu  Enter Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 62: AMD Mem Configuration Status menu - Channel level

Menu Item	**Description**	**Default Config**	**Options**
AMD Mem Configuration Status	Shows the summary information about the DDR memory configuration status	N/A	[Socket 0] [Socket 1]

3.4 Security Menu

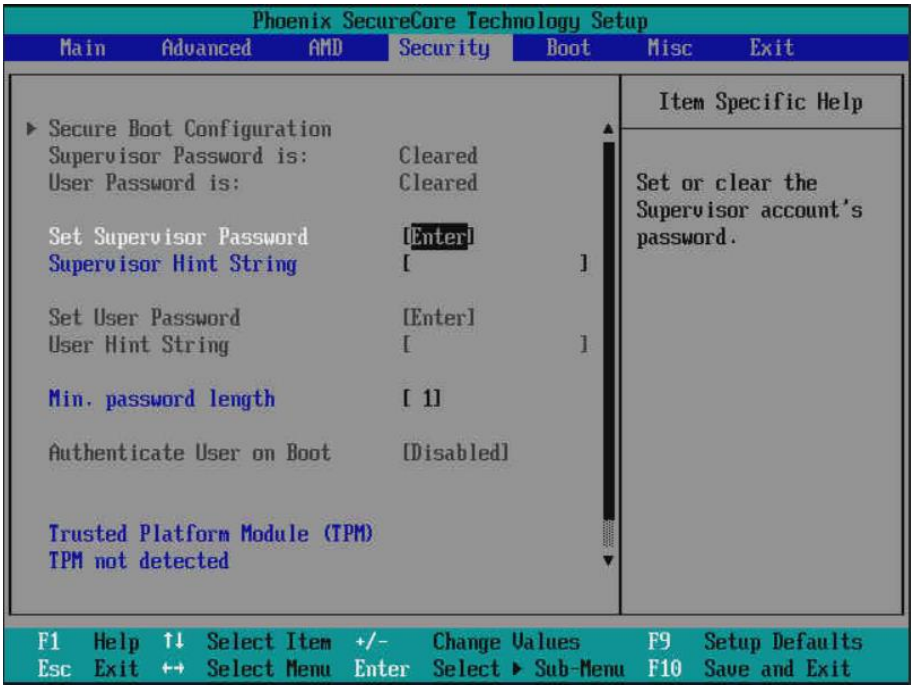

```
                    Phoenix SecureCore Technology Setup
  Main     Advanced     AMD     Security     Boot     Misc     Exit

 ▸ Secure Boot Configuration                             Item Specific Help
   Supervisor Password is:       Cleared
   User Password is:             Cleared                 Set or clear the
                                                         Supervisor account's
   Set Supervisor Password       [Enter]                 password.
   Supervisor Hint String        [           ]

   Set User Password             [Enter]
   User Hint String              [           ]

   Min. password length          [ 1]

   Authenticate User on Boot     [Disabled]


   Trusted Platform Module (TPM)
   TPM not detected

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 63: Security Menu

3.4.1 Set Supervisor Password

This option is used to configure a security password to the Supervisor user to access the UEFI settings. Once the password is configured, the “Secure Boot Configuration” menu is then enabled to be accessed. Please refer to Section 3.4.7 Secure Boot Configuration for additional information.

In order to configure the password, highlight the option “Set Supervisor Password”, then hit <ENTER>. Type the password as required by the menus, and hit <ENTER> to confirm each entry. After the password is set, Save and Exit (3.7.2 Exit Saving Changes). The next time the UEFI menu is going to be accessed, the password will be required.

Menu Item	Description	Default Config	Options
Set Supervisor Password	Set a security password to the Supervisor user (user with full access to the Secure Boot menu) to access UEFI configuration. The maximum number of characters for setting the password is 20.	N/A	Enter a password.

3.4.2 Supervisor Hint String

This option is used to configure a hint string related to the supervisor password set in Section 3.4 Security Menu. In order to configure the hint string, highlight the option “Supervisor Hint String”, then type the hint string and hit <ENTER> to confirm. After the hint string is set, Save and Exit (3.7.2 Exit Saving Changes). The next time the UEFI menu is going to be accessed, the password hint will be available.

The hint string can be shown by selecting the option “See Password Hint”, whenever the password is requested, as shown in the images below.

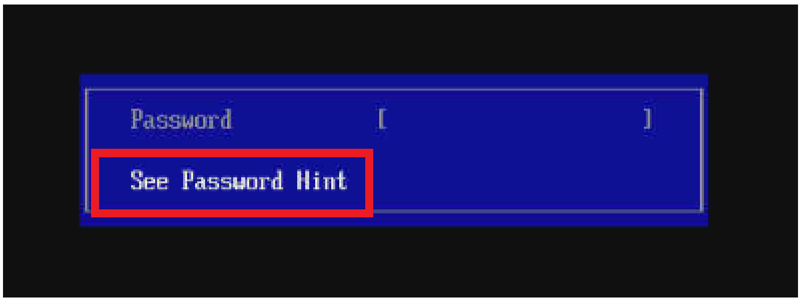

Figure 64: Password Hint menu

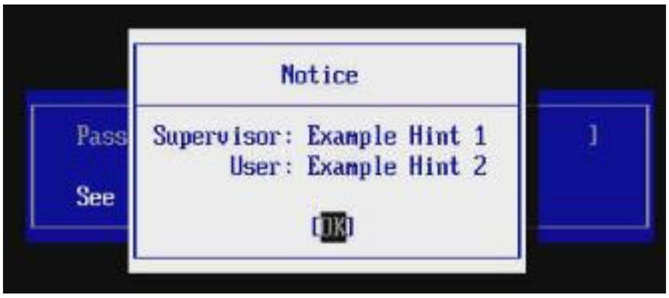

Notice

Supervisor: Example Hint 1
User: Example Hint 2

Figure 65: Password Hint being shown

Menu Item	Description	Default Config	Options
Supervisor Hint String	Set a hint string to help remember the supervisor password.	N/A	Enter a hint string.

3.4.3 Set User Password

This option is used to configure a security password to the regular user to access the UEFI settings. The user configured with this password does not have privileges to access the Secure Boot Menu - Section 3.4.7 Secure Boot Configuration. Only the supervisor user, with the password configured as explained in Section 3.4 Security Menu, can access the Secure Boot Configuration.

In order to configure the password, highlight the option “Set User Password”, then hit <ENTER>. Type the password as required by the menus, and hit <ENTER> to confirm each entry. After the password is set, Save and Exit (3.7.2 Exit Saving Changes). The next time the UEFI menu is going to be accessed, the password will be required.

Menu Item	Description	Default Config	Options
Set Supervisor Password	Set a security password to the regular user (user with no access to Secure Boot menu) to access UEFI configuration.	N/A	Enter a password.

3.4.4 User Hint String

This option is used to configure a hint string related to the user password set in Section 3.4.3 Set User Password. In order to configure the hint string, highlight the option “User Hint String”, then type the hint string and hit <ENTER> to confirm. After the hint string is set, Save and Exit (3.7.2 Exit Saving Changes). The next time the UEFI menu is going to be accessed, the password hint will be available.

The hint string can be shown by selecting the option “See Password Hint”, whenever the password is requested, as shown in the images below.

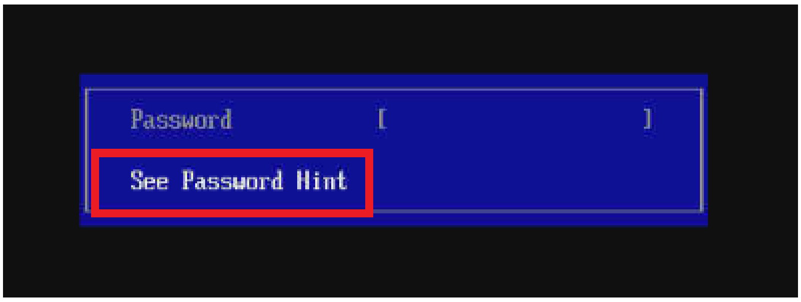

Figure 66: Password Hint menu

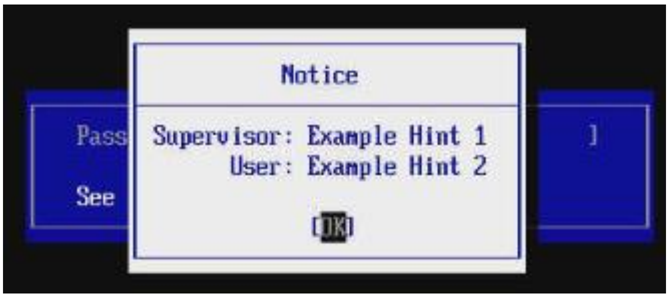
Notice

Supervisor: Example Hint 1
User: Example Hint 2

Figure 67: Password Hint being shown

Menu Item	Description	Default Config	Options
Supervisor Hint String	Set a hint string to help remember the user password.	N/A	Enter a hint string.

3.4.5 Min. Password Length

Menu Item	Description	Default Config	Options
Min. password length	Sets the minimum number of characters acceptable for configuring a password. The acceptable values are in the range from 1 up to 20 characters.	1	[<value in decimal>]

3.4.6 Authenticate User on Boot

Menu Item	Description	Default Config	Options
Authenticate User on Boot	Enables/Disables the password request for the user to access the UEFI menu.	Disabled	[Enabled] [Disabled]

3.4.7 Secure Boot Configuration

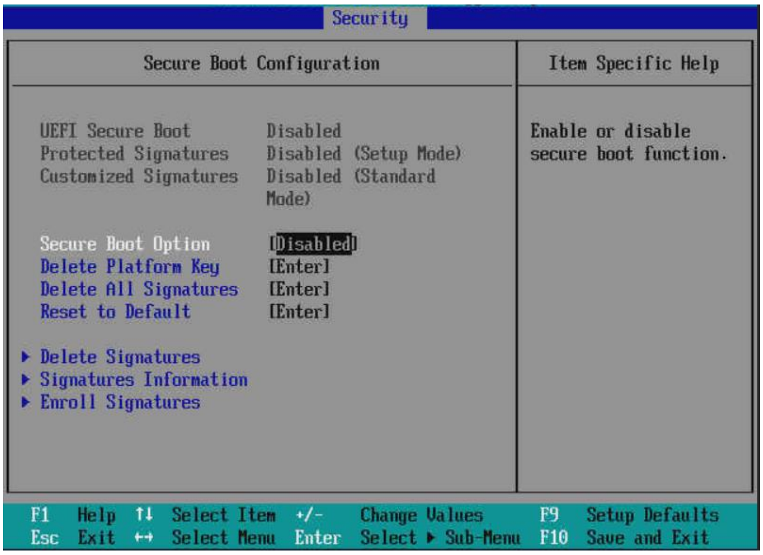

