# Intel® Platform Innovation Framework for EFI Driver Execution Environment Core Interface Specification (DXE CIS)

## A Foundation Specification

---

Version 0.91.01

December 8, 2004

Information in this document is provided in connection with Intel. products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

This document as well as the software described in it is furnished under license and may only be used or copied in accordance with the terms of the license. The information in this manual is furnished for informational use only, is subject to change without notice, and should not be construed as a commitment by Intel Corporation to update or revise the information or document. Intel Corporation assumes no responsibility or liability for any errors or inaccuracies that may appear in this document or any software that may be provided in association with this document.

This document provides website addresses for certain third party websites. The referenced sites are not under the control of Intel and Intel is not responsible for the content of any referenced site or any link contained in a referenced site. Intel does not endorse companies or products for sites which it references. If you decide to access any of the third party sites referenced in this document, you do this entirely at your own risk.

*Other names and brands may be claimed as the property of others.

Intel, the Intel logo, Itanium and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Copyright ©2002-2005 Intel Corporation**.** All Rights Reserved.

# Revision History

| Revision | Revision History | Date |
|---|---|---|
| 0.9 | First public release. | 9/16/03 |
| | Added "A Foundation Specification" line to the title page. No other changes, so the revision number and date were not changed. | 6/30/04 |
| 0.91 | Correctly labeled second parameter as “Out” in EFI_TIMER_ARCH_PROTOCOL.GetTimerPeriod().<br>Removed event registration fom the EFI Runtime Architectural Protocol. | 12/08/04 |
| 0.91.01 | Remove Intel Confidential | 06/26/08 |

intel

# Contents


**1 Introduction ..... 11**
- Overview ..... 11
- Organization of the DXE CIS ..... 11
- Target Audience ..... 12
- Conventions Used in This Document ..... 12
  - Data Structure Descriptions ..... 12
  - Protocol Descriptions ..... 13
  - Procedure Descriptions ..... 13
  - Instruction Descriptions ..... 14
  - Pseudo-Code Conventions ..... 14
  - Typographic Conventions ..... 14

**2 Overview ..... 17**
- Driver Execution Environment (DXE) Phase ..... 17
- EFI System Table ..... 18
  - Overview ..... 18
  - EFI Boot Services Table ..... 19
  - EFI Runtime Services Table ..... 20
  - DXE Services Table ..... 20
- DXE Foundation ..... 21
- DXE Dispatcher ..... 21
- DXE Drivers ..... 21
- DXE Architectural Protocols ..... 22

**3 Boot Manager ..... 25**
- Boot Manager ..... 25

**4 EFI System Table ..... 27**
- Introduction ..... 27
- EFI Image Entry Point ..... 27
  - EFI_IMAGE_ENTRY_POINT ..... 28
- EFI Table Header ..... 30
- EFI System Table ..... 31
- EFI Boot Services Table ..... 34
  - EFI_BOOT_SERVICES ..... 34
- EFI Runtime Services Table ..... 40
  - EFI_RUNTIME_SERVICES ..... 40
- EFI Configuration Table ..... 43
- DXE Services Table ..... 45
  - DXE_SERVICES ..... 45
- EFI Image Entry Point Examples ..... 48
  - EFI Application Example ..... 48
  - Non-EFI Driver Model Example (Resident in Memory) ..... 51
  - Non-EFI Driver Model (Nonresident in Memory) ..... 52

EFI Driver Model Example ..... 53
EFI Driver Model Example (Unloadable) ..... 54
EFI Driver Model Example (Multiple Instances) ..... 56

**5 Services - Boot Services ..... 59**

EFI 1.10 Boot Services ..... 59
Extensions to EFI 1.10 Boot Services ..... 61
CreateEvent() ..... 61
LoadImage() ..... 65

**6 Services - Runtime Services ..... 69**

EFI 1.10 Runtime Services ..... 69
Additional Runtime Services ..... 70
Status Code Services ..... 70
ReportStatusCode() ..... 71

**7 Services - DXE Services ..... 75**

Introduction ..... 75
Global Coherency Domain Services ..... 75
Overview ..... 75
Global Coherency Domain (GCD) Services Overview ..... 75
GCD Memory Resources ..... 75
GCD I/O Resources ..... 77
Global Coherency Domain Services ..... 79
AddMemorySpace() ..... 80
AllocateMemorySpace() ..... 83
FreeMemorySpace() ..... 86
RemoveMemorySpace() ..... 88
GetMemorySpaceDescriptor() ..... 90
SetMemorySpaceAttributes() ..... 93
GetMemorySpaceMap() ..... 95
AddIoSpace() ..... 97
AllocateIoSpace() ..... 99
FreeIoSpace() ..... 102
RemoveIoSpace() ..... 104
GetIoSpaceDescriptor() ..... 106
GetIoSpaceMap() ..... 108
Dispatcher Services ..... 110
Dispatcher Services ..... 110
Dispatch() ..... 111
Schedule() ..... 112
Trust() ..... 113
ProcessFirmwareVolume() ..... 114

**8 Protocols - Device Path Protocol ..... 117**

Introduction ..... 117
Firmware Volume File Path Media Device Path ..... 118

**9 DXE Foundation ..... 119**
Introduction ..... 119
Hand-Off Block (HOB) List ..... 120
DXE Foundation Data Structures ..... 121
Required DXE Foundation Components ..... 122
Handing Control to DXE Dispatcher ..... 124
DXE Foundation Entry Point ..... 125
DXE Foundation Entry Point ..... 125
DXE_ENTRY_POINT ..... 126
Dependencies ..... 127
EFI Boot Services Table ..... 127
EFI Boot Services Dependencies ..... 127
SetTimer() ..... 128
RaiseTPL() ..... 128
RestoreTPL() ..... 129
SetWatchdogTimer() ..... 129
Stall() ..... 129
GetNextMonotonicCount() ..... 129
CalculateCrc32() ..... 129
EFI Runtime Services Table ..... 130
EFI Runtime Services Dependencies ..... 130
GetVariable() ..... 130
GetNextVariableName() ..... 130
SetVariable() ..... 130
GetTime() ..... 131
SetTime() ..... 131
GetWakeupTime() ..... 131
SetWakeupTime() ..... 131
SetVirtualAddressMap() ..... 131
ConvertPointer() ..... 132
ResetSystem() ..... 132
GetNextHighMonotonicCount() ..... 132
ReportStatusCode() ..... 132
DXE Services Table ..... 132
DXE Services Dependencies ..... 132
GetMemorySpaceDescriptor() ..... 133
SetMemorySpaceAttributes() ..... 133
GetMemorySpaceMap() ..... 133
HOB Translations ..... 134
HOB Translations Overview ..... 134
PHIT HOB ..... 134
CPU HOB ..... 134
Resource Descriptor HOBs ..... 135
Firmware Volume HOBs ..... 135
Memory Allocation HOBs ..... 136
GUID Extension HOBs ..... 136

**10 DXE Dispatcher ..... 137**
Introduction ..... 137
Requirements ..... 138
The a priori File ..... 139
EFI_APRIORI_GUID ..... 140
Dependency Expressions ..... 140
Dependency Expressions Overview ..... 140
Dependency Expression Instruction Set ..... 140
BEFORE ..... 142
AFTER ..... 143
PUSH ..... 144
AND ..... 145
OR ..... 146
NOT ..... 147
TRUE ..... 148
FALSE ..... 149
END ..... 150
SOR ..... 151
Dependency Expression with No Dependencies ..... 152
Empty Dependency Expressions ..... 152
Dependency Expression Reverse Polish Notation (RPN) ..... 154
DXE Dispatcher State Machine ..... 155
DXE Dispatcher State Machine ..... 155
Example Orderings ..... 157
Security Considerations ..... 159

**11 DXE Drivers ..... 161**
Introduction ..... 161
Classes of DXE Drivers ..... 161
Basic ..... 161
Early DXE Drivers ..... 161
DXE Drivers That Follow the EFI Driver Model ..... 162
Additional ..... 162
Additional Classifications ..... 162

**12 DXE Architectural Protocols ..... 163**
Introduction ..... 163
Boot Device Selection (BDS) Architectural Protocol ..... 166
EFI_BDS_ARCH_PROTOCOL ..... 166
EFI_BDS_ARCH_PROTOCOL.Entry() ..... 167
CPU Architectural Protocol ..... 168
EFI_CPU_ARCH_PROTOCOL ..... 168
EFI_CPU_ARCH_PROTOCOL.FlushDataCache() ..... 171
EFI_CPU_ARCH_PROTOCOL.EnableInterrupt() ..... 173
EFI_CPU_ARCH_PROTOCOL.DisableInterrupt() ..... 174
EFI_CPU_ARCH_PROTOCOL.GetInterruptState() ..... 175
EFI_CPU_ARCH_PROTOCOL.Init() ..... 176

EFI_CPU_ARCH_PROTOCOL. RegisterInterruptHandler() ....177
EFI_CPU_ARCH_PROTOCOL.GetTimerValue() ....179
EFI_CPU_ARCH_PROTOCOL.SetMemoryAttributes() ....181
Metronome Architectural Protocol ....183
EFI_METRONOME_ARCH_PROTOCOL....183
EFI_METRONOME_ARCH_PROTOCOL.WaitForTick() ....184
Monotonic Counter Architectural Protocol ....185
EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL....185
Real Time Clock Architectural Protocol ....186
EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL ....186
Reset Architectural Protocol....187
EFI_RESET_ARCH_PROTOCOL ....187
Runtime Architectural Protocol....188
Runtime Architectural Protocol ....188
EFI_RUNTIME_ARCH_PROTOCOL....188
EFI_RUNTIME_ARCH_PROTOCOL.RegisterImage()...**Error! Bookmark not defined.**
EFI_RUNTIME_ARCH_PROTOCOL.RegisterEvent() ...**Error! Bookmark not defined.**
Security Architectural Protocol ....194
EFI_SECURITY_ARCH_PROTOCOL ....194
EFI_SECURITY_ARCH_PROTOCOL. FileAuthenticationState() ....196
Status Code Architectural Protocol....198
EFI_STATUS_CODE_ARCH_PROTOCOL....198
Timer Architectural Protocol ....199
EFI_TIMER_ARCH_PROTOCOL....199
EFI_TIMER_ARCH_PROTOCOL.RegisterHandler() ....201
EFI_TIMER_ARCH_PROTOCOL.SetTimerPeriod() ....203
EFI_TIMER_ARCH_PROTOCOL.GetTimerPeriod()....204
EFI_TIMER_ARCH_PROTOCOL. GenerateSoftInterrupt() ....205
Variable Architectural Protocol ....206
EFI_VARIABLE_ARCH_PROTOCOL....206
Variable Write Architectural Protocol....207
EFI_VARIABLE_WRITE_ARCH_PROTOCOL ....207
Watchdog Timer Architectural Protocol ....208
Watchdog Timer Architectural Protocol....208
EFI_WATCHDOG_TIMER_ARCH_PROTOCOL ....208
EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. RegisterHandler() ....210
EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. SetTimerPeriod() ....212
EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. GetTimerPeriod()....213

**13 Returned Status Codes.... 215**

Returned Status Codes....215
EFI_STATUS Codes Ranges....215
EFI_STATUS Success Codes (High Bit Clear)....215
EFI_STATUS Error Codes (High Bit Set) ....216
EFI_STATUS Warning Codes (High Bit Clear)....217

**14 Dependency Expression Grammar ..... 219**
Dependency Expression Grammar ..... 219
Example Dependency Expression BNF Grammar ..... 219
Sample Dependency Expressions ..... 219

## Figures

Figure 2.1. Framework Firmware Phases ..... 18
Figure 7.1. GCD Memory State Transitions ..... 77
Figure 7.2. GCD I/O State Transitions ..... 78
Figure 9.1. HOB List ..... 120
Figure 9.2. EFI System Table and Related Components ..... 121
Figure 9.3. DXE Foundation Components ..... 122
Figure 10.1. DXE Driver States ..... 155
Figure 10.2. Sample Firmware Volume ..... 157
Figure 12.1. DXE Architectural Protocols ..... 164

## Tables

Table 1.1. Organization of the DXE CIS ..... 11
Table 5.1. Boot Services in the EFI 1.10 Specification ..... 59
Table 6.1. EFI 1.10 Runtime Services ..... 69
Table 6.2. Status Code Runtime Services ..... 70
Table 7.1. Global Coherency Domain Services ..... 79
Table 7.2. Dispatcher Services ..... 110
Table 8.1. Firmware Volume File Path Media Device Path ..... 118
Table 9.1. Boot Service Dependencies ..... 127
Table 9.2. Runtime Service Dependencies ..... 130
Table 9.3. DXE Service Dependencies ..... 133
Table 9.4. Resource Descriptor HOB to GCD Type Mapping ..... 135
Table 10.1. Dependency Expression Opcode Summary ..... 141
Table 10.2. DXE Dispatcher Orderings ..... 158

# 1 Introduction

## Overview

This specification defines the core code and services that are required for an implementation of the Driver Execution Environment (DXE) phase of the Intel® Platform Innovation Framework for EFI (hereafter referred to as the "Framework"). This DXE Core Interface Specification (CIS) does the following:

- Describes the basic components of the DXE phase
- Provides code definitions for services and functions that are architecturally required by the *Intel® Platform Innovation Framework for EFI Architecture Specification*
- Presents a set of backward-compatible extensions to the *EFI 1.10 Specification*
- Describes the machine preparation that is required for subsequent phases of firmware execution

See Organization of the DXE CIS for more information.