```
                              Security
        Secure Boot Configuration                  Item Specific Help

   UEFI Secure Boot          Disabled               Enable or disable
   Protected Signatures      Disabled (Setup Mode)  secure boot function.
   Customized Signatures     Disabled (Standard
                             Mode)

   Secure Boot Option        [Disabled]
   Delete Platform Key       [Enter]
   Delete All Signatures     [Enter]
   Reset to Default          [Enter]

 ▸ Delete Signatures
 ▸ Signatures Information
 ▸ Enroll Signatures

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

Figure 68: Secure Boot Configuration Menu

3.4.7.1 Secure Boot Option

Menu Item	Description	Default Config	Options
Secure Boot Option	Enables or disables the Secure Boot function.	Disabled	[Enabled] [Disabled]

3.4.7.2 Delete Platform Key

Deletes the Platform Key (PK), keeping the other signatures (KEK, db, dbx) unchanged. When the Platform Key is deleted, the system is changed to setup mode and the Secure Boot is forced to disabled state.

Menu Item	Description	Default Config	Options
Delete Platform Key	Deletes the Platform Key (PK).	N/A	Confirm [Yes] [No]

3.4.7.3 Delete All Signatures

Deletes the Platform Key (PK), the Key Exchange Key (KEK), the allowed signatures database (db) and the forbidden signatures database (dbx). When all the signatures are deleted, the system is changed to setup mode and the Secure Boot is forced to disabled state.

Menu Item	Description	Default Config	Options
Delete All Signatures	Deletes all signatures (PK, KEK, db and dbx).	N/A	Confirm [Yes] [No]

3.4.7.4 Reset to Default

Resets the Platform Key (PK), the Key Exchange Key (KEK), the allowed signatures database (db) and the forbidden signatures database (dbx) to default manufacturing values. When the signatures are reset, the system is changed to Custom Mode.

Menu Item	Description	Default Config	Options
Reset do Default	Resets all signatures (PK, KEK, db and dbx) to default manufacturing values.	N/A	Confirm [Yes] [No]

3.4.7.5 Delete Signatures

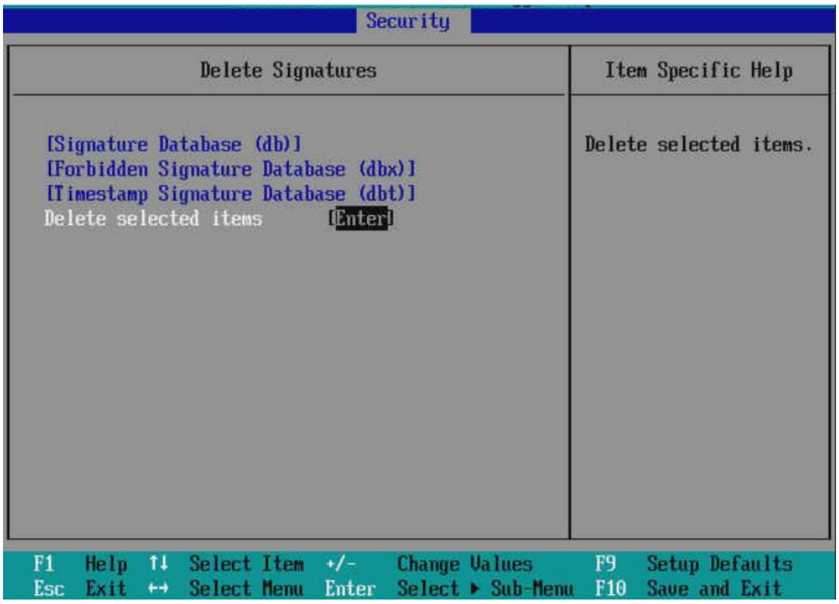

Security

Delete Signatures

[Signature Database (db)]
[Forbidden Signature Database (dbx)]
[Timestamp Signature Database (dbt)]
Delete selected items [Enter]

Item Specific Help

Delete selected items.

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 69: Delete Signatures Menu

Deletes one or more signatures manually selected by the user. In order to delete one or more signatures, the user may follow the procedure below:

1) Access the “Delete Signatures” option.
2) Select the signature you would like to delete and hit <ENTER>.
3) Highlight the option “Select” and hit <ENTER>.
4) Repeat the process for all signatures you would like to delete.
5) Finally, scroll down till the end of the screen, or hit <END>, highlight the option “Delete selected items” and hit <ENTER>.

Menu Item	Description	Default Config	Options
Delete signatures	Deletes one or more signatures manually selected by the user.	N/A	Check the procedure described above.