## Organization of the DXE CIS

This DXE Foundation Interface Specification (CIS) is organized as listed below. Because the DXE Foundation is just one component of a Framework-based firmware solution, there are a number of additional specifications that are referred to throughout this document:

- For references to other Framework specifications, click on the hyperlink in the page or navigate through the table of contents (TOC) in the left navigation pane to view the referenced specification.
- For references to non-Framework specifications, see References in the Interoperability and Component Specifications help system.

**Table 1.1. Organization of the DXE CIS**

| Book | Description |
|---|---|
| Overview | Describes the major components of DXE, including the boot manager, firmware core, protocols, and requirements. |
| Boot Manager | Describes the boot manager, which is used to load EFI drivers, EFI applications, and EFI OS loaders. |
| EFI System Table | Describes the EFI System Table that is passed to every EFI driver and EFI application. |
| Services - Boot Services | Contains the definitions of the fundamental services that are present in an EFI-compliant system before an OS is booted. |
| Services - Runtime Services | Contains definitions for the fundamental services that are present in an EFI-compliant system before and after an OS is booted. |
| Services - DXE Services | Contains definitions for the fundamental services that are present in a DXE-compliant system before an OS is booted. |

| Book | Description |
|---|---|
| Protocols - Device Path Protocol | Defines the device path extensions required by the DXE Foundation. |
| DXE Foundation | Describes the DXE Foundation that consumes HOBs, Firmware Volumes, and DXE Architectural Protocols to produce an EFI System Table, EFI Boot Services, EFI Runtime Services, and the DXE Services. |
| DXE Dispatcher | Describes the DXE Dispatcher that is responsible for loading and executing DXE drivers from Firmware Volumes. |
| DXE Drivers | Describes the different classes of DXE drivers that may be stored in Firmware Volumes. |
| DXE Architectural Protocols | Describes the Architectural Protocols that are produced by DXE drivers. They are also consumed by the DXE Foundation to produce the EFI Boot Services, EFI Runtime Services, and DXE Services. |
| Returned Status Codes | Lists success, error, and warning codes returned by DXE and EFI interfaces. |
| Dependency Expression Grammar | Describes the BNF grammar for a tool that can convert a text file containing a dependency expression into a dependency section of a DXE driver stored in a Firmware Volume. |

## Target Audience

This document is intended for the following readers:

- IHVs and OEMs who will be implementing DXE drivers that are stored in firmware volumes.
- BIOS developers, either those who create general-purpose BIOS and other firmware products or those who modify these products for use in Intel® architecture–based products.

## Conventions Used in This Document

This document uses the typographic and illustrative conventions described below.

### Data Structure Descriptions

Intel® processors based on 32-bit Intel® architecture (IA-32) are “little endian” machines. This distinction means that the low-order byte of a multibyte data item in memory is at the lowest address, while the high-order byte is at the highest address. Processors of the Intel® Itanium® processor family may be configured for both “little endian” and “big endian” operation. All implementations designed to conform to this specification will use “little endian” operation.

In some memory layout descriptions, certain fields are marked *reserved*. Software must initialize such fields to zero and ignore them when read. On an update operation, software must preserve any reserved field.

The data structures described in this document generally have the following format:

**STRUCTURE NAME:** The formal name of the data structure.

**Summary:** A brief description of the data structure.

**Prototype:** A “C-style” type declaration for the data structure.

**Parameters:** A brief description of each field in the data structure prototype.

**Description:** A description of the functionality provided by the data structure, including any limitations and caveats of which the caller should be aware.

**Related Definitions:** The type declarations and constants that are used only by this data structure.

### Protocol Descriptions

The protocols described in this document generally have the following format:

**Protocol Name:** The formal name of the protocol interface.

**Summary:** A brief description of the protocol interface.

**GUID:** The 128-bit Globally Unique Identifier (GUID) for the protocol interface.

**Protocol Interface Structure:** A “C-style” data structure definition containing the procedures and data fields produced by this protocol interface.

**Parameters:** A brief description of each field in the protocol interface structure.

**Description:** A description of the functionality provided by the interface, including any limitations and caveats of which the caller should be aware.

**Related Definitions:** The type declarations and constants that are used in the protocol interface structure or any of its procedures.

### Procedure Descriptions

The procedures described in this document generally have the following format:

**ProcedureName():** The formal name of the procedure.

**Summary:** A brief description of the procedure.

**Prototype:** A “C-style” procedure header defining the calling sequence.

**Parameters:** A brief description of each field in the procedure prototype.

**Description:** A description of the functionality provided by the interface, including any limitations and caveats of which the caller should be aware.

**Related Definitions:** The type declarations and constants that are used only by this procedure.

**Status Codes Returned:** A description of any codes returned by the interface. The procedure is required to implement any status codes listed in this table. Additional error codes may be returned, but they will not be tested by standard compliance tests, and any software that uses the procedure cannot depend on any of the extended error codes that an implementation may provide.

## Instruction Descriptions

A dependency expression instruction description generally has the following format:

**InstructionName** The formal name of the instruction.

**SYNTAX:** A brief description of the instruction.

**DESCRIPTION:** A description of the functionality provided by the instruction accompanied by a table that details the instruction encoding.

**OPERATION:** Details the operations performed on operands.

**BEHAVIORS AND RESTRICTIONS:** An item-by-item description of the behavior of each operand involved in the instruction and any restrictions that apply to the operands or the instruction.

## Pseudo-Code Conventions

Pseudo code is presented to describe algorithms in a more concise form. None of the algorithms in this document are intended to be compiled directly. The code is presented at a level corresponding to the surrounding text.

In describing variables, a *list* is an unordered collection of homogeneous objects. A *queue* is an ordered list of homogeneous objects. Unless otherwise noted, the ordering is assumed to be First In First Out (FIFO).

Pseudo code is presented in a C-like format, using C conventions where appropriate. The coding style, particularly the indentation style, is used for readability and does not necessarily comply with an implementation of the *Extensible Firmware Interface Specification*.

## Typographic Conventions

This document uses the typographic and illustrative conventions described below:

| | |
|---|---|
| Plain text | The normal text typeface is used for the vast majority of the descriptive text in a specification. |
| Plain text (blue) | In the online help version of this specification, any plain text that is underlined and in blue indicates an active link to the cross-reference. Click on the word to follow the hyperlink. Note that these links are *not* active in the PDF of the specification. |

| | |
|---|---|
| **Bold** | In text, a **Bold** typeface identifies a processor register name. In other instances, a **Bold** typeface can be used as a running head within a paragraph. |
| *Italic* | In text, an *Italic* typeface can be used as emphasis to introduce a new term or to indicate a manual or specification name. |
| **`BOLD Monospace`** | Computer code, example code segments, and all prototype code segments use a **`BOLD Monospace`** typeface with a dark red color. These code listings normally appear in one or more separate paragraphs, though words or segments can also be embedded in a normal text paragraph. |
| **`Bold Monospace`** | In the online help version of this specification, words in a **`Bold Monospace`** typeface that is underlined and in blue indicate an active hyperlink to the code definition for that function or type definition. Click on the word to follow the hyperlink. Note that these links are *not* active in the PDF of the specification. Also, these inactive links in the PDF may instead have a **`Bold Monospace`** appearance that is underlined but in dark red. Again, these links are not active in the PDF of the specification. |
| *`Italic Monospace`* | In code or in text, words in *`Italic Monospace`* indicate placeholder names for variable information that must be supplied (i.e., arguments). |
| `Plain Monospace` | In code, words in a `Plain Monospace` typeface that is a dark red color but is not bold or italicized indicate pseudo code or example code. These code segments typically occur in one or more separate paragraphs. |

See the master Framework glossary in the Framework Interoperability and Component Specifications help system for definitions of terms and abbreviations that are used in this document or that might be useful in understanding the descriptions presented in this document.

See the master Framework references in the Interoperability and Component Specifications help system for a complete list of the additional documents and specifications that are required or suggested for interpreting the information presented in this document.

The Framework Interoperability and Component Specifications help system is available at the following URL:

http://www.intel.com/technology/framework/spec.htm

# 2
# Overview

## Driver Execution Environment (DXE) Phase

The Driver Execution Environment (DXE) phase is where most of the system initialization is performed. Pre-EFI Initialization (PEI), the phase prior to DXE, is responsible for initializing permanent memory in the platform so that the DXE phase can be loaded and executed. The state of the system at the end of the PEI phase is passed to the DXE phase through a list of position-independent data structures called *Hand-Off Blocks* (HOBs). HOBs are described in detail in the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification*.

There are several components in the DXE phase:

- DXE Foundation
- DXE Dispatcher
- A set of DXE drivers

The DXE Foundation produces a set of Boot Services, Runtime Services, and DXE Services. The DXE Dispatcher is responsible for discovering and executing DXE drivers in the correct order. The DXE drivers are responsible for initializing the processor, chipset, and platform components as well as providing software abstractions for system services, console devices, and boot devices. These components work together to initialize the platform and provide the services required to boot an operating system. The DXE phase and Boot Device Selection (BDS) phases work together to establish consoles and attempt the booting of operating systems. The DXE phase is terminated when an operating system is successfully booted. The DXE Foundation is composed of boot services code, so no code from the DXE Foundation itself is allowed to persist into the OS runtime environment. Only the runtime data structures allocated by the DXE Foundation and services and data structured produced by runtime DXE drivers are allowed to persist into the OS runtime environment.

The figure below shows the phases that a platform with Framework firmware will execute.

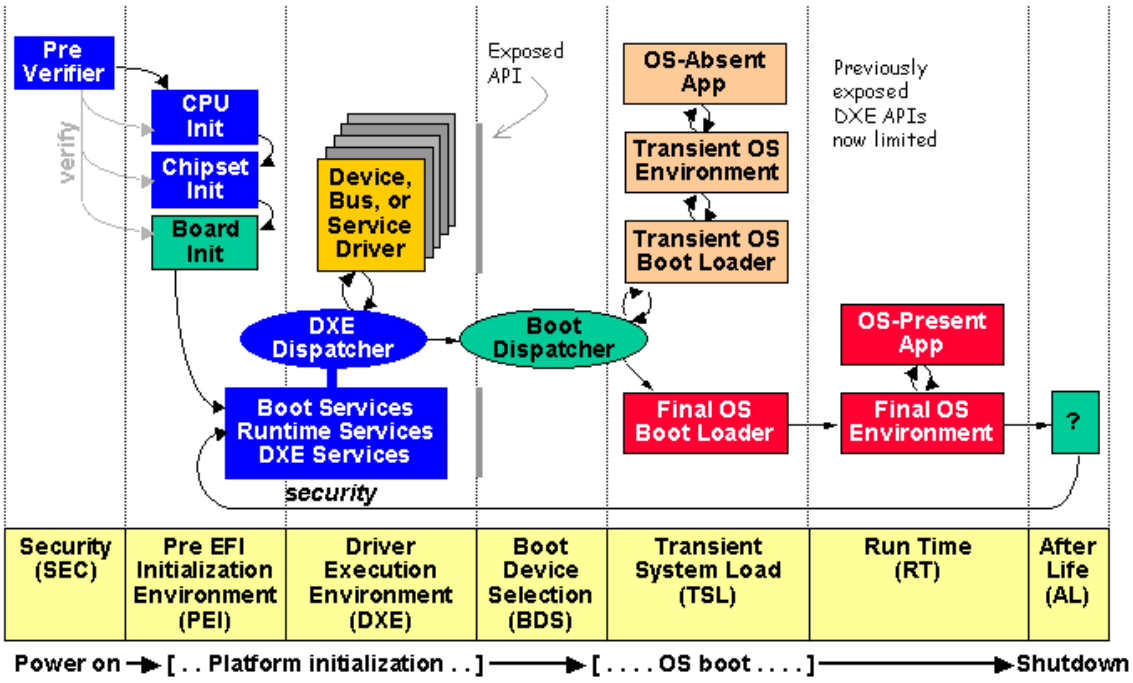


**Figure 2.1. Framework Firmware Phases**

In a Framework firmware implementation, the phase executed prior to DXE is PEI. This specification covers the transition from the PEI to the DXE phase, the DXE phase, and the DXE phase's interaction with the BDS phase. The DXE phase does not require a PEI phase to be executed. The only requirement for the DXE phase to execute is the presence of a valid HOB list. There are many different implementations that can produce a valid HOB list for the DXE phase to execute. The PEI phase in a Framework firmware implementation is just one of many possible implementations.

# EFI System Table

## Overview

The EFI System Table is passed to every executable component in the DXE phase. The EFI System Table contains a pointer to the following:

- EFI Boot Services Table
- EFI Runtime Services Table

It also contains pointers to the console devices and their associated I/O protocols. In addition, the EFI System Table contains a pointer to the EFI Configuration Table, and this table contains a list of GUID/pointer pairs. The EFI Configuration Table may include tables such as the DXE Services Table, HOB list, ACPI table, SMBIOS table, and SAL System table.

The EFI Boot Services Table contains services to access the contents of the handle database. The handle database is where protocol interfaces produced by drivers are registered. Other drivers can use the EFI Boot Services to look up these services produced by other drivers.

All of the services available in the DXE phase may be accessed through a pointer to the EFI System Table.

## EFI Boot Services Table

The following is a brief summary of the services that are available through the EFI Boot Services Table. These services are described in detail in the *EFI 1.10 Specification.* This DXE CIS makes a few minor, backward-compatible extensions to these services.

| | |
|---|---|
| ***Task Priority Services:*** | Provides services to increase or decrease the current task priority level. This can be used to implement simple locks and to disable the timer interrupt for short periods of time. These services depend on the CPU Architectural Protocol. |
| ***Memory Services:*** | Provides services to allocate and free pages in 4 KB increments and allocate and free pool on byte boundaries. It also provides a service to retrieve a map of all the current physical memory usage in the platform. |
| ***Event and Timer Services:*** | Provides services to create events, signal events, check the status of events, wait for events, and close events. One class of events is timer events, and that class supports periodic timers with variable frequencies and one-shot timers with variable durations. These services depend on the CPU Architectural Protocol, the Timer Architectural Protocol, the Metronome Architectural Protocol, and the Watchdog Timer Architectural Protocol. |
| ***Protocol Handler Services:*** | Provides services to add and remove handles from the handle database. It also provides services to add and remove protocols from the handles in the handle database. Additional services are available that allow any component to lookup handles in the handle database, and open and close protocols in the handle database. |
| ***Image Services:*** | Provides services to load, start, exit, and unload images using the PE/COFF image format. These services use the services of the Security Architectural Protocol if it is present. |
| ***Driver Support Services:*** | Provides services to connect and disconnect drivers to devices in the platform. These services are used by the BDS phase to either connect all drivers to all devices, or to connect only the minimum number of drivers to devices required to establish the consoles and boot an operating system. The minimal connect strategy is one possible mechanism to reduce boot time. |

## EFI Runtime Services Table

The following is a brief summary of the services that are available through the EFI Runtime Services Table. These services are described in detail in the *EFI 1.10 Specification*. One additional runtime service, Status Code Services, is described in this specification.

| | |
|---|---|
| ***Variable Services :*** | Provides services to look up, add, and remove environment variables from nonvolatile storage. These services depend on the Variable Architectural Protocol and the Variable Write Architectural Protocol. |
| ***Real Time Clock Services:*** | Provides services to get and set the current time and date. It also provides services to get and set the time and date of an optional wake-up timer. These services depend on the Real Time Clock Architectural Protocol. |
| ***Reset Services:*** | Provides services to shut down or reset the platform. These services depend on the Reset Architectural Protocol. |
| ***Status Code Services:*** | Provides services to send status codes to a system log or a status code reporting device. These services depend on the Status Code Architectural Protocol. |
| ***Virtual Memory Services:*** | Provides services that allow the runtime DXE components to be converted from a physical memory map to a virtual memory map. These services can only be called once in physical mode. Once the physical to virtual conversion has been performed, these services cannot be called again. These services depend on the Runtime Architectural Protocol. |

## DXE Services Table

The following is a brief summary of the services that are available through the DXE Services Table. These are new services that are available in boot service time and are required only by the DXE Foundation and DXE drivers.

| | |
|---|---|
| ***Global Coherency Domain Services:*** | Provides services to manage I/O resources, memory-mapped I/O resources, and system memory resources in the platform. These services are used to dynamically add and remove these resources from the processor's global coherency domain. |
| ***Dispatcher Services:*** | Provides services to manage DXE drivers that are being dispatched by the DXE Dispatcher. |

## DXE Foundation

The DXE Foundation is a boot service image that is responsible for producing the following:

- EFI Boot Services
- EFI Runtime Services
- DXE Services

The DXE Foundation consumes a HOB list and the services of the DXE Architectural Protocols to produce the full complement of EFI Boot Services, EFI Runtime Services, and DXE Services. The HOB list is described in detail in the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification.*

The DXE Foundation is an implementation of EFI. The DXE Foundation defined in this specification is backward compatible with the *EFI 1.10 Specification*. As a result, both the DXE Foundation and DXE drivers share many of the attributes of EFI images. Because this specification makes extensions to the standard EFI interfaces, DXE images will not be functional on EFI systems that are not compliant with this DXE CIS. However, EFI images must be functional on all EFI-compliant systems including those that are compliant with the DXE CIS.

## DXE Dispatcher

The DXE Dispatcher is one component of the DXE Foundation. This component is required to discover DXE drivers stored in firmware volumes and execute them in the proper order. The proper order is determine by a combination of an *a priori* file that is optionally stored in the firmware volume and the dependency expressions that are part of the DXE drivers. The dependency expression tells the DXE Dispatcher the set of services that a particular DXE driver requires to be present for the DXE driver to execute. The DXE Dispatcher does not allow a DXE driver to execute until all of the DXE driver's dependencies have been satisfied. After all of the DXE drivers have been loaded and executed by the DXE Dispatcher, control is handed to the BDS Architectural Protocol that is responsible for implementing a boot policy that is compliant with the EFI Boot Manager described in the *EFI 1.10 Specification*.

## DXE Drivers

The DXE drivers are required to initialize the processor, chipset, and platform. They are also required to produce the DXE Architectural Protocols and any additional protocol services required to produce I/O abstractions for consoles and boot devices.

# DXE Architectural Protocols

The following is a brief summary of the DXE Architectural Protocols. The DXE Foundation is abstracted from the platform through the DXE Architectural Protocols. The DXE Architectural Protocols manifest the platform-specific components of the DXE Foundation. DXE drivers that are loaded and executed by the DXE Dispatcher component of the DXE Foundation must produce these protocols.

| | |
|---|---|
| ***Security Architectural Protocol:*** | Allows the DXE Foundation to authenticate files stored in firmware volumes before they are used. |
| ***CPU Architectural Protocol:*** | Provides services to manage caches, manage interrupts, retrieve the processor's frequency, and query any processor-based timers. |
| ***Metronome Architectural Protocol:*** | Provides the services required to perform very short calibrated stalls. |
| ***Timer Architectural Protocol:*** | Provides the services required to install and enable the heartbeat timer interrupt required by the timer services in the DXE Foundation. |
| ***BDS Architectural Protocol:*** | Provides an entry point that the DXE Foundation calls once after all of the DXE drivers have been dispatched from all of the firmware volumes. This entry point is the transition from the DXE phase to the Boot Device Selection (BDS) phase, and it is responsible for establishing consoles and enabling the boot devices required to boot an OS. |
| ***Watchdog Timer Architectural Protocol:*** | Provides the services required to enable and disable a watchdog timer in the platform. |
| ***Runtime Architectural Protocol:*** | Provides the services required to convert all runtime services and runtime drivers from physical mappings to virtual mappings. |
| ***Variable Architectural Protocol:*** | Provides the services to retrieve environment variables and set volatile environment variables. |
| ***Variable Write Architectural Protocol:*** | Provides the services to set nonvolatile environment variables. |
| ***Monotonic Counter Architectural Protocol:*** | Provides the services required by the DXE Foundation to manage a 64-bit monotonic counter. |
| ***Reset Architectural Protocol:*** | Provides the services required to reset or shutdown the platform. |

| | |
|---|---|
| ***Status Code Architectural Protocol:*** | Provides the services to send status codes from the DXE Foundation or DXE drivers to a log or device. |
| ***Real Time Clock Architectural Protocol:*** | Provides the services to retrieve and set the current time and date as well as the time and date of an optional wake-up timer. |

# 3
# Boot Manager

## Boot Manager

The Boot Manager in DXE executes after all the DXE drivers whose dependencies have been satisfied have been dispatched by the DXE Dispatcher. At that time, control is handed to the Boot Device Selection (BDS) phase of execution. The BDS phase is responsible for implementing the platform boot policy. System firmware that is compliant with this specification must implement the boot policy specified in the Boot Manager chapter of the *EFI 1.10 Specification*. This boot policy provides flexibility that allows system vendors to customize the user experience during this phase of execution.

The BDS phase is implemented as part of the BDS Architectural Protocol. The DXE Foundation will hand control to the BDS Architectural Protocol after all of the DXE drivers whose dependencies have been satisfied have been loaded and executed by the DXE Dispatcher. The BDS phase is responsible for the following:

- Initializing console devices
- Loading device drivers
- Attempting to load and execute boot selections

If the BDS phase cannot make forward progress, it will reinvoke the DXE Dispatcher to see if the dependencies of any additional DXE drivers have been satisfied since the last time the DXE Dispatcher was invoked.

# 4
# EFI System Table

## Introduction

The topics in this book describe the following:

- The entry point to a DXE or EFI image
- The parameters that are passed to that entry point
- Examples of how the various table examples are presented in the EFI environment

There are four types of EFI images that can be loaded and executed by EFI firmware:

- EFI applications
- EFI OS loaders
- DXE drivers
- EFI drivers

There are no differences in the entry point for these four image types.

## EFI Image Entry Point

Two parameters are passed to the entry point of an EFI image:

- The image handle of the EFI image being executed
- A pointer to the EFI System Table

The EFI System Table contains pointers to the following:

- Active console devices
- EFI Boot Services Table
- EFI Runtime Services Table
- List of EFI Configuration Tables such as the DXE Services Table, HOB list, ACPI table, SMBIOS table, and SAL System Table

# EFI_IMAGE_ENTRY_POINT

## Summary

This function is the main entry point for a DXE or EFI image. This entry point is the same for EFI applications, EFI OS loaders, DXE drivers, and EFI drivers including both device drivers and bus drivers.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_IMAGE_ENTRY_POINT) (
  IN  EFI_HANDLE          ImageHandle,
  IN  EFI_SYSTEM_TABLE    *SystemTable
  );
```

## Parameters

*ImageHandle*

The firmware allocated handle for the EFI image.

*SystemTable*

A pointer to the EFI System Table.

## Description

This function is the entry point to an EFI image. An EFI image is loaded and relocated in system memory by the EFI Boot Service **LoadImage()**. An EFI image is invoked through the EFI Boot Service **StartImage()**.

The first argument is the image's image handle. The second argument is a pointer to the image's system table. The system table contains the standard output and input handles, plus pointers to the **EFI_BOOT_SERVICES** and **EFI_RUNTIME_SERVICES** tables. The service tables contain the entry points in the firmware for accessing the core EFI system functionality. The handles in the system table are used to obtain basic access to the console. In addition, the EFI system table contains pointers to other standard tables that a loaded image may use if the associated pointers are initialized to nonzero values. Examples of such tables are DXE Services, HOB List, ACPI, SMBIOS, and SAL System Table.

The *ImageHandle* is a firmware-allocated handle that is used to identify the image on various functions. The handle also supports one or more protocols that the image can use. All images support the **EFI_LOADED_IMAGE** protocol that returns the source location of the image, the memory location of the image, the load options for the image, etc. The exact **EFI_LOADED_IMAGE** structure is defined in the *EFI 1.10 Specification*.

If the EFI image is an EFI application, then the EFI application executes and either returns or calls the EFI Boot Services **Exit()**. An EFI application is always unloaded from memory when it exits, and its return status is returned to the component that started the EFI application.

If the EFI image is an EFI OS loader, then the EFI OS loader executes and either returns, calls the EFI Boot Service **Exit()**, or calls the EFI Boot Service **ExitBootServices()**. If the EFI

OS Loader returns or calls **`Exit()`**, then the load of the OS has failed, and the EFI OS loader is unloaded from memory and control is returned to the component that attempted to boot the EFI OS loader. If **`ExitBootServices()`** is called, then the OS loader has taken control of the platform, and EFI will not regain control of the system until the platform is reset. One method of resetting the platform is through the EFI Runtime Service **`ResetSystem()`**.

If the EFI image is an EFI driver, then the EFI driver executes and either returns or calls the EFI Boot Service **`Exit()`**. If an EFI driver returns an error, then the driver is unloaded from memory. If the EFI driver returns **`EFI_SUCCESS`**, then it stays resident in memory. If the EFI driver does **not** follow the EFI Driver Model, then it performs any required initialization and installs its protocol services before returning. If the EFI driver *does* follow the EFI Driver Model, then the entry point is not allowed to touch any device hardware. Instead, the entry point is required to create and install the **`EFI_DRIVER_BINDING_PROTOCOL`** (defined in the *EFI 1.10 Specification*) on the *`ImageHandle`* of the EFI driver. If this process is completed, then **`EFI_SUCCESS`** is returned. If the resources are not available to complete the driver initialization, then **`EFI_OUT_OF_RESOURCES`** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The driver was initialized. |
| EFI_OUT_OF_RESOURCES | The request could not be completed due to a lack of resources. |
| EFI_UNLOAD_IMAGE | The driver was initialized, and the driver should be unloaded from memory |
| Other error codes | The driver failed to initialize, and the driver should be unloaded from memory. |

# EFI Table Header

## Summary

Data structure that precedes all of the standard EFI table types.

## Related Definitions

```
typedef struct {
  UINT64    Signature;
  UINT32    Revision;
  UINT32    HeaderSize;
  UINT32    CRC32;
  UINT32    Reserved;
} EFI_TABLE_HEADER;
```

## Parameters

*Signature*

A 64-bit signature that identifies the type of table that follows. Unique signatures have been generated for the EFI System Table, the EFI Boot Services Table, and the EFI Runtime Services Table.

*Revision*

The revision of the EFI specification to which this table conforms. The upper 16 bits of this field contain the major revision value, and the lower 16 bits contain the minor revision value. The minor revision values are limited to the range of 00..99.

*HeaderSize*

The size in bytes of the entire table including the **EFI_TABLE_HEADER**.

*CRC32*

The 32-bit CRC for the entire table. This value is computed by setting this field to 0 and computing the 32-bit CRC for *HeaderSize* bytes.

*Reserved*

Reserved field that must be set to 0.

## NOTE

*The size of the EFI System Table, EFI Runtime Services Table, and EFI Boot Services Table might increase over time. It is very important to always use the HeaderSize field of* **EFI_TABLE_HEADER** *to determine the size of these tables.*

## Description

The data type **EFI_TABLE_HEADER** is the data structure that precedes all of the standard EFI table types. It includes a signature that is unique for each table type, a revision of the table that may be updated as extensions are added to the EFI table types, and a 32-bit CRC so a consumer of an EFI table type can validate the contents of the EFI table.

# EFI System Table

## Summary

Contains pointers to the runtime and boot services tables.

## Related Definitions