3.4.7.6 Signatures Information

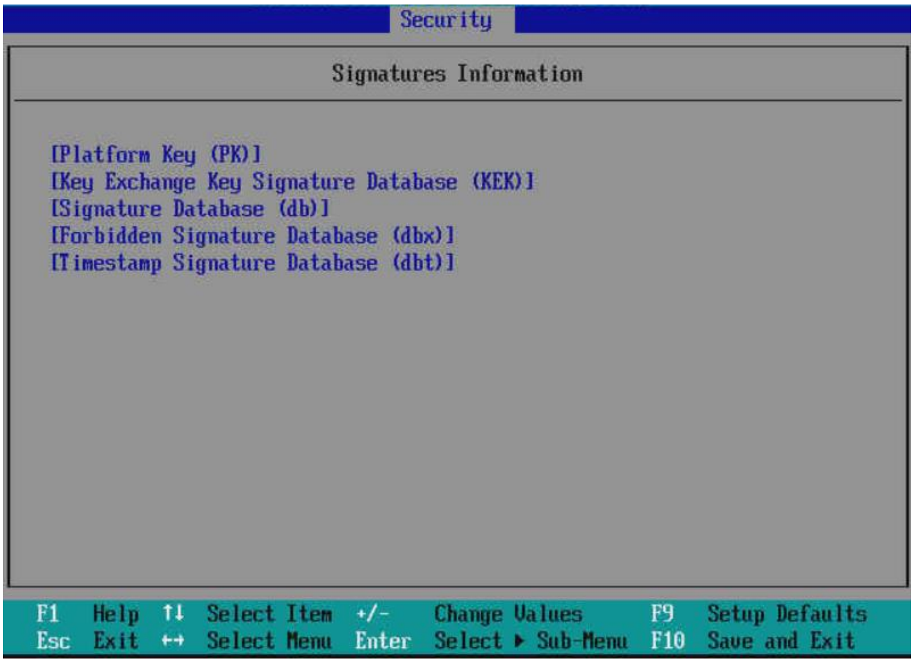

Security

Signatures Information

[Platform Key (PK)]
[Key Exchange Key Signature Database (KEK)]
[Signature Database (db)]
[Forbidden Signature Database (dbx)]
[Timestamp Signature Database (dbt)]

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 70: Signatures Information Menu

Provides information about the currently available signatures for Platform Key (PK), Key Exchange Key (KEK), allowed signatures database (db) and forbidden signatures database (dbx).

Menu Item	Description	Default Config	Options
Signatures Information	Provides information about the currently available signatures.	N/A	N/A

3.4.7.7 Enroll Signatures

Allows the user to manually enroll signatures. The process to performing such a task is described below:

1) Access the “Enroll Signatures” option. The file explorer will be shown.
2) Navigate through the desired file explorer and select the signature file you would like to enroll, then hit <ENTER>. The file extension must be “EFI”, “CER” or “DER”.
3) Select the database you would like to add the certificate to: db, dbx or dbt, then hit <ENTER>.

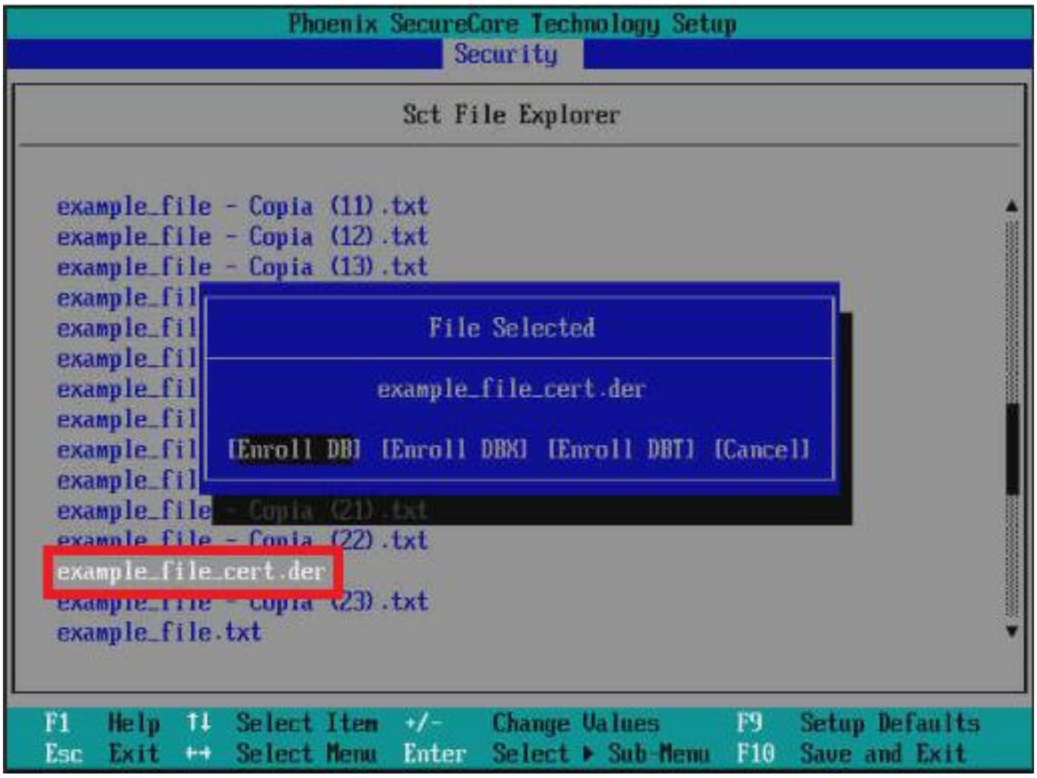
Phoenix SecureCore Technology Setup
Security
Sct File Explorer
example_file - Copia (11).txt
example_file - Copia (12).txt
example_file - Copia (13).txt
File Selected
example_file_cert.der
[Enroll DB] [Enroll DBX] [Enroll DBT] [Cancel]
example_file - Copia (21).txt
example_file - Copia (22).txt
example_file_cert.der
example_file - Copia (23).txt
example_file.txt
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Once the certificate is successfully configured, it can be viewed by accessing the “Signatures Information” option, as explained in section 3.4.7.6 Signatures Information.

Menu Item	**Description**	**Default Config**	**Options**
Enroll Signatures	Allows the user to manually enroll signatures.	N/A	Check the procedure described above.

3.4.8 Trusted Platform Module (TPM)

The TPM is a module used for improving the security level of the system. The TPM module for the DM-SV01 system is optional and so it is sold separately. The settings described in this chapter are valid only if the TPM device is assembled in the DM-SV01 motherboard.

3.4.8.1 Current Selected TPM Device

Configures the TPM device mode of operation. This option can be accessed only if the optional TPM hardware module is populated in the motherboard and if it has been previously activated as explained in section 3.2.1.1 TPM. The possible configurations are TPM 1.2 and TPM 2.0.