```
#define EFI_SYSTEM_TABLE_SIGNATURE      0x5453595320494249
#define EFI_SYSTEM_TABLE_REVISION       ((1<<16) | (10))
#define EFI_1_10_SYSTEM_TABLE_REVISION  ((1<<16) | (10))
#define EFI_1_02_SYSTEM_TABLE_REVISION  ((1<<16) | (02))

typedef struct {
  EFI_TABLE_HEADER              Hdr;
  CHAR16                        *FirmwareVendor;
  UINT32                        FirmwareRevision;
  EFI_HANDLE                    ConsoleInHandle;
  SIMPLE_INPUT_INTERFACE        *ConIn;
  EFI_HANDLE                    ConsoleOutHandle;
  SIMPLE_TEXT_OUTPUT_INTERFACE  *ConOut;
  EFI_HANDLE                    StandardErrorHandle;
  SIMPLE_TEXT_OUTPUT_INTERFACE  *StdErr;
  EFI_RUNTIME_SERVICES          *RuntimeServices;
  EFI_BOOT_SERVICES             *BootServices;
  UINTN                         NumberOfTableEntries;
  EFI_CONFIGURATION_TABLE       *ConfigurationTable;
} EFI_SYSTEM_TABLE;
```

## Parameters

*Hdr*

The table header for the EFI System Table. This header contains the **EFI_SYSTEM_TABLE_SIGNATURE** and **EFI_SYSTEM_TABLE_REVISION** values along with the size of the **EFI_SYSTEM_TABLE** structure and a 32-bit CRC to verify that the contents of the EFI System Table are valid.

*FirmwareVendor*

A pointer to a null terminated Unicode string that identifies the vendor that produces the system firmware for the platform.

*FirmwareRevision*

A firmware vendor specific value that identifies the revision of the system firmware for the platform.

*ConsoleInHandle*

The handle for the active console input device. This handle must support the **SIMPLE_INPUT_PROTOCOL**. This handle is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*ConIn*

A pointer to the **SIMPLE_INPUT_PROTOCOL** interface that is associated with *ConsoleInHandle*. This interface is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*ConsoleOutHandle*

The handle for the active console output device. This handle must support the **SIMPLE_TEXT_OUTPUT_PROTOCOL**. This handle is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*ConOut*

A pointer to the **SIMPLE_TEXT_OUTPUT_PROTOCOL** interface that is associated with *ConsoleOutHandle*. This interface is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*StandardErrorHandle*

The handle for the active standard error console device. This handle must support the **SIMPLE_TEXT_OUTPUT_PROTOCOL**. This handle is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*StdErr*

A pointer to the **SIMPLE_TEXT_OUTPUT_PROTOCOL** interface that is associated with *StandardErrorHandle*. This interface is only valid after the BDS phase has connected the console devices, and before **ExitBootServices()** is called.

*RuntimeServices*

A pointer to the EFI Runtime Services Table.

*BootServices*

A pointer to the EFI Boot Services Table.

*NumberOfTableEntries*

The number of EFI Configuration Tables in the buffer *ConfigurationTable*.

*ConfigurationTable*

A pointer to the EFI Configuration Tables. The number of entries in the table is *NumberOfTableEntries*.

## Description

The EFI System Table contains pointers to the runtime and boot services tables. Except for the table header, all elements in the service tables are prototypes of function pointers to functions as defined in the following books:

- Services - Boot Services
- Services - Runtime Services

Prior to a call to **`ExitBootServices()`**, all of the fields of the EFI System Table are valid. After an operating system has taken control of the platform with a call to **`ExitBootServices()`**, only the following fields are valid:

- *`Hdr`*
- *`FirmwareVendor`*
- *`FirmwareRevision`*
- *`RuntimeServices`*
- *`NumberOfTableEntries`*
- *`ConfigurationTable`*

# EFI Boot Services Table

## EFI_BOOT_SERVICES

### Summary

Contains a table header and pointers to all of the boot services.

### Related Definitions

```
#define EFI_BOOT_SERVICES_SIGNATURE     0x56524553544f4f42
#define EFI_BOOT_SERVICES_REVISION      ((1<<16) | (10))

typedef struct {
  EFI_TABLE_HEADER                  Hdr;

  //
  // Task Priority Services
  //
  EFI_RAISE_TPL                     RaiseTPL;
  EFI_RESTORE_TPL                   RestoreTPL;

  //
  // Memory Services
  //
  EFI_ALLOCATE_PAGES                AllocatePages;
  EFI_FREE_PAGES                    FreePages;
  EFI_GET_MEMORY_MAP                GetMemoryMap;
  EFI_ALLOCATE_POOL                 AllocatePool;
  EFI_FREE_POOL                     FreePool;

  //
  // Event & Timer Services
  //
  EFI_CREATE_EVENT                  CreateEvent;
  EFI_SET_TIMER                     SetTimer;
  EFI_WAIT_FOR_EVENT                WaitForEvent;
  EFI_SIGNAL_EVENT                  SignalEvent;
  EFI_CLOSE_EVENT                   CloseEvent;
  EFI_CHECK_EVENT                   CheckEvent;
```

```
//
// Protocol Handler Services
//
EFI_INSTALL_PROTOCOL_INTERFACE   InstallProtocolInterface;
EFI_REINSTALL_PROTOCOL_INTERFACE ReinstallProtocolInterface;
EFI_UNINSTALL_PROTOCOL_INTERFACE UninstallProtocolInterface;
EFI_HANDLE_PROTOCOL              HandleProtocol;
EFI_HANDLE_PROTOCOL              PCHandleProtocol;
EFI_REGISTER_PROTOCOL_NOTIFY     RegisterProtocolNotify;
EFI_LOCATE_HANDLE                LocateHandle;
EFI_LOCATE_DEVICE_PATH           LocateDevicePath;
EFI_INSTALL_CONFIGURATION_TABLE  InstallConfigurationTable;

//
// Image Services
//
EFI_IMAGE_LOAD                   LoadImage;
EFI_IMAGE_START                  StartImage;
EFI_EXIT                         Exit;
EFI_IMAGE_UNLOAD                 UnloadImage;
EFI_EXIT_BOOT_SERVICES           ExitBootServices;

//
// Miscellaneous Services
//
EFI_GET_NEXT_MONOTONIC_COUNT     GetNextMonotonicCount;
EFI_STALL                        Stall;
EFI_SET_WATCHDOG_TIMER           SetWatchdogTimer;

//
// Driver Support Services
//
EFI_CONNECT_CONTROLLER           ConnectController;
EFI_DISCONNECT_CONTROLLER        DisconnectController;

//
// Open and Close Protocol Services
//
EFI_OPEN_PROTOCOL                OpenProtocol;
EFI_CLOSE_PROTOCOL               CloseProtocol;
EFI_OPEN_PROTOCOL_INFORMATION    OpenProtocolInformation;
```

```
    //
    // Extended Protocol Handler Services
    //
    EFI_PROTOCOLS_PER_HANDLE          ProtocolsPerHandle;
    EFI_LOCATE_HANDLE_BUFFER          LocateHandleBuffer;
    EFI_LOCATE_PROTOCOL               LocateProtocol;

    EFI_INSTALL_MULTIPLE_PROTOCOL_INTERFACES
                              InstallMultipleProtocolInterfaces;
    EFI_UNINSTALL_MULTIPLE_PROTOCOL_INTERFACES
                              IninstallMultipleProtocolInterfaces;

    //
    // 32-bit CRC Services
    //
    EFI_CALCULATE_CRC32               CalculateCrc32;

    //
    // Memory Utility Services
    //
    EFI_COPY_MEM                      CopyMem;
    EFI_SET_MEM                       SetMem;

} EFI_BOOT_SERVICES;
```

## Parameters

*Hdr*

The table header for the EFI Boot Services Table. This header contains the **EFI_BOOT_SERVICES_SIGNATURE** and **EFI_BOOT_SERVICES_REVISION** values along with the size of the **EFI_BOOT_SERVICES_TABLE** structure and a 32-bit CRC to verify that the contents of the EFI Boot Services Table are valid.

*RaiseTPL*

Raises the task priority level.

*RestoreTPL*

Restores/lowers the task priority level.

*AllocatePages*

Allocates pages of a particular type.

*FreePages*

Frees allocated pages.

*GetMemoryMap*

Returns the current boot services memory map and memory map key.

*AllocatePool*

Allocates a pool of a particular type.

*FreePool*

Frees allocated pool.

*CreateEvent*

Creates a general-purpose event structure. See the **CreateEvent()** function description in this document.

*SetTimer*

Sets an event to be signaled at a particular time.

*WaitForEvent*

Stops execution until an event is signaled.

*SignalEvent*

Signals an event.

*CloseEvent*

Closes and frees an event structure.

*CheckEvent*

Checks whether an event is in the signaled state.

*InstallProtocolInterface*

Installs a protocol interface on a device handle.

*ReinstallProtocolInterface*

Reinstalls a protocol interface on a device handle.

*UninstallProtocolInterface*

Removes a protocol interface from a device handle.

*HandleProtocol*

Queries a handle to determine if it supports a specified protocol.

*PCHandleProtocol*

Reserved. Must be **NULL**.

*RegisterProtocolNotify*

Registers an event that is to be signaled whenever an interface is installed for a specified protocol.

*LocateHandle*

Returns an array of handles that support a specified protocol.

*LocateDevicePath*

Locates all devices on a device path that support a specified protocol and returns the handle to the device that is closest to the path.

*InstallConfigurationTable*

Adds, updates, or removes a configuration table from the EFI System Table.

*LoadImage*

Loads an EFI image into memory. See the **LoadImage()** function description in this document.

*StartImage*

Transfers control to a loaded image's entry point.

*Exit*

Exits the image's entry point.

*UnloadImage*

Unloads an image.

*ExitBootServices*

Terminates boot services.

*GetNextMonotonicCount*

Returns a monotonically increasing count for the platform.

*Stall*

Stalls the processor.

*SetWatchdogTimer*

Resets and sets a watchdog timer used during boot services time.

*ConnectController*

Uses a set of precedence rules to find the best set of drivers to manage a controller.

*DisconnectController*

Informs a set of drivers to stop managing a controller.

*OpenProtocol*

Adds elements to the list of agents consuming a protocol interface.

*CloseProtocol*

Removes elements from the list of agents consuming a protocol interface.

*OpenProtocolInformation*

Retrieve the list of agents that are currently consuming a protocol interface.

*ProtocolsPerHandle*

Retrieves the list of protocols installed on a handle. The return buffer is automatically allocated.

*LocateHandleBuffer*

Retrieves the list of handles from the handle database that meet the search criteria. The return buffer is automatically allocated.

*LocateProtocol*

Finds the first handle in the handle database the supports the requested protocol.

*InstallMultipleProtocolInterfaces*

Installs one or more protocol interfaces onto a handle.

*UninstallMultipleProtocolInterfaces*

Uninstalls one or more protocol interfaces from a handle.

*CalculateCrc32*

Computes and returns a 32-bit CRC for a data buffer.

*CopyMem*

Copies the contents of one buffer to another buffer.

*SetMem*

Fills a buffer with a specified value.

## Description

The EFI Boot Services Table contains a table header and pointers to all of the boot services. Except for the table header, all elements in the EFI Boot Services Tables are prototypes of function pointers to functions as defined in Services - Boot Services. The function pointers in this table are not valid after the operating system has taken control of the platform with a call to **ExitBootServices()**.

# EFI Runtime Services Table

## EFI_RUNTIME_SERVICES

### Summary

Contains a table header and pointers to all of the runtime services.

### Related Definitions

```
#define EFI_RUNTIME_SERVICES_SIGNATURE  0x56524553544e5552
#define EFI_RUNTIME_SERVICES_REVISION   ((1<<16) | (10))

typedef struct {
  EFI_TABLE_HEADER                Hdr;

  //
  // Time Services
  //
  EFI_GET_TIME                    GetTime;
  EFI_SET_TIME                    SetTime;
  EFI_GET_WAKEUP_TIME             GetWakeupTime;
  EFI_SET_WAKEUP_TIME             SetWakeupTime;

  //
  // Virtual Memory Services
  //
  EFI_SET_VIRTUAL_ADDRESS_MAP     SetVirtualAddressMap;
  EFI_CONVERT_POINTER             ConvertPointer;

  //
  // Variable Services
  //
  EFI_GET_VARIABLE                GetVariable;
  EFI_GET_NEXT_VARIABLE_NAME      GetNextVariableName;
  EFI_SET_VARIABLE                SetVariable;

  //
  // Miscellaneous Services
  //
  EFI_GET_NEXT_HIGH_MONO_COUNT    GetNextHighMonotonicCount;
  EFI_RESET_SYSTEM                ResetSystem;

  //
  // Status Code Services
  //
  EFI_REPORT_STATUS_CODE          ReportStatusCode;

} EFI_RUNTIME_SERVICES;
```

## Parameters

*Hdr*

The table header for the EFI Runtime Services Table. This header contains the **EFI_RUNTIME_SERVICES_SIGNATURE** and **EFI_RUNTIME_SERVICES_REVISION** values along with the size of the **EFI_RUNTIME_SERVICES_TABLE** structure and a 32-bit CRC to verify that the contents of the EFI Runtime Services Table are valid.

*GetTime*

Returns the current time and date and the time-keeping capabilities of the platform.

*SetTime*

Sets the current local time and date information.

*GetWakeupTime*

Returns the current wake-up alarm clock setting.

*SetWakeupTime*

Sets the system wake-up alarm clock time.

*SetVirtualAddressMap*

Used by an OS loader to convert from physical addressing to virtual addressing.

*ConvertPointer*

Used by EFI components to convert internal pointers when switching to virtual addressing.

*GetVariable*

Returns the value of a variable.

*GetNextVariableName*

Enumerates the current variable names.

*SetVariable*

Sets the value of a variable.

*GetNextHighMonotonicCount*

Returns the next high 32 bits of the platform's monotonic counter.

*ResetSystem*

Resets the entire platform.

*ReportStatusCode*

Provides an interface that a software module can call to report a status code. See the **ReportStatusCode()** function description in this document.

## Description

The EFI Runtime Services Table contains a table header and pointers to all of the runtime services. Except for the table header, all elements in the EFI Runtime Services Tables are prototypes of function pointers to functions as defined Services - Runtime Services. Unlike the EFI Boot Services Table, this table and the function pointers it contains are valid after the operating system has taken control of the platform with a call to **`ExitBootServices()`**. If a call to **`SetVirtualAddressMap()`** is made by the OS, then the function pointers in this table are fixed up to point to the new virtually mapped entry points.

# EFI Configuration Table

## Summary

The *ConfigurationTable* field of the EFI System Table points to a list of GUID/pointer pairs. The lists of GUIDs below are required for OS and firmware interoperability. Other GUIDs may be defined as required by different IBV, OEMs, IHVs, and OSVs.

## Related Definitions

```
typedef struct{
  EFI_GUID                          VendorGuid;
  VOID                              *VendorTable;
} EFI_CONFIGURATION_TABLE;
```

## Parameters

*VendorGuid*

The 128-bit GUID value that uniquely identifies the EFI Configuration Table. See GUID Definitions below for GUID values defined by this specification.

*VendorTable*

A pointer to the table associated with *VendorGuid*.

## Description

The EFI Configuration Table is the *ConfigurationTable* field in the EFI System Table. This table contains a set of GUID/pointer pairs. Each element of this table is described by this **EFI_CONFIGURATION_TABLE** structure. The number of types of configuration tables is expected to grow over time, which is why a GUID is used to identify the configuration table type. The EFI Configuration Table may contain at most once instance of each table type.

## GUID Definitions

```
#define DXE_SERVICES_TABLE_GUID  \
{0x5ad34ba,0x6f02,0x4214,0x95,0x2e,0x4d,0xa0,0x39,0x8e,0x2b,0xb9}

#define HOB_LIST_GUID  \
{0x7739f24c,0x93d7,0x11d4,0x9a,0x3a,0x0,0x90,0x27,0x3f,0xc1,0x4d}

#define ACPI_20_TABLE_GUID  \
{0x8868e871,0xe4f1,0x11d3,0xbc,0x22,0x0,0x80,0xc7,0x3c,0x88,0x81}

#define ACPI_TABLE_GUID \
{0xeb9d2d30,0x2d88,0x11d3,0x9a,0x16,0x0,0x90,0x27,0x3f,0xc1,0x4d}
```

```
#define SAL_SYSTEM_TABLE_GUID \
{0xeb9d2d32,0x2d88,0x11d3,0x9a,0x16,0x0,0x90,0x27,0x3f,0xc1,0x4d}

#define SMBIOS_TABLE_GUID \
{0xeb9d2d31,0x2d88,0x11d3,0x9a,0x16,0x0,0x90,0x27,0x3f,0xc1,0x4d}

#define MPS_TABLE_GUID \
{0xeb9d2d2f,0x2d88,0x11d3,0x9a,0x16,0x0,0x90,0x27,0x3f,0xc1,0x4d}
```

# DXE Services Table

## DXE_SERVICES

### Summary

Contains a table header and pointers to all of the DXE-specific services.

### Related Definitions

```
#define DXE_SERVICES_SIGNATURE  0x565245535f455844
#define DXE_SERVICES_REVISION   ((0<<16) | (90)

typedef struct {
  EFI_TABLE_HEADER                Hdr;

  //
  // Global Coherency Domain Services
  //
  EFI_ADD_MEMORY_SPACE            AddMemorySpace;
  EFI_ALLOCATE_MEMORY_SPACE       AllocateMemorySpace;
  EFI_FREE_MEMORY_SPACE           FreeMemorySpace;
  EFI_REMOVE_MEMORY_SPACE         RemoveMemorySpace;
  EFI_GET_MEMORY_SPACE_DESCRIPTOR GetMemorySpaceDescriptor;
  EFI_SET_MEMORY_SPACE_ATTRIBUTES SetMemorySpaceAttributes;
  EFI_GET_MEMORY_SPACE_MAP        GetMemorySpaceMap;
  EFI_ADD_IO_SPACE                AddIoSpace;
  EFI_ALLOCATE_IO_SPACE           AllocateIoSpace;
  EFI_FREE_IO_SPACE               FreeIoSpace;
  EFI_REMOVE_IO_SPACE             RemoveIoSpace;
  EFI_GET_IO_SPACE_DESCRIPTOR     GetIoSpaceDescriptor;
  EFI_GET_IO_SPACE_MAP            GetIoSpaceMap;

  //
  // Dispatcher Services
  //
  EFI_DISPATCH                    Dispatch;
  EFI_SCHEDULE                    Schedule;
  EFI_TRUST                       Trust;

  //
  // Service to process a single firmware volume found in a
capsule
  //
  EFI_PROCESS_FIRMWARE_VOLUME     ProcessFirmwareVolume;
} DXE_SERVICES;
```

## Parameters

*Hdr*

The table header for the DXE Services Table. This header contains the **DXE_SERVICES_SIGNATURE** and **DXE_SERVICES_REVISION** values along with the size of the **DXE_SERVICES_TABLE** structure and a 32-bit CRC to verify that the contents of the DXE Services Table are valid.

*AddMemorySpace*

Adds reserved memory, system memory, or memory-mapped I/O resources to the global coherency domain of the processor. See the **AddMemorySpace()** function description in this document.

*AllocateMemorySpace*

Allocates nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. See the **AllocateMemorySpace()** function description in this document.

*FreeMemorySpace*

Frees nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. See the **FreeMemorySpace()** function description in this document.

*RemoveMemorySpace*

Removes reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. See the **RemoveMemorySpace()** function description in this document.

*GetMemorySpaceDescriptor*

Retrieves the descriptor for a memory region containing a specified address. See the **GetMemorySpaceDescriptor()** function description in this document.

*SetMemorySpaceAttributes*

Modifies the attributes for a memory region in the global coherency domain of the processor. See the **SetMemorySpaceAttributes()** function description in this document.

*GetMemorySpaceMap*

Returns a map of the memory resources in the global coherency domain of the processor. See the **GetMemorySpaceMap()** function description in this document.

*AddIoSpace*

Adds reserved I/O or I/O resources to the global coherency domain of the processor. See the **AddIoSpace()** function description in this document.

*AllocateIoSpace*

Allocates nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor. See the **AllocateIoSpace()** function description in this document.

*FreeIoSpace*

Frees nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor. See the **FreeIoSpace()** function description in this document.

*RemoveIoSpace*

Removes reserved I/O or I/O resources from the global coherency domain of the processor. See the **RemoveIoSpace()** function description in this document.

*GetIoSpaceDescriptor*

Retrieves the descriptor for an I/O region containing a specified address. See the **GetIoSpaceDescriptor()** function description in this document.

*GetIoSpaceMap*

Returns a map of the I/O resources in the global coherency domain of the processor. See the **GetIoSpaceMap()** function description in this document.

*Dispatch*

Loads and executed DXE drivers from firmware volumes. See the **Dispatch()** function description in this document.

*Schedule*

Clears the Schedule on Request (SOR) flag for a component that is stored in a firmware volume. See the **Schedule()** function description in this document.

*Trust*

Promotes a file stored in a firmware volume from the untrusted to the trusted state. See the **Trust()** function description in this document.

*ProcessFirmwareVolume*

Creates a firmware volume handle for a firmware volume that is present in system memory. See the **ProcessFirmwareVolume()** function description in this document.

## Description

The EFI DXE Services Table contains a table header and pointers to all of the DXE-specific services. Except for the table header, all elements in the DXE Services Tables are prototypes of function pointers to functions as defined in Services - DXE Services.

# EFI Image Entry Point Examples

## EFI Application Example

The following example shows the EFI image entry point for an EFI application. This application makes use of the EFI System Table, EFI Boot Services Table, EFI Runtime Services Table, and DXE Services Table.

```
EFI_GUID  gEfiDxeServicesTableGuid = DXE_SERVICES_TABLE_GUID;

EFI_SYSTEM_TABLE            *gST;
EFI_BOOT_SERVICES           *gBS;
EFI_RUNTIME_SERVICES        *gRT;
DXE_SERVICES                *gDS;

EfiApplicationEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  UINTN                             Index;
  BOOLEAN                           Result;
  EFI_STATUS                        Status;
  EFI_TIME                          *Time;
  UINTN                             NumberOfDescriptors;
  EFI_GCD_MEMORY_SPACE_DESCRIPTOR   MemorySpaceDescriptor;

  gST = SystemTable;
  gBS = gST->BootServices;
  gRT = gST->RuntimeServices;

  gDS = NULL;
  for (Index = 0; Index < gST->NumberOfTableEntries; Index++) {
    Result = EfiCompareGuid (
               &gEfiDxeServicesTableGuid,
               &(gST->ConfigurationTable[Index].VendorGuid)
               );
    if (Result) {
      gDS = gST->ConfigurationTable[Index].VendorTable;
    }
  }
  if (gDS == NULL) {
    return EFI_NOT_FOUND;
  }
```

```
//
// Use EFI System Table to print “Hello World” to the active console
// output device.
//
Status = gST->ConOut->OutputString (gST->ConOut, L”Hello World\n\r”);
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Use EFI Boot Services Table to allocate a buffer to store the
// current time and date.
//
Status = gBS->AllocatePool (
                EfiBootServicesData,
                sizeof (EFI_TIME),
                (VOID **)&Time
                );
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Use the EFI Runtime Services Table to get the current
// time and date.
//
Status = gRT->GetTime (&Time, NULL)
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Use EFI Boot Services to free the buffer that was used to store
// the current time and date.
//
Status = gBS->FreePool (Time);
if (EFI_ERROR (Status)) {
  return Status;
}

//
// Use the DXE Services Table to get the current GCD Memory Space Map
//
Status = gDS->GetMemorySpaceMap (
                &NumberOfDescriptors,
                &MemorySpaceMap
                );
if (EFI_ERROR (Status)) {
  return Status;
}
```

```
  //
  // Use EFI Boot Services to free the buffer that was used to store
  // the GCD Memory Space Map.
  //
  Status = gBS->FreePool (MemorySpaceMap);
  if (EFI_ERROR (Status)) {
    return Status;
  }

  return Status;
}
```

## Non-EFI Driver Model Example (Resident in Memory)

The following example shows the EFI image entry point for an EFI driver that does not follow the *EFI Driver Model*. Because this driver returns **EFI_SUCCESS**, it will stay resident in memory after it exits.

```
EFI_GUID  gEfiDxeServicesTableGuid = DXE_SERVICES_TABLE_GUID;

EFI_SYSTEM_TABLE            *gST;
EFI_BOOT_SERVICES           *gBS;
EFI_RUNTIME_SERVICES        *gRT;
DXE_SERVICES                *gDS;

EfiDriverEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  UINTN    Index;
  BOOLEAN  Result;

  gST = SystemTable;
  gBS = gST->BootServices;
  gRT = gST->RuntimeServices;

  gDS = NULL;
  for (Index = 0; Index < gST->NumberOfTableEntries; Index++) {
    Result = EfiCompareGuid (
               &gEfiDxeServicesTableGuid,
               &(gST->ConfigurationTable[Index].VendorGuid)
               );
    if (Result) {
      gDS = gST->ConfigurationTable[Index].VendorTable;
    }
  }
  if (gDS == NULL) {
    return EFI_UNLOAD_IMAGE;
  }

  //
  // Implement driver initialization here.
  //

  return EFI_SUCCESS;
}
```

## Non-EFI Driver Model (Nonresident in Memory)

The following example shows the EFI image entry point for an EFI driver that also does not follow the *EFI Driver Model*. Because this driver returns the error code **EFI_UNLOAD_IMAGE**, it will not stay resident in memory after it exits.

```
EFI_GUID  gEfiDxeServicesTableGuid = DXE_SERVICES_TABLE_GUID;

EFI_SYSTEM_TABLE            *gST;
EFI_BOOT_SERVICES           *gBS;
EFI_RUNTIME_SERVICES        *gRT;
DXE_SERVICES                *gDS;

EfiDriverEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  UINTN    Index;
  BOOLEAN  Result;

  gST = SystemTable;
  gBS = gST->BootServices;
  gRT = gST->RuntimeServices;

  gDS = NULL;
  for (Index = 0; Index < gST->NumberOfTableEntries; Index++) {
    Result = EfiCompareGuid (
               &gEfiDxeServicesTableGuid,
               &(gST->ConfigurationTable[Index].VendorGuid)
               );
    if (Result) {
      gDS = gST->ConfigurationTable[Index].VendorTable;
    }
  }
  if (gDS == NULL) {
    return EFI_UNLOAD_IMAGE;
  }

  //
  // Implement driver initialization here.
  //

  return EFI_UNLOAD_IMAGE;
}
```

## EFI Driver Model Example

The following is an *EFI Driver Model* example that shows the driver initialization routine for the ABC device controller that is on the XYZ bus. The **EFI_DRIVER_BINDING_PROTOCOL** is defined in Chapter 9 of the *EFI 1.10 Specification.* The function prototypes for the **AbcSupported()**, **AbcStart()**, and **AbcStop()** functions are defined in Section 9.1 of the *EFI 1.10 Specification.* This function saves the driver's image handle and a pointer to the EFI Boot Services Table in global variables, so that the other functions in the same driver can have access to these values. It then creates an instance of the **EFI_DRIVER_BINDING_PROTOCOL** and installs it onto the driver's image handle.