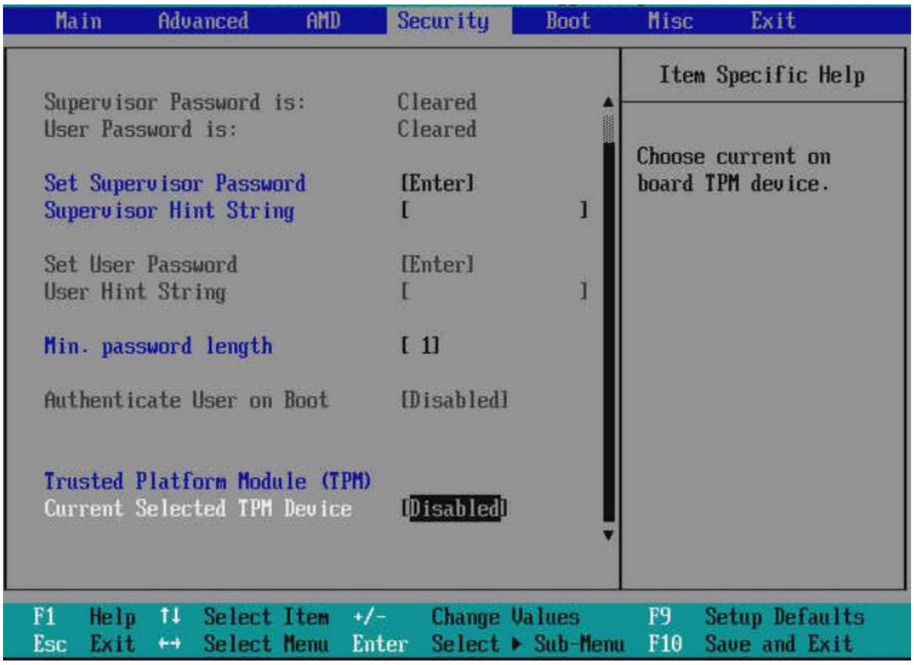

```
   Main     Advanced     AMD     Security     Boot     Misc     Exit

  Supervisor Password is:        Cleared           Item Specific Help
  User Password is:              Cleared
                                                   Choose current on
  Set Supervisor Password        [Enter]           board TPM device.
  Supervisor Hint String         [         ]

  Set User Password              [Enter]
  User Hint String               [         ]

  Min. password length           [ 1]

  Authenticate User on Boot      [Disabled]


  Trusted Platform Module (TPM)
  Current Selected TPM Device    [Disabled]

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 71: TPM configuration

Menu Item	Description	Default Config	Options
Current Selected TPM Device	Configures the TPM's mode of operation in TPM 1.2 or TPM 2.0.	Disabled	[Disabled] [TPM 2] [TPM 1.2]

3.4.8.2 TPM Configuration

3.4.8.2.1 TPM Action

The TPM Action menu is used to send a command to the TPM device. This menu becomes available as soon as the TPM is configured as TPM 1.2 or TPM 2.0, as explained in section 3.4.8 Trusted Platform Module (TPM).

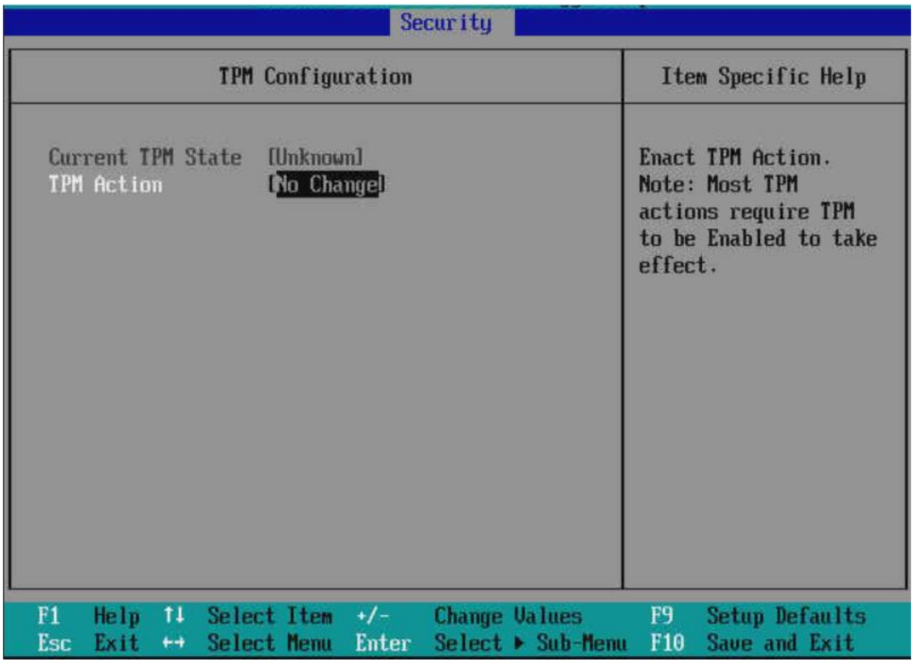

```
                              Security
            TPM Configuration                 Item Specific Help

   Current TPM State    [Unknown]           Enact TPM Action.
   TPM Action           [No Change]         Note: Most TPM
                                            actions require TPM
                                            to be Enabled to take
                                            effect.

 F1   Help  ↑↓  Select Item  +/-   Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter Select ▸ Sub-Menu   F10  Save and Exit
```

Figure 72: TPM Action menu

The following options are available:

- No change
- Enable
- Enable and Activate
- Set Owner Install, with state=True
- Set Owner Install, with state=False
- Enable, Activate, and Set Owner Install with state=True
- Enable, Activate, and Set Owner Install with state=False
- Require PP for provisioning
- Require PP for clear
- Do not Require PP for clear
- Enable, Activate and Clear
- Enable, Activate, Clear, Enable, and Activate

Please see below a brief explanation of each available item:

- **No change**: default value. No command is sent to the TPM device.
- **Enable**: enables the TPM device, making most features available.
- **Activate**: differently from the enable action, the activation command can be done only through physical presence, and requires a restart to take action.
- **Set Owner Install, with state=True**
- **Set Owner Install, with state=False**
- **Require PP**: when a function requires Physical Presence (PP), it means it is only possible to perform such an operation by physically and directly accessing the UEFI menu by an authorized user.
- **Clear**: clears all the ownership information of the TPM device, removing all keys and related data.

Menu Item	Description	Default Config	Options
TPM Action	Sends a command to the TPM device to perform the selected operation.	No change	[No change] [Enable] [Enable and Activate] [Set Owner Install, with state=True] [Set Owner Install, with state=False] [Enable, Activate, and Set Owner Install with state=True] [Enable, Activate, and Set Owner Install with state=False] [Require PP for provisioning] [Require PP for clear] [Do not Require PP for clear] [Enable, Activate and Clear] [Enable, Activate, Clear, Enable, and Activate]

When the TPM component has already been activated and UEFI recognizes its current state, the menu items for TPM configuration changes and the below options become available.

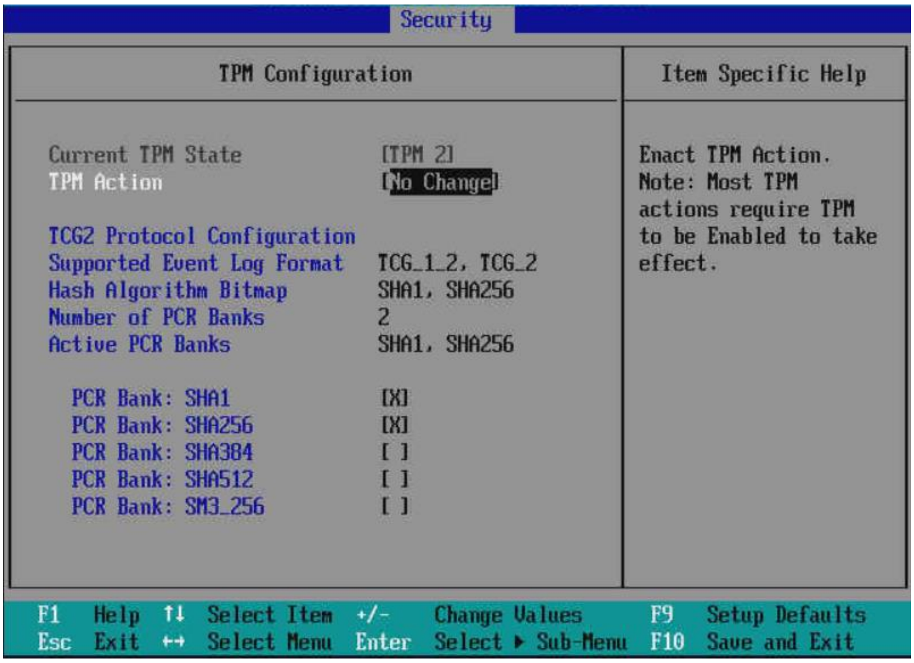

```
TPM Configuration                                   Item Specific Help

Current TPM State              [TPM 2]              Enact TPM Action.
TPM Action                     [No Change]          Note: Most TPM
                                                    actions require TPM
TCG2 Protocol Configuration                         to be Enabled to take
Supported Event Log Format     TCG_1_2, TCG_2       effect.
Hash Algorithm Bitmap          SHA1, SHA256
Number of PCR Banks            2
Active PCR Banks               SHA1, SHA256

  PCR Bank: SHA1               [X]
  PCR Bank: SHA256             [X]
  PCR Bank: SHA384             [ ]
  PCR Bank: SHA512             [ ]
  PCR Bank: SM3_256            [ ]
```

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 73: TPM Configuration

- No change
- TPM2 HierarchyControl (TPM_RH_OWNER_YES, TPM_RH_ENDORSEMENT YES)
- TPM2 HierarchyControl (TPM_RH_OWNER_NO, TPM_RH_ENDORSEMENT NO)
- TPM2 ClearControl (NO) + Clear
- TPM2 PCR_Allocate (Algorithm IDs)
- TPM2 Change EPS
- TPM2 LogAllDigests
- TPM2 HierarchyControl (TPM_RH_OWNER_NO, TPM_RH_ENDORSEMENT YES)

Menu Item	Description	Default Config	Options
TPM Action	Sends a command to the TPM device to perform the selected operation.	No change	[No change] [TPM2 HierarchyControl (TPM_RH_OWNER_YES, TPM_RH_ENDORSEMENT YES)] [TPM2 HierarchyControl (TPM_RH_OWNER_NO, TPM_RH_ENDORSEMENT NO)] [TPM2 ClearControl (NO) + Clear] [TPM2 PCR_Allocate (Algorithm IDs)] [TPM2 Change EPS] [TPM2 LogAllDigests] [TPM2 HierarchyControl (TPM_RH_OWNER_NO, TPM_RH_ENDORSEMENT YES)]

The items available for configuration are described below:

- **HierarchyControl - TPM_RH_OWNER**: configures the authorization to use the storage root hierarchy key (YES or NO).
- **HierarchyControl - TPM_RH_ENDORSEMENT**: configures the authorization to use the endorsement root hierarchy key (YES or NO).
- **Clear**: clears all the ownership information of the TPM device, removing all keys and related data.
- **PCR_Allocate**: configures the enabling of one or more PCR banks to be allocated by TPM. When the PCR_allocate option is configured, the user can input the values of the PCR backs to be allocated in the menu called “TPM2 Operation Parameter”.
- **Change EPS**: allows the user to reset the Endorsement Primary Seed of the TPM, which is the root to generate all the other seeds, keys and additional values inside the TPM.
- **Log All Digests**: enables the logging of the digest or, in other words, the result of the hashing algorithm in the TPM’s PCR banks.

3.4.8.2.2 PCR Bank (Hashing algorithms configuration)

Allows the user to enable/disable specific hashing algorithms to be used in the TPM’s PCR banks.

Menu Item	Description	Default Config	Options
PCR Bank	Enables/disables specific hashing algorithms to be used in the TPM’s PCR banks.	SHA1 and SHA256 Enabled	[Enable/Disable SHA1] [Enable/Disable SHA256] [Enable/Disable SHA384] [Enable/Disable SHA512] [Enable/Disable SM3_256]

3.5 Boot Menu

3.5.1 Boot Priority Order

The boot Priority Order menu allows the user to select the devices that the operating system will try to boot from. The BIOS will try to boot from the device with the priority number 1; if not possible, then it tries to boot from device numbered as 2 and so on. The figure 74 below shows the Boot Priority Order screen.

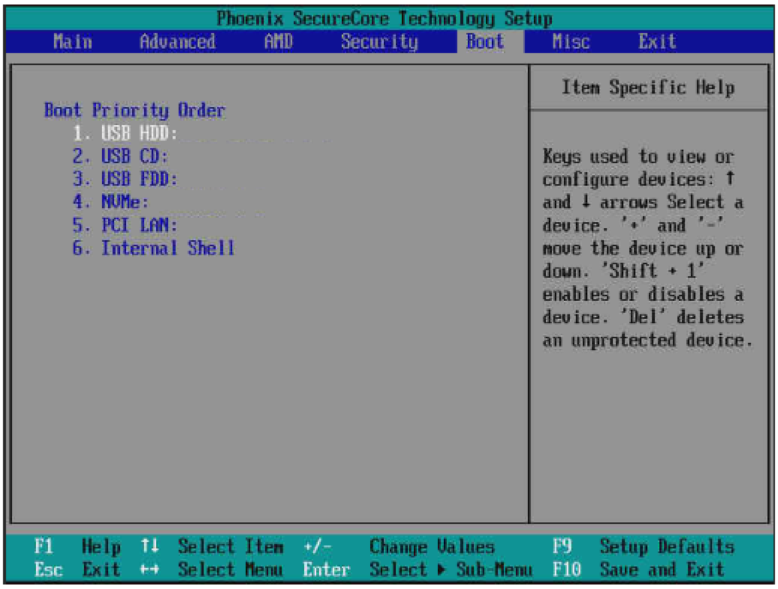

```
                    Phoenix SecureCore Technology Setup
   Main     Advanced     AMD     Security     Boot     Misc     Exit

  Boot Priority Order                          Item Specific Help
     1. USB HDD:
     2. USB CD:                              Keys used to view or
     3. USB FDD:                             configure devices: ↑
     4. NVMe:                                and ↓ arrows Select a
     5. PCI LAN:                             device. '+' and '-'
     6. Internal Shell                       move the device up or
                                             down. 'Shift + 1'
                                             enables or disables a
                                             device. 'Del' deletes
                                             an unprotected device.

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 74: Boot Priority Order screen

The boot priority order can be changed by the following procedure:

1) Use the arrow keys to select the boot device you'd like to change.
2) Press <+> key if you want to move the device up in the boot priority order, or Press <-> key if you want to move the device down in the boot priority order.

Users can also use the <Shift + 1> key to disable a specific device. Disabled devices are ignored during the boot process, independent of its position in the boot priority order. If a device is disabled, it will be marked with an exclamation point (!) in the left side of the priority number, as shown in figure 75. Disabled devices can be enabled again by using the <Shift + 1> key as well.

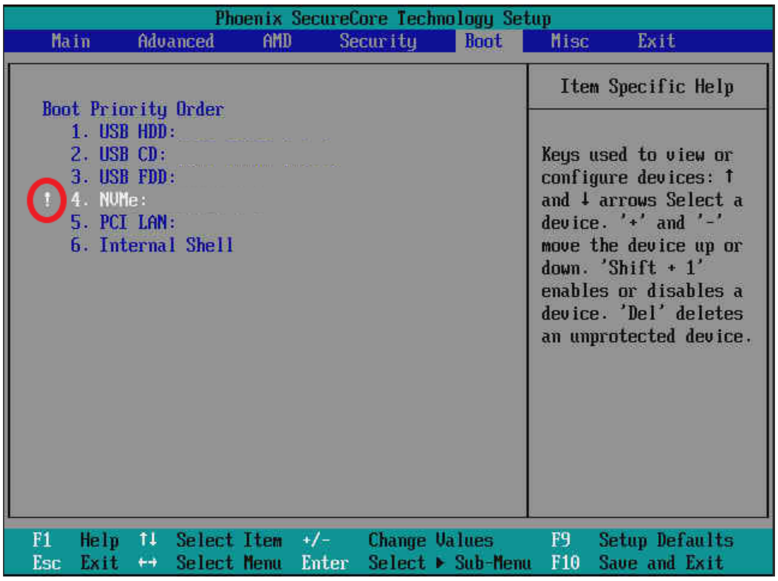

```
                    Phoenix SecureCore Technology Setup
   Main     Advanced     AMD     Security     Boot     Misc     Exit

  Boot Priority Order                          Item Specific Help
     1. USB HDD:
     2. USB CD:                              Keys used to view or
     3. USB FDD:                             configure devices: ↑
    !4. NVMe:                                and ↓ arrows Select a
     5. PCI LAN:                             device. '+' and '-'
     6. Internal Shell                       move the device up or
                                             down. 'Shift + 1'
                                             enables or disables a
                                             device. 'Del' deletes
                                             an unprotected device.

 F1   Help  ↑↓  Select Item  +/-    Change Values      F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu  F10  Save and Exit
```

Figure 75: Boot Option disabled

3.6 Misc Menu

The Misc (miscellaneous) menu is a variable screen, which presents to the user configuration options from plug-in devices connected to the server, when these devices provide this functionality.