```
extern EFI_GUID                     gEfiDriverBindingProtocolGuid;
EFI_BOOT_SERVICES                   *gBS;
static EFI_DRIVER_BINDING_PROTOCOL  mAbcDriverBinding = {
  AbcSupported,
  AbcStart,
  AbcStop,
  0x10,
  NULL,
  NULL
};

AbcEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  EFI_STATUS  Status;

  gBS = SystemTable->BootServices;

  mAbcDriverBinding->ImageHandle         = ImageHandle;
  mAbcDriverBinding->DriverBindingHandle = ImageHandle;

  Status = gBS->InstallMultipleProtocolInterfaces(
                 &mAbcDriverBinding->DriverBindingHandle,
                 &gEfiDriverBindingProtocolGuid, &mAbcDriverBinding,
                 NULL
                 );
  return Status;
}
```

## EFI Driver Model Example (Unloadable)

The following is the same *EFI Driver Model* example as in EFI Driver Model Example, except that it also includes the code required to allow the driver to be unloaded through the boot service **`Unload()`**. Any protocols installed or memory allocated in **`AbcEntryPoint()`** must be uninstalled or freed in the **`AbcUnload()`**. The **`AbcUnload()`** function first checks to see how many controllers this driver is currently managing. If the number of controllers is greater than zero, then this driver cannot be unloaded at this time, so an error is returned.

```
extern EFI_GUID                     gEfiLoadedImageProtocolGuid;
extern EFI_GUID                     gEfiDriverBindingProtocolGuid;
EFI_BOOT_SERVICES                   *gBS;
static EFI_DRIVER_BINDING_PROTOCOL  mAbcDriverBinding = {
  AbcSupported,
  AbcStart,
  AbcStop,
  1,
  NULL,
  NULL
};

EFI_STATUS
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  );

AbcEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  EFI_STATUS                 Status;
  EFI_LOADED_IMAGE_PROTOCOL  *LoadedImage;

  gBS = SystemTable->BootServices;

  Status = gBS->OpenProtocol (
                ImageHandle,
                &gEfiLoadedImageProtocolGuid,
                &LoadedImage,
                ImageHandle,
                NULL,
                EFI_OPEN_PROTOCOL_GET_PROTOCOL
                );
  if (EFI_ERROR (Status)) {
    return Status;
  }
  LoadedImage->Unload = AbcUnload;
```

```
  mAbcDriverBinding->ImageHandle         = ImageHandle;
  mAbcDriverBinding->DriverBindingHandle = ImageHandle;

  Status = gBS->InstallMultipleProtocolInterfaces(
                 &mAbcDriverBinding->DriverBindingHandle,
                 &gEfiDriverBindingProtocolGuid, &mAbcDriverBinding,
                 NULL
                 );

  return Status;
}

EFI_STATUS
AbcUnload (
  IN EFI_HANDLE  ImageHandle
  )

{
  EFI_STATUS  Status;
  UINTN       Count;

  Status = LibGetManagedControllerHandles (ImageHandle, &Count, NULL);
  if (EFI_ERROR (Status)) {
    return Status;
  }

  if (Count > 0) {
    return EFI_ACCESS_DENIED;
  }

  Status = gBS->UninstallMultipleProtocolInterfaces (
                 ImageHandle,
                 &gEfiDriverBindingProtocolGuid, &mAbcDriverBinding,
                 NULL
                 );
  return Status;
}
```

## EFI Driver Model Example (Multiple Instances)

The following is the same as the first *EFI Driver Model* example, except that it produces three **EFI_DRIVER_BINDING_PROTOCOL** instances. The first one is installed onto the driver's image handle. The other two are installed onto newly created handles.

```
extern EFI_GUID                     gEfiDriverBindingProtocolGuid;
EFI_BOOT_SERVICES                   *gBS;

static EFI_DRIVER_BINDING_PROTOCOL  mAbcDriverBindingA = {
  AbcSupportedA,
  AbcStartA,
  AbcStopA,
  1,
  NULL,
  NULL
};

static EFI_DRIVER_BINDING_PROTOCOL  mAbcDriverBindingB = {
  AbcSupportedB,
  AbcStartB,
  AbcStopB,
  1,
  NULL,
  NULL
};

static EFI_DRIVER_BINDING_PROTOCOL  mAbcDriverBindingC = {
  AbcSupportedC,
  AbcStartC,
  AbcStopC,
  1,
  NULL,
  NULL
};

AbcEntryPoint(
  IN EFI_HANDLE        ImageHandle,
  IN EFI_SYSTEM_TABLE  *SystemTable
  )

{
  EFI_STATUS  Status;

  gBS = SystemTable->BootServices;
```

```
  //
  // Install mAbcDriverBindingA onto ImageHandle
  //
  mAbcDriverBindingA->ImageHandle         = ImageHandle;
  mAbcDriverBindingA->DriverBindingHandle = ImageHandle;



  Status = gBS->InstallMultipleProtocolInterfaces(
                  &mAbcDriverBindingA->DriverBindingHandle,
                  &gEfiDriverBindingProtocolGuid, &mAbcDriverBindingA,
                  NULL
                  );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Install mAbcDriverBindingB onto a newly created handle
  //
  mAbcDriverBindingB->ImageHandle         = ImageHandle;
  mAbcDriverBindingB->DriverBindingHandle = NULL;

  Status = gBS->InstallMultipleProtocolInterfaces(
                  &mAbcDriverBindingB->DriverBindingHandle,
                  &gEfiDriverBindingProtocolGuid, &mAbcDriverBindingB,
                  NULL
                  );
  if (EFI_ERROR (Status)) {
    return Status;
  }

  //
  // Install mAbcDriverBindingC onto a newly created handle
  //
  mAbcDriverBindingC->ImageHandle         = ImageHandle;
  mAbcDriverBindingC->DriverBindingHandle = NULL;

  Status = gBS->InstallMultipleProtocolInterfaces(
                  &mAbcDriverBindingC->DriverBindingHandle,
                  &gEfiDriverBindingProtocolGuid, &mAbcDriverBindingC,
                  NULL
                  );

  return Status;
}
```

# 5
# Services - Boot Services

## EFI 1.10 Boot Services

The table below lists all the boot services that are documented in the *EFI 1.10 Specification*. See the *EFI 1.10 Specification* for a detailed description for each of these boot services.

This DXE CIS defines backward-compatible extensions to the following services:

- **`CreateEvent()`**
- **`LoadImage()`**

The details of these extensions are contained in the following topics. The extension to **`CreateEvent()`** is a candidate for inclusion in a future revision of the EFI specification.

**Table 5.1. Boot Services in the EFI 1.10 Specification**

| Name | Type | Description |
|---|---|---|
| CreateEvent | Boot | Creates a general-purpose event structure. |
| CloseEvent | Boot | Closes and frees an event structure. |
| SignalEvent | Boot | Signals an event. |
| WaitForEvent | Boot | Stops execution until an event is signaled. |
| CheckEvent | Boot | Checks whether an event is in the signaled state. |
| SetTimer | Boot | Sets an event to be signaled at a particular time. |
| RaiseTPL | Boot | Raises the task priority level. |
| RestoreTPL | Boot | Restores/lowers the task priority level. |
| AllocatePages | Boot | Allocates pages of a particular type. |
| FreePages | Boot | Frees allocated pages. |
| GetMemoryMap | Boot | Returns the current boot services memory map and memory map key. |
| AllocatePool | Boot | Allocates a pool of a particular type. |
| FreePool | Boot | Frees allocated pool. |
| InstallProtocolInterface | Boot | Installs a protocol interface on a device handle. |
| UninstallProtocolInterface | Boot | Removes a protocol interface from a device handle. |
| ReinstallProtocolInterface | Boot | Reinstalls a protocol interface on a device handle. |
| RegisterProtocolNotify | Boot | Registers an event that is to be signaled whenever an interface is installed for a specified protocol. |
| LocateHandle | Boot | Returns an array of handles that support a specified protocol. |
| HandleProtocol | Boot | Queries a handle to determine if it supports a specified protocol. |

| Name | Type | Description |
|---|---|---|
| LocateDevicePath | Boot | Locates all devices on a device path that support a specified protocol and returns the handle to the device that is closest to the path. |
| OpenProtocol | Boot | Adds elements to the list of agents consuming a protocol interface. |
| CloseProtocol | Boot | Removes elements from the list of agents consuming a protocol interface. |
| OpenProtocolInformation | Boot | Retrieve the list of agents that are currently consuming a protocol interface. |
| ConnectController | Boot | Uses a set of precedence rules to find the best set of drivers to manage a controller. |
| DisconnectController | Boot | Informs a set of drivers to stop managing a controller. |
| ProtocolsPerHandle | Boot | Retrieves the list of protocols installed on a handle. The return buffer is automatically allocated. |
| LocateHandleBuffer | Boot | Retrieves the list of handles from the handle database that meet the search criteria. The return buffer is automatically allocated. |
| LocateProtocol | Boot | Finds the first handle in the handle database the supports the requested protocol. |
| InstallMultipleProtocolInterfaces | Boot | Installs one or more protocol interfaces onto a handle. |
| UninstallMultipleProtocolInterfaces | Boot | Uninstalls one or more protocol interfaces from a handle. |
| LoadImage | Boot | Loads an EFI image into memory. |
| StartImage | Boot | Transfers control to a loaded image's entry point. |
| UnloadImage | Boot | Unloads an image. |
| EFI_IMAGE_ENTRY_POINT | Boot | Prototype of an EFI image's entry point. |
| Exit | Boot | Exits the image's entry point. |
| ExitBootServices | Boot | Terminates boot services. |
| SetWatchDogTimer | Boot | Resets and sets a watchdog timer used during boot services time. |
| Stall | Boot | Stalls the processor. |
| CopyMem | Boot | Copies the contents of one buffer to another buffer. |
| SetMem | Boot | Fills a buffer with a specified value. |
| GetNextMonotonicCount | Boot | Returns a monotonically increasing count for the platform. |
| InstallConfigurationTable | Boot | Adds, updates, or removes a configuration table from the EFI System Table. |
| CalculateCrc32 | Boot | Computes and returns a 32-bit CRC for a data buffer. |

# Extensions to EFI 1.10 Boot Services

## CreateEvent()

### Summary

Creates an event. This function has been extended from the **CreateEvent()** Boot Service defined in the *EFI 1.10 Specification*. The event types **EFI_EVENT_NOTIFY_SIGNAL_ALL** and **EFI_EVENT_SIGNAL_READY_TO_BOOT** have been added to this service.

### Prototype

```
EFI_STATUS
CreateEvent (
  IN UINT32             Type,
  IN EFI_TPL            NotifyTpl,
  IN EFI_EVENT_NOTIFY   NotifyFunction,
  IN VOID               *NotifyContext,
  OUT EFI_EVENT         *Event
  );
```

### Parameters

*Type*

The type of event to create and its mode and attributes. The **#define** statements in "Related Definitions" below can be used to specify an event's mode and attributes.

*NotifyTpl*

The task priority level of event notifications. Type **EFI_TPL** is defined in **RaiseTPL()** in the *EFI 1.10 Specification.*

*NotifyFunction*

Pointer to the event's notification function. Type **EFI_EVENT_NOTIFY** is defined in "Related Definitions" below.

*NotifyContext*

Pointer to the notification function's context; corresponds to parameter *Context* in the notification function.

*Event*

Pointer to the newly created event if the call succeeds; undefined otherwise. Type **EFI_EVENT** is defined in "Related Definitions" below.

### Description

The **CreateEvent()** function creates a new event of type *Type* and returns it in the location referenced by *Event*. The event's notification function, context, and task priority level are specified by *NotifyFunction*, *NotifyContext*, and *NotifyTpl*, respectively.

Events exist in one of two states, “waiting” or “signaled." When an event is created, firmware puts it in the "waiting” state. When the event is signaled, firmware changes its state to “signaled” and, if **EFI_EVENT_NOTIFY_SIGNAL** is specified, places a call to its notification function in a FIFO queue. There is a queue for each of the “basic” task priority levels defined in the *EFI 1.10 Specification* (**EFI_TPL_APPLICATION**, **EFI_TPL_CALLBACK**, and **EFI_TPL_NOTIFY**). The functions in these queues are invoked in FIFO order, starting with the highest priority level queue and proceeding to the lowest priority queue that is unmasked by the current TPL. If the current TPL is equal to or greater than the queued notification, it will wait until the TPL is lowered via **RestoreTPL()**.

In a general sense, there are two “types” of events, synchronous and asynchronous. Asynchronous events are closely related to timers and are used to support periodic or timed interruption of program execution. This capability is typically used with device drivers. For example, a network device driver that needs to poll for the presence of new packets could create an event whose type includes **EFI_EVENT_TIMER** and then call the **SetTimer()** function. When the timer expires, the firmware signals the event.

Synchronous events have no particular relationship to timers. Instead, they are used to ensure that certain activities occur following a call to a specific interface function. One example of this is the cleanup that needs to be performed in response to a call to the **ExitBootServices()** function. **ExitBootServices()** can clean up the firmware since it understands firmware internals, but it cannot clean up on behalf of drivers that have been loaded into the system. The drivers have to do that themselves by creating an event whose type is **EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES** and whose notification function is a function within the driver itself. Then, when **ExitBootServices()** has finished its cleanup, it signals each event of type **EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES**.

Another example of the use of synchronous events occurs when an event of type **EFI_EVENT_SIGNAL_VIRTUAL_ADDRESS_CHANGE** is used in conjunction with the **SetVirtualAddressmap()** function in Chapter 6 of the *EFI 1.10 Specification.*

The **EFI_EVENT_NOTIFY_WAIT** and **EFI_EVENT_NOTIFY_SIGNAL** flags are exclusive. If neither flag is specified, the caller does not require any notification concerning the event and the *NotifyTpl*, *NotifyFunction*, and *NotifyContext* parameters are ignored. If **EFI_EVENT_NOTIFY_WAIT** is specified, then the event is signaled and its notify function is queued whenever a consumer of the event is waiting for it (via **WaitForEvent()** or **CheckEvent()**). If the **EFI_EVENT_NOTIFY_SIGNAL** flag is specified then the event’s notify function is queued whenever the event is signaled.

**NOTE**

*Because its internal structure is unknown to the caller, Event cannot be modified by the caller. The only way to manipulate it is to use the published event interfaces.*

## Related Definitions