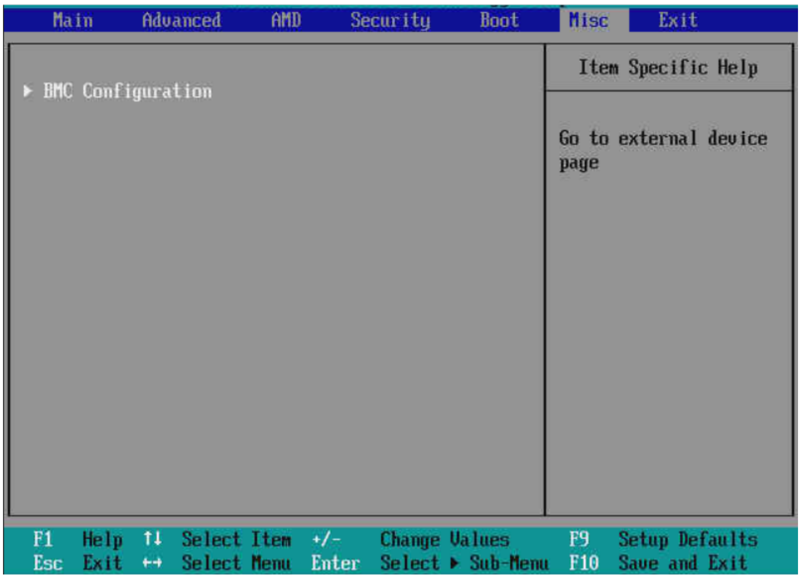

Main Advanced AMD Security Boot Misc Exit

▸ BMC Configuration

Item Specific Help

Go to external device page

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 76: Misc Menu

Several plug-in devices, such as NVMe disks, Ethernet adapters or other PCIe cards, have UEFI compatible software recorded on its own volatile memory devices (Option ROM). When the UEFI is loaded, it reads the contents of these ROMs and loads the UEFI compatible SW provided by the device manufacturer. Some of these SWs provide configuration menus to the user, to configure some card-specific settings. The configuration menus provided by these Option ROMs FW are shown to the user in the Misc Menu. The figure below shows the menu screen when a Network Adapter Card Option ROM is available.

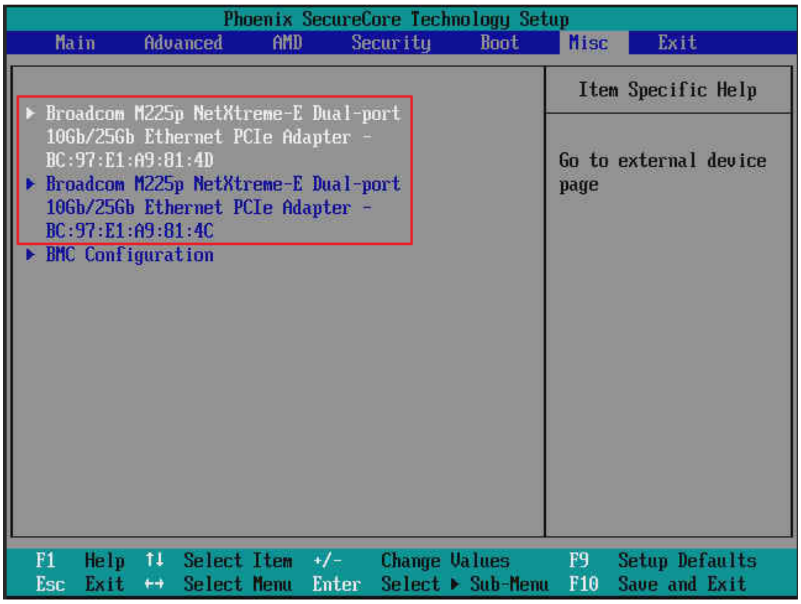

Phoenix SecureCore Technology Setup

Main Advanced AMD Security Boot Misc Exit

▸ Broadcom M225p NetXtreme-E Dual-port 10Gb/25Gb Ethernet PCIe Adapter - BC:97:E1:A9:81:4D
▸ Broadcom M225p NetXtreme-E Dual-port 10Gb/25Gb Ethernet PCIe Adapter - BC:97:E1:A9:81:4C
▸ BMC Configuration

Item Specific Help

Go to external device page

F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

Figure 77: Misc menu with NIC

3.6.1 BMC Configuration

The BMC Configuration menu shows some information about the server's inventory and also provides some specific settings controlled by the BMC.

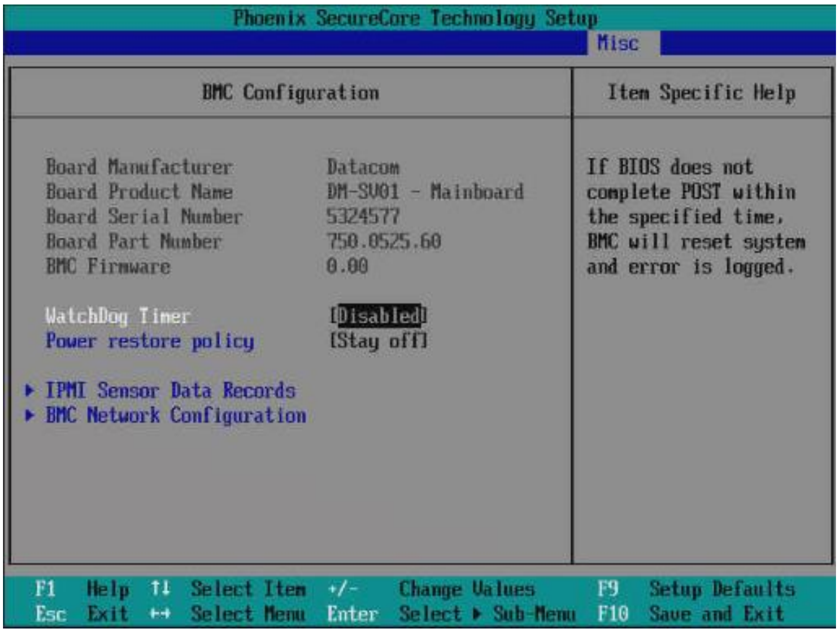

```
                    Phoenix SecureCore Technology Setup
                                                       Misc
            BMC Configuration                      Item Specific Help

  Board Manufacturer        Datacom                If BIOS does not
  Board Product Name        DM-SV01 - Mainboard    complete POST within
  Board Serial Number       5324577                the specified time,
  Board Part Number         750.0525.60            BMC will reset system
  BMC Firmware              0.00                   and error is logged.

  WatchDog Timer            [Disabled]
  Power restore policy      [Stay off]

▶ IPMI Sensor Data Records
▶ BMC Network Configuration

 F1   Help  ↑↓  Select Item  +/-    Change Values       F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▶ Sub-Menu   F10  Save and Exit
```

Figure 78: BMC Configuration menu

3.6.1.1 WatchDog Timer

This option sets a watchdog timer to wait for BIOS to complete the Power On Self Test (POST). If the watchdog timer is enabled and the POST is not completed within the specified time, then BMC will reset the system and log an error.

Menu Item	Description	Default Config	Options
Watchdog Timer	Timer to wait for BIOS to complete POST.	Disabled	[Enabled] [Disabled]

3.6.1.2 Power restore policy

The power restore policy option defines the behavior of the CPU turn on/off functionality when the system is powered on after an event of power loss has occurred before. Three different behaviors can be configured:

1) Stay Off: the BMC initializes and keeps the CPUs turned off. User should manually send a Power On command to turn the CPUs on.
2) Power On: the BMC initializes and automatically turns the CPUs on.
3) Restored to last state: the BMC initializes and restores the CPUs to the last state when the power loss event has occurred:
    a) If the CPUs state was “turned off” when power loss has occurred, BMC behaves like the “Stay Off” configuration.

b) If the CPUS state was “turned on” when power loss has occurred, BMC behaves like the “Stay On” configuration.

Menu Item	Description	Default Config	Options
Power Restore Policy	Defines the power on policy when Power is applied to the system.	Stay Off	[Stay Off] [Restored to last state] [Power on]

3.6.1.3 IPMI Sensor Data Records

The IPMI Sensor Data Records shows, in real time, the measurement values for all sensors present in the system (temperature, power supplies and FANs). The menu is read only, there are no options to configure.