```
//*******************************************************
// EFI_EVENT
//*******************************************************
typedef VOID   *EFI_EVENT


//*******************************************************
// Event Types
//*******************************************************
// These types can be “ORed” together as needed – for example,
// EFI_EVENT_TIMER might be “ORed” with EFI_EVENT_NOTIFY_WAIT or
// EFI_EVENT_NOTIFY_SIGNAL.
#define EFI_EVENT_TIMER                         0x80000000
#define EFI_EVENT_RUNTIME                       0x40000000
#define EFI_EVENT_RUNTIME_CONTEXT               0x20000000

#define EFI_EVENT_NOTIFY_WAIT                   0x00000100
#define EFI_EVENT_NOTIFY_SIGNAL                 0x00000200
#define EFI_EVENT_NOTIFY_SIGNAL_ALL             0x00000400

#define EFI_EVENT_SIGNAL_READY_TO_BOOT          0x00000203
#define EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES     0x00000201
#define EFI_EVENT_SIGNAL_VIRTUAL_ADDRESS_CHANGE 0x60000202
#define EFI_EVENT_SIGNAL_LEGACY_BOOT            0x00000204
```

Following is a description of the fields in the above definition.

| | |
|---|---|
| EFI_EVENT_TIMER | The event is a timer event and may be passed to **SetTimer()**. Note that timers only function during boot services time. |
| EFI_EVENT_RUNTIME | The event is allocated from runtime memory. If an event is to be signaled after the call to **ExitBootServices()**, the event's data structure and notification function need to be allocated from runtime memory. For more information, see **SetVirtualAddressMap()** in Services - Runtime Services. |
| EFI_EVENT_RUNTIME_CONTEXT | The event's *NotifyContext* pointer points to a runtime memory address. See the above discussion of **EFI_EVENT_RUNTIME**. |
| EFI_EVENT_NOTIFY_WAIT | The event's *NotifyFunction* is to be invoked whenever the event is being waited on via **WaitForEvent()** or **CheckEvent()**. |
| EFI_EVENT_NOTIFY_SIGNAL | The event's *NotifyFunction* is to be invoked whenever the event is signaled via **SignalEvent()**. |

| | |
|---|---|
| EFI_EVENT_NOTIFY_SIGNAL_ALL | Used to signal all events of a specified type. For example, this bit may be used with **EFI_EVENT_SIGNAL_READY_TO_BOOT**. |
| EFI_EVENT_SIGNAL_READY_TO_BOOT | This event is to be notified by the system when the EFI Boot Manager is about to load and execute a boot option. |
| EFI_EVENT_SIGNAL_EXIT_BOOT_SERVICES | This event is to be notified by the system when **ExitBootServices()** is invoked. This type cannot be used with any other EVT bit type. The notification function for this event is not allowed to use the Memory Allocation Services, or call any functions that use the Memory Allocation Services, because these services modify the current memory map. |
| EFI_EVENT_SIGNAL_VIRTUAL_ADDRESS_CHANGE | The event is to be notified by the system when **SetVirtualAddressMap()** is performed. This type cannot be used with any other EVT bit type. See the discussion above of **EFI_EVENT_RUNTIME**. |
| EFI_EVENT_SIGNAL_LEGACY_BOOT | This event is to be notified by the system when the EFI Boot Manager is about to boot a legacy boot option. Events of this type are notified just before INT19 is invoked. |

```
//*******************************************************
// EFI_EVENT_NOTIFY
//*******************************************************
typedef
VOID
(EFIAPI *EFI_EVENT_NOTIFY) (
  IN EFI_EVENT  Event,
  IN VOID       *Context
  );
```

*Event*

Event whose notification function is being invoked. Type **EFI_EVENT** is defined above.

*Context*

Pointer to the notification function's context, which is implementation dependent. *Context* corresponds to *NotifyContext* in **CreateEvent()**.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The event structure was created. |
| EFI_INVALID_PARAMETER | One of the parameters has an invalid value. |
| EFI_OUT_OF_RESOURCES | The event could not be allocated. |

## LoadImage()

### Summary

Loads an EFI image into memory. This function has been extended from the **LoadImage()** Boot Service defined in the *EFI 1.10 Specification* to allow EFI images to be loaded from files stored in firmware volumes. It also validates the image using the services of the Security Architectural Protocol.

### Prototype

```
EFI_STATUS
LoadImage (
  IN BOOLEAN           BootPolicy,
  IN EFI_HANDLE        ParentImageHandle,
  IN EFI_DEVICE_PATH   *FilePath,
  IN VOID              *SourceBuffer  OPTIONAL ,
  IN UINTN             SourceSize,
  OUT EFI_HANDLE       *ImageHandle
  );
```

### Parameters

*BootPolicy*

If **TRUE**, indicates that the request originates from the boot manager, and that the boot manager is attempting to load *FilePath* as a boot selection. Ignored if *SourceBuffer* is not **NULL**.

*ParentImageHandle*

The caller's image handle. Type **EFI_HANDLE** is defined in the **InstallProtocolInterface()** function description in the *EFI 1.10 Specification*. This field is used to initialize the *ParentHandle* field of the **LOADED_IMAGE** protocol for the image that is being loaded.

*FilePath*

The specific file path from which the image is loaded. Type **EFI_DEVICE_PATH** is defined in the **LocateDevicePath()** function description in the *EFI 1.10 Specification*.

*SourceBuffer*

If not **NULL**, a pointer to the memory location containing a copy of the image to be loaded.

*SourceSize*

The size in bytes of *SourceBuffer*. Ignored if *SourceBuffer* is **NULL**.

*ImageHandle*

Pointer to the returned image handle that is created when the image is successfully loaded. Type **EFI_HANDLE** is defined in the **InstallProtocolInterface()** function description in the *EFI 1.10 Specification.*

## Description

The **LoadImage()** function loads an EFI image into memory and returns a handle to the image. The supported subsystem values in the PE image header are listed in "Related Definitions" below. The image is loaded in one of two ways. If *SourceBuffer* is not **NULL**, the function is a memory-to-memory load in which *SourceBuffer* points to the image to be loaded and *SourceSize* indicates the image's size in bytes. *FilePath* specifies where the image specified by *SourceBuffer* and *SourceSize* was loaded. In this case, the caller has copied the image into *SourceBuffer* and can free the buffer once loading is complete.

If *SourceBuffer* is **NULL**, the function is a file copy operation that uses the **EFI_FIRMWARE_VOLUME_PROTOCOL**, followed by the **SIMPLE_FILE_SYSTEM_PROTOCOL** and then the **LOAD_FILE_PROTOCOL** to access the file referred to by *FilePath*. In this case, the *BootPolicy* flag is passed to the **LOAD_FILE.LoadFile()** function and is used to load the default image responsible for booting when the *FilePath* only indicates the device. For more information see the discussion of the Load File Protocol in Chapter 11 of the *EFI 1.10 Specification.*

Regardless of the type of load (memory-to-memory or file copy), the function relocates the code in the image while loading it.

The image is also validated using the **FileAuthenticationState()** service of the Security Architectural Protocol (SAP). If the SAP returns the status **EFI_SUCCESS**, then the load operation is completed normally. If the SAP returns the status **EFI_SECURITY_VIOLATION**, then the load operation is completed normally, and the **EFI_SECURITY_VIOLATION** status is returned. In this case, the caller is not allowed to start the image until some platform specific policy is executed to protect the system while executing untrusted code. If the SAP returns the status **EFI_ACCESS_DENIED**, then the image should never be trusted. In this case, the image is unloaded from memory, and **EFI_ACCESS_DENIED** is returned.

Once the image is loaded, firmware creates and returns an **EFI_HANDLE** that identifies the image and supports the **LOADED_IMAGE_PROTOCOL**. The caller may fill in the image's "load options" data, or add additional protocol support to the handle before passing control to the newly loaded image by calling **StartImage()**. Also, once the image is loaded, the caller either starts it by calling **StartImage()** or unloads it by calling **UnloadImage()**.

## Related Definitions

```
//*************************************************************
// Supported subsystem values
//*************************************************************

#define EFI_IMAGE_SUBSYSTEM_EFI_APPLICATION          10
#define EFI_IMAGE_SUBSYSTEM_EFI_BOOT_SERVICE_DRIVER  11
#define EFI_IMAGE_SUBSYSTEM_EFI_RUNTIME_DRIVER       12
#define EFI_IMAGE_SUBSYSTEM_SAL_RUNTIME_DRIVER       13
```

Following is a description of the fields in the above definition.

| | |
|---|---|
| EFI_IMAGE_SUBSYSTEM_EFI_APPLICATION | The image is loaded into memory of type **EfiLoaderCode**, and the memory is freed when the application exits. |
| EFI_IMAGE_SUBSYSTEM_EFI_BOOT_SERVICE_DRIVER | The image is loaded into memory of type **EfiBootServicesCode**. If the image exits with an error code, then the memory for the image is free. If the image exits with **EFI_SUCCESS**, then the memory for the image is not freed. |
| EFI_IMAGE_SUBSYSTEM_EFI_RUNTIME_DRIVER | The image is loaded into memory of type **EfiRuntimeServicesCode**. If the image exits with an error code, then the memory for the image is free. If the image exits with **EFI_SUCCESS**, then the memory for the image is not freed. Images of this type are automatically converted from physical addresses to virtual address when the Runtime Service **SetVirtualAddressMap()** is called. |
| EFI_IMAGE_SUBSYSTEM_SAL_RUNTIME_DRIVER | The image is loaded into memory of type **EfiRuntimeServicesCode**. If the image exits with an error code, then the memory for the image is free. If the image exits with **EFI_SUCCESS**, then the memory for the image is not freed. Images of this type are not converted from physical to virtual addresses when the Runtime Service **SetVirtualAddressMap()** is called. |

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The image was loaded into memory. |
| EFI_SECURITY_VIOLATION | The image was loaded into memory, but the current security policy dictates that the image should not be executed at this time. |
| EFI_ACCESS_DENIED | The image was not loaded into memory because the current security policy dictates that the image should never be executed. |
| EFI_NOT_FOUND | The *FilePath* was not found. |
| EFI_INVALID_PARAMETER | One of the parameters has an invalid value. |
| EFI_UNSUPPORTED | The image type is not supported, or the device path cannot be parsed to locate the proper protocol for loading the file. |
| EFI_OUT_OF_RESOURCES | Image was not loaded due to insufficient resources. |
| EFI_LOAD_ERROR | Image was not loaded because the image format was corrupt or not understood. |
| EFI_DEVICE_ERROR | Image was not loaded because the device returned a read error. |

# 6
# Services - Runtime Services

## EFI 1.10 Runtime Services

The table below lists all the runtime services that are documented in the *EFI 1.10 Specification*. See the *EFI 1.10 Specification* for a detailed description for each of these runtime services.

This DXE CIS defines one additional runtime service:

- Status Code Services

The details of this additional service are contained in the following topics. This service is a candidate for inclusion in a future revision of the EFI specification.

**Table 6.1. EFI 1.10 Runtime Services**

| Name | Type | Description |
|---|---|---|
| GetVariable | Runtime | Returns the value of a variable. |
| GetNextVariableName | Runtime | Enumerates the current variable names. |
| SetVariable | Runtime | Sets the value of a variable. |
| GetTime | Runtime | Returns the current time and date, and the time-keeping capabilities of the platform. |
| SetTime | Runtime | Sets the current local time and date information. |
| GetWakeupTime | Runtime | Returns the current wake-up alarm clock setting. |
| SetWakeupTime | Runtime | Sets the system wake-up alarm clock time. |
| SetVirtualAddressMap | Runtime | Used by an OS loader to convert from physical addressing to virtual addressing. |
| ConvertPointer | Runtime | Used by EFI components to convert internal pointers when switching to virtual addressing. |
| ResetSystem | Runtime | Resets the entire platform. |
| GetNextHighMonotonicCount | Runtime | Returns the next high 32 bits of the platform's monotonic counter. |

# Additional Runtime Services

## Status Code Services

The table below lists the runtime services that are used to report status codes. These services are candidates for inclusion in a future revision of the EFI specification.

**Table 6.2. Status Code Runtime Services**

| Name | Type | Description |
|---|---|---|
| ReportStatusCode | Runtime | Reports status codes at boot services time and runtime. |

## ReportStatusCode()

### Summary

Provides an interface that a software module can call to report a status code.

### Prototype