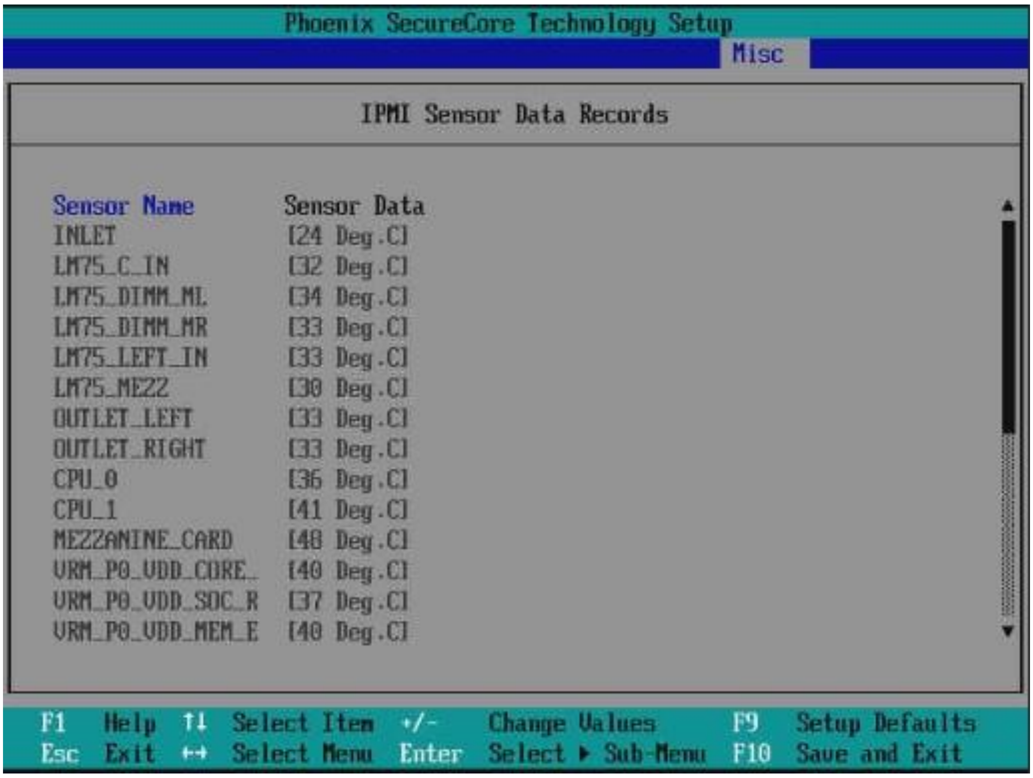
Phoenix SecureCore Technology Setup
Misc
F1 Help ↑↓ Select Item +/- Change Values F9 Setup Defaults
Esc Exit ↔ Select Menu Enter Select ▸ Sub-Menu F10 Save and Exit

IPMI Sensor Data Records

Sensor Name	Sensor Data
INLET	[24 Deg.C]
LM75_C_IN	[32 Deg.C]
LM75_DIMM_ML	[34 Deg.C]
LM75_DIMM_MR	[33 Deg.C]
LM75_LEFT_IN	[33 Deg.C]
LM75_MEZZ	[30 Deg.C]
OUTLET_LEFT	[33 Deg.C]
OUTLET_RIGHT	[33 Deg.C]
CPU_0	[36 Deg.C]
CPU_1	[41 Deg.C]
MEZZANINE_CARD	[48 Deg.C]
VRM_P0_VDD_CORE_	[40 Deg.C]
VRM_P0_VDD_SOC_R	[37 Deg.C]
VRM_P0_VDD_MEM_E	[40 Deg.C]

3.6.1.4 BMC Network Configuration

The BMC Network Configuration menu allows the user to configure the Baseboard Management Controller (BMC) network settings, without the need to log in the BMC itself.

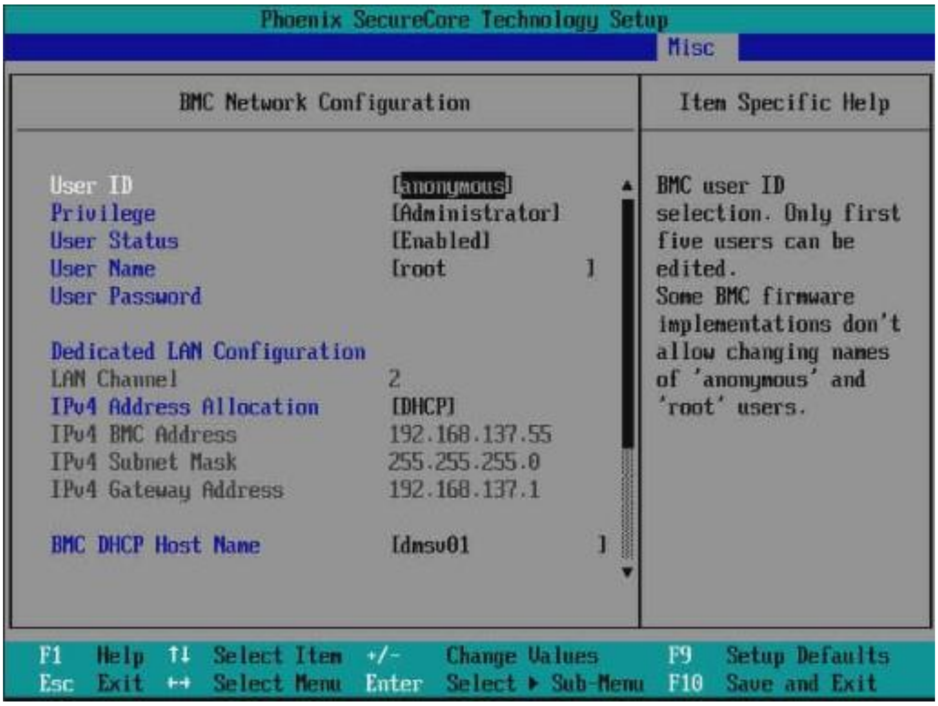

```
                  Phoenix SecureCore Technology Setup
                                                  Misc

        BMC Network Configuration               Item Specific Help

  User ID                    [anonymous]      BMC user ID
  Privilege                  [Administrator]  selection. Only first
  User Status                [Enabled]        five users can be
  User Name                  [root        ]   edited.
  User Password                               Some BMC firmware
                                              implementations don't
  Dedicated LAN Configuration                 allow changing names
  LAN Channel                2                of 'anonymous' and
  IPv4 Address Allocation    [DHCP]           'root' users.
  IPv4 BMC Address           192.168.137.55
  IPv4 Subnet Mask           255.255.255.0
  IPv4 Gateway Address       192.168.137.1

  BMC DHCP Host Name         [dmsv01       ]

 F1   Help  ↑↓  Select Item  +/-    Change Values        F9   Setup Defaults
 Esc  Exit  ↔   Select Menu  Enter  Select ▸ Sub-Menu    F10  Save and Exit
```

There are two LAN Channels available to configure:

- LAN Channel 1: BMC Out Of Band (OOB) Management Ethernet Port (eth1)
- LAN Channel 2: BMC NC-SI (Network Controller Sideband Interface) Management Ethernet Port (eth0)

The LAN Channel 2 is available only if a mezzanine Ethernet adapter with support to NC-SI is assembled in the DM-SV01 equipment.

Both Channels mentioned above have a specific menu to configure the network settings. The configuration items are shown in the table below.

Menu Item	Description	Default Config	Options
User ID	This option allows the creation of a new user to access the BMC. The user configuration can be accessed in the BMC web GUI, through the menu “Access > Local users”.	N/A	[user1] [user2] [user3] [user4] [user5]
Privilege	Configure the privilege level of the user selected in the “User ID”.	N/A	[NoAccess] [ReadOnly] [Operator] [Administrator]
User Status	Enables or disables the user selected in the “User ID”.	N/A	[Enabled] [Disabled]
User Name	Configures the username for the current selected User ID.	N/A	String
User Password	Configures a password for the current selected User ID.	N/A	String
IPv4 Address	Selects the method to configure the IP Address in the	DHCP	[DHCP]

Allocation	Ethernet Port: DHCP or Static.		[Static]
IPv4 BMC Address	Field to configure the IP Address for the BMC Ethernet Port.	N/A	N/A
IPv4 Subnet Mask	Field to configure the Subnet Mask for the BMC Ethernet Port.	N/A	N/A
IPv4 Gateway Address	Field to configure the Gateway IP Address for the BMC Ethernet Port.	N/A	N/A
BMC DHCP Host Name	The hostname reported to the DHCP server when asking for an IP address.	dmsv01	String

The BMC network settings in the UEFI menu reflects the state of the “Configuration > Network Settings” menu available in the BMC web GUI.

For additional information about the Users configuration, please refer to the “Access > Local Users” menu in the BMC web GUI.

3.7 Exit Menu

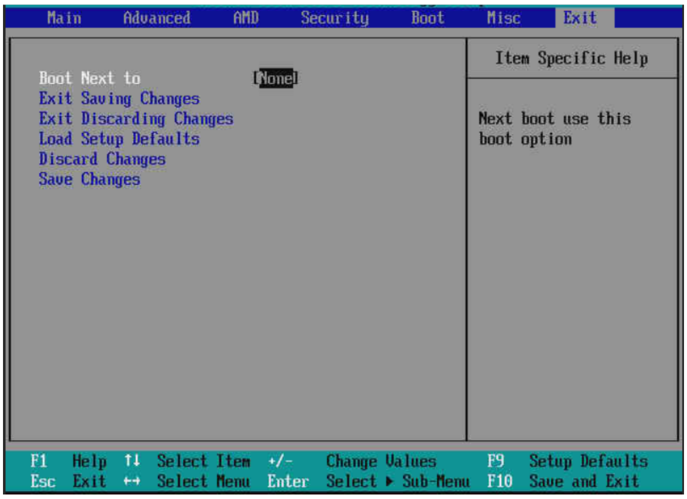

Boot Next to [None]

Next boot use this boot option

Figure 79: Exit Menu

3.7.1 Boot Next to

Menu Item	Description	Default Config	Options
Power Restore Policy	Defines the power on policy when Power is applied to the system.	Stay Off	[Stay Off] [Restored to last state] [Power on]

3.7.2 Exit Saving Changes

Command used to save all changes performed in the UEFI menus, and then exit the UEFI setup. When the setup is exited, the system resets and boots according to the configured boot option, as explained in section 3.5 Boot Menu.

Menu Item	Description	Default Config	Options
Exit Saving Changes	Command used to save all changes performed in the UEFI menus and exit setup.	N/A	Confirm [Yes] [No]

3.7.3 Exit Discarding Changes

Command used to discard all changes performed in the UEFI menus, and then exit the UEFI setup. When the setup is exited, the system resets and boots according to the configured boot option, as explained in section 3.5.1 Boot Priority Order.

Menu Item	Description	Default Config	Options
Exit Discarding Changes	Command used to discard all changes performed in the UEFI menus and exit setup.	N/A	Confirm [Yes] [No]

3.7.4 Load Setup Defaults

Command used to load the default values for all the UEFI menus. The command itself does not save the changes or resets the system. After using the “Load Setup Defaults” option, the user should use Exit Saving Changes to apply the modifications or Exit Discarding Changes to undo.

Menu Item	Description	Default Config	Options
Load Setup Defaults	Loads the default values for all the UEFI menus.	N/A	Confirm [Yes] [No]

3.7.5 Discard Changes

Discards any changes performed in the UEFI menus, restoring the original values of the current boot time, without rebooting the system afterwards. This option does not load the setup defaults, it just restores the values that have been changed in the current access to the UEFI setup.

Menu Item	Description	Default Config	Options

Discard Changes	Discards any changes performed in the UEFI menus, restoring the original values of the current boot time.	N/A	Confirm [Yes] [No]

3.7.6 Save Changes

Saves any changes performed in the UEFI menus in the current boot time, without rebooting the system afterwards.

Menu Item	**Description**	**Default Config**	**Options**
Save Changes	Saves any changes performed in the UEFI menus in the current boot time, without rebooting the system afterwards.	N/A	Confirm [Yes] [No]

4 References

(1) “Socket SP3 Platform NUMA Topology for AMD Family 17h Models 30h–3Fh”, available at *https://developer.amd.com/wp-content/resources/56338_1.00_pub.pdf*

(2) “Workload Tuning Guide for AMD EPYC™ 7002 Series Processor Based Servers”, available at *https://developer.amd.com/wp-content/resources/56745_0.80.pdf*

(3) “AMD I/O Virtualization Technology (IOMMU) Specification”, available at *https://www.amd.com/system/files/TechDocs/48882_IOMMU.pdf*

(4) “AMD SP3 Family 17h Models 30h–3Fh Preferred IO Usage Guide”, available at https://developer.amd.com/wp-content/resources/56570_0.93.pdf

(5) “Processor Programming Reference (PPR) for Family 19h Model 01h, Revision B1 Processors”, available at *https://www.amd.com/system/files/TechDocs/55898_B1_pub_0.50.zip*

(6) “AMD MEMORY ENCRYPTION”, available at *https://developer.amd.com/wordpress/media/2013/12/AMD_Memory_Encryption_Whitepaper_v7-Public.pdf*

(7) “DM-SV01 Server - Product Manual”.

(8) “AMD64 Architecture Programmer’s Manual Volume 2: System Programming”, available at *https://www.amd.com/system/files/TechDocs/24593.pdf*

(9) Article: “UEFI boot: how does that actually work, then?” available at *this link*.

(10) “DM-SV01 Server - BMC User Manual”.

This document comprises 139 pages.

Revision History:

Date	**Description**
2022/06/03	First official release (1.0)
2022/08/22	Release 1.1 - Updated menus after BIOS release upgrade to DM_SV01_R200: - Added chapters 3.3.2.3 up to 3.3.2.9 (SRIS) > SRIS mode debug and suboptions > SRIS Autodetect - Added options below in chapter 3.3.3.3.1.2 (DRAM self refresh settings): > SubUrgRefLowerBound > UrgRefLimit > DRAM Maximum Activate Count > DRAM Refresh Rate > Self-Refresh Exit Staggering - Added chapters 3.3.3.3.4.5 up to 3.3.3.3.4.7 (MBIST additional settings): AMD > AMD CBS > UMC Common Options > Memory MBIST > Memory Healing BIST > Mem BIST Test Select > Mem BIST Post Package Repair Type - Added chapters 3.3.3.4.15 up to 3.3.3.4.20 (NBIO miscellaneous additional settings): AMD > AMD CBS > NBIO Common Options > CAC Weight Adjustment > EDC Control Throttle > SRIS > Compliance Loopback > Multi Upstream Auto Speed Change > Multi Auto Speed Change On Last Rate - Added chapters 3.3.3.4.21.24 up to 3.3.3.4.21.26 (EDC tracking settings): AMD > AMD CBS > NBIO Common Options > SMU Common Options > EDC Current Tracking > EDC Tracking Current Threshold > EDC Tracking Report Interval - Added chapter 3.3.3.5.5.1 (boot timer enable/disable): AMD > AMD CBS > FCH Common Options > Miscellaneous Options > Boot Timer Enable - NTB in chapter 3.3.3.6: added warning for non-supported feature.
2022/12/12	Release 1.1 - Updated menus after BIOS release upgrade to DM_SV01_R203:

	- Updated section 3.1.4.7 Console Redirection: > Changed console redirection default setting to “disabled”.

Annex A - BIOS/UEFI release notes

A.1 Release Datacom_SV01_R203 X64

BIOS/UEFI release: Datacom_SV01_R203 X64	
Changes when comparing to the previous release (R202)	**Related section**
Setup - Console redirection disabled by default, making the virtual media operate faster and improving the user navigation experience in the BIOS menus.	3.1.4.7 Console Redirection
Setup - “Continue C.R. after POST” disabled by default. The reason is the same as the first item.	3.1.4.7 Console Redirection
Bug - Correct buffer length for supporting up to 20-byte password.	3.4.1 Set Supervisor Password
Bug - Avoid delete of empty boot option.	3.5 Boot Menu
Improvement - Report DIMMs of CPU P1 as not present when CPU P1 is not present.	Not applicable.
Bug - Fix wrong status of power restore policy in BIOS. Now, the “Stay off”, “Power on” and “Restore to last state” options are working as expected.	3.6.1.2 Power restore policy
Bug - Enable the UART FIFO before enabling UART ports, solving the issue of not being possible to navigate in BIOS menus by using the serial over LAN (SOL) screen from BMC.	Not applicable.
Bug - Fix FreeBSD screen output when console redirection is disabled.	3.1.4.7 Console Redirection