```
EFI_STATUS
(EFIAPI *EFI_REPORT_STATUS_CODE) (
  IN EFI_STATUS_CODE_TYPE     Type,
  IN EFI_STATUS_CODE_VALUE    Value,
  IN UINT32                   Instance,
  IN EFI_GUID                 *CallerId  OPTIONAL,
  IN EFI_STATUS_CODE_DATA     *Data      OPTIONAL
  );
```

### Parameters

*Type*

Indicates the type of status code being reported. Type **EFI_STATUS_CODE_TYPE** is defined in "Related Definitions" below.

*Value*

Describes the current status of a hardware or software entity. This included information about the class and subclass that is used to classify the entity as well as an operation. For progress codes, the operation is the current activity. For error codes, it is the exception. For debug codes, it is not defined at this time. Type **EFI_STATUS_CODE_VALUE** is defined in "Related Definitions" below. Specific values are discussed in the *Intel® Platform Innovation Framework for EFI Status Code Specification*.

*Instance*

The enumeration of a hardware or software entity within the system. A system may contain multiple entities that match a class/subclass pairing. The instance differentiates between them. An instance of 0 indicates that instance information is unavailable, not meaningful, or not relevant. Valid instance numbers start with 1.

*CallerId*

This optional parameter may be used to identify the caller. This parameter allows the status code driver to apply different rules to different callers. Type **EFI_GUID** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification*.

*Data*

This optional parameter may be used to pass additional data. Type **EFI_STATUS_CODE_DATA** is defined in "Related Definitions" below. The contents of this data type may have additional GUID-specific data. The standard GUIDs and their associated data structures are defined in the *Intel® Platform Innovation Framework for EFI Status Code Specification*.

## Description

Various software modules including drivers can call this function to report a status code. No disposition of the status code is guaranteed. The **ReportStatusCode()** function may choose to log the status code, but this action is not required.

It is possible that this function may get called at **EFI_TPL_LEVEL_HIGH**. Therefore, this function cannot call any protocol interface functions or services (including memory allocation) that are not guaranteed to work at **EFI_TPL_LEVEL_HIGH**. It should be noted that **SignalEvent()** could be called by this function because it works at any TPL including **EFI_TPL_LEVEL_HIGH**. It is possible for an implementation to use events to log the status codes when the TPL level is reduced.

**ReportStatusCode()** function can perform other implementation specific work, but that is not specified in the architecture document.

In case of an error, the caller can specify the severity. In most cases, the entity that reports the error may not have a platform wide view and may not be able to accurately assess the impact of the error condition. The DXE driver that produces the Status Code Architectural Protocol, **EFI_STATUS_CODE_ARCH_PROTOCOL**, is responsible for assessing the true severity level based on the reported severity and other information. This DXE driver may perform platform specific actions based on the type and severity of the status code being reported.

If *Data* is present, the Status Code Architectural Protocol driver treats it as read only data. The Status Code Architectural Protocol driver must copy *Data* to a local buffer in an atomic operation before performing any other actions. This is necessary to make this function re-entrant. The size of the local buffer may be limited. As a result, some of the *Data* can be lost. The size of the local buffer should at least be 256 bytes in size. Larger buffers will reduce the probability of losing part of the *Data*. Note than multiple status codes may be reported at elevated TPL levels before the TPL level is reduced. Allocating multiple local buffers may reduce the probability losing status codes at elevated TPL levels. If all of the local buffers are consumed, then this service may not be able to perform the platform specific action required by the status code being reported. As a result, if all the local buffers are consumed, the behavior of this service is undefined.

If the *CallerId* parameter is not **NULL**, then it is required to point to a constant GUID. In other words, the caller may not reuse or release the buffer pointed to by *CallerId*.

## Related Definitions

```
//
// Status Code Type Definition
//
typedef UINT32 EFI_STATUS_CODE_TYPE;

//
// A Status Code Type is made up of the code type and severity
// All values masked by EFI_STATUS_CODE_RESERVED_MASK are
// reserved for use by this specification.
//
#define EFI_STATUS_CODE_TYPE_MASK       0x000000FF
#define EFI_STATUS_CODE_SEVERITY_MASK   0xFF000000
#define EFI_STATUS_CODE_RESERVED_MASK   0x00FFFF00

//
// Definition of code types, all other values masked by
// EFI_STATUS_CODE_TYPE_MASK are reserved for use by
// this specification.
//
#define EFI_PROGRESS_CODE               0x00000001
#define EFI_ERROR_CODE                  0x00000002
#define EFI_DEBUG_CODE                  0x00000003

//
// Definitions of severities, all other values masked by
// EFI_STATUS_CODE_SEVERITY_MASK are reserved for use by
// this specification.
// Uncontained errors are major errors that could not contained
// to the specific component that is reporting the error
// For example, if a memory error was not detected early enough,
// the bad data could be consumed by other drivers.
//
#define EFI_ERROR_MINOR                 0x40000000
#define EFI_ERROR_MAJOR                 0x80000000
#define EFI_ERROR_UNRECOVERED           0x90000000
#define EFI_ERROR_UNCONTAINED           0xa0000000

//
// Status Code Value Definition
//
typedef UINT32 EFI_STATUS_CODE_VALUE;

//
// A Status Code Value is made up of the class, subclass, and
// an operation.
//
#define EFI_STATUS_CODE_CLASS_MASK      0xFF000000
```

```
#define EFI_STATUS_CODE_SUBCLASS_MASK   0x00FF0000
#define EFI_STATUS_CODE_OPERATION_MASK  0x0000FFFF

//
// Definition of Status Code extended data header.
// The data will follow HeaderSize bytes from the beginning of
// the structure and is Size bytes long.
//
typedef struct {
  UINT16    HeaderSize;
  UINT16    Size;
  EFI_GUID  Type;
} EFI_STATUS_CODE_DATA;
```

*HeaderSize*

The size of the structure. This is specified to enable future expansion.

*Size*

The size of the data in bytes. This does not include the size of the header structure.

*Type*

The GUID defining the type of the data. The standard GUIDs and their associated data structures are defined in the *Intel® Platform Innovation Framework for EFI Status Code Specification*.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The function completed successfully |
| EFI_DEVICE_ERROR | The function should not be completed due to a device error. |

# 7
# Services - DXE Services

## Introduction

This chapter describes the services from the DXE Services Table. These services include the following:

- Global Coherency Domain (GCD) Services
- Dispatcher Services

The GCD Services are used to manage the system memory, memory-mapped I/O, and I/O resources present in a platform. The Dispatcher Services are used to invoke the DXE Dispatcher and modify the state of a DXE driver that is being tracked by the DXE Dispatcher.

## Global Coherency Domain Services

### Overview

### Global Coherency Domain (GCD) Services Overview

The Global Coherency Domain (GCD) Services are used to manage the memory and I/O resources visible to the boot processor. These resources are managed in two different maps:

- GCD memory space map
- GCD I/O space map

If memory or I/O resources are added, removed, allocated, or freed, then the GCD memory space map and GCD I/O space map are updated. GCD Services are also provided to retrieve the contents of these two resource maps.

The GCD Services can be broken up into two groups. The first manages the memory resources visible to the boot processor, and the second manages the I/O resources visible to the boot processor. Not all processor types support I/O resources, so the management of I/O resources may not be required. However, since system memory resources and memory-mapped I/O resources are required to execute the DXE environment, the management of memory resources is always required.

### GCD Memory Resources

The Global Coherency Domain (GCD) Services used to manage memory resources include the following:

- **`AddMemorySpace()`**
- **`AllocateMemorySpace()`**
- **`FreeMemorySpace()`**
- **`RemoveMemorySpace()`**
- **`SetMemorySpaceAttributes()`**

The GCD Services used to retrieve the GCD memory space map include the following:

- **GetMemorySpaceDescriptor()**
- **GetMemorySpaceMap()**

The GCD memory space map is initialized from the HOB list that is passed to the entry point of the DXE Foundation. One HOB type describes the number of address lines that are used to access memory resources. This information is used to initialize the state of the GCD memory space map. Any memory regions outside this initial region are not available to any of the GCD Services that are used to manage memory resources. The GCD memory space map is designed to describe the memory address space with as many as 64 address lines. Each region in the GCD memory space map can begin and end on a byte boundary. There are additional HOB types that describe the location of system memory, the location memory mapped I/O, the location of firmware devices, the location of firmware volumes, the location of reserved regions, and the location of system memory regions that were allocated prior to the execution of the DXE Foundation. The DXE Foundation must parse the contents of the HOB list to guarantee that memory regions reserved prior to the execution of the DXE Foundation are honored. As a result, the GCD memory space map must reflect the memory regions described in the HOB list. The GCD memory space map provides the DXE Foundation with the information required to initialize the memory services such as **AllocatePages()**, **FreePages()**, **AllocatePool()**, **FreePool()**, and **GetMemoryMap()**. See the *EFI 1.10 Specification* for definitions of these services.

A memory region described by the GCD memory space map can be in one of several different states:

- Nonexistent memory
- System memory
- Memory-mapped I/O
- Reserved memory

These memory regions can be allocated and freed by DXE drivers executing in the DXE environment. In addition, a DXE driver can attempt to adjust the caching attributes of a memory region. The figure below shows the possible state transitions for each byte of memory in the GCD memory space map. The transitions are labeled with the GCD Service that can move the byte from one state to another. The GCD services are required to merge similar memory regions that are adjacent to each other into a single memory descriptor, which reduces the number of entries in the GCD memory space map.

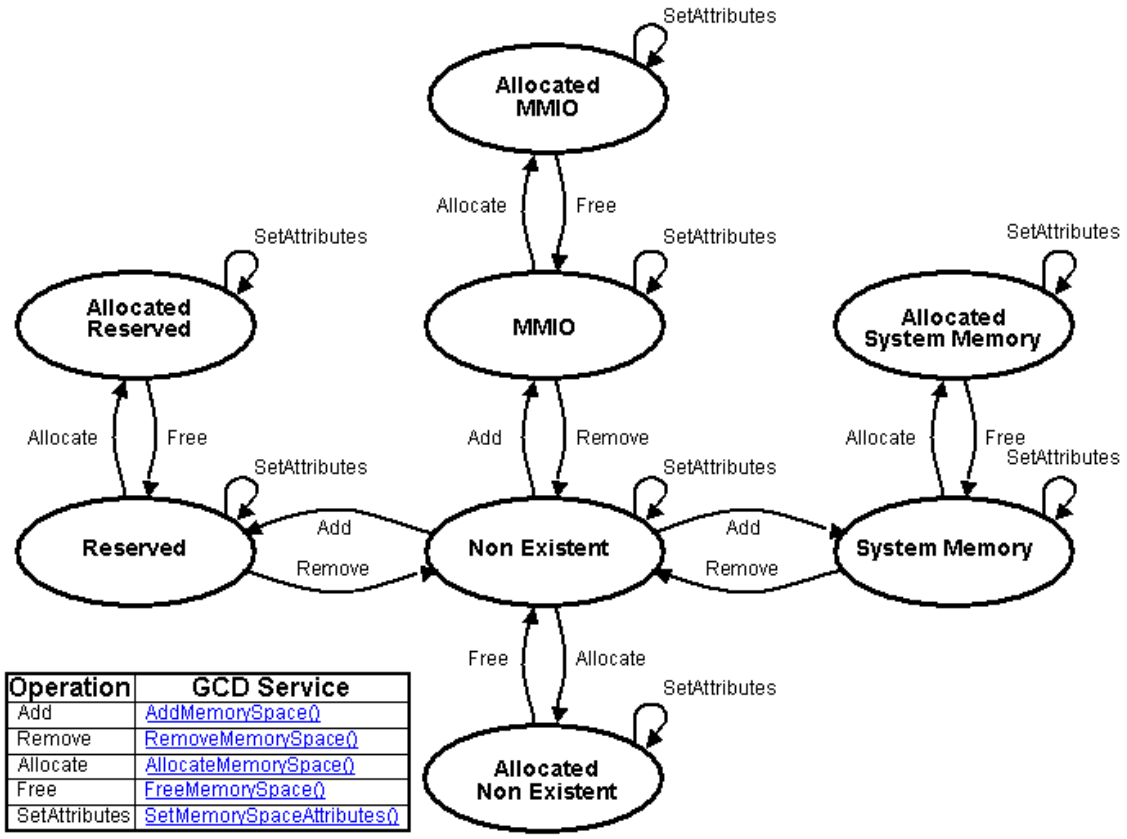


**Figure 7.1. GCD Memory State Transitions**

## GCD I/O Resources

The Global Coherency Domain (GCD) Services used to manage I/O resources include the following:

- **AddIoSpace()**
- **AllocateIoSpace()**
- **FreeIoSpace()**
- **RemoveIoSpace()**

The GCD Services used to retrieve the GCD I/O space map include the following:

- **GetIoSpaceDescriptor()**
- **GetIoSpaceMap()**

The GCD I/O space map is initialized from the HOB list that is passed to the entry point of the DXE Foundation. One HOB type describes the number of address lines that are used to access I/O resources. This information is used to initialize the state of the GCD I/O space map. Any I/O regions outside this initial region are not available to any of the GCD Services that are used to manage I/O resources. The GCD I/O space map is designed to describe the I/O address space with

as many as 64 address lines. Each region in the GCD I/O space map can being and end on a byte boundary.

An I/O region described by the GCD I/O space map can be in several different states. These include nonexistent I/O, I/O, and reserved I/O. These I/O regions can be allocated and freed by DXE drivers executing in the DXE environment. The figure below shows the possible state transitions for each byte of I/O in the GCD I/O space map. The transitions are labeled with the GCD Service that can move the byte from one state to another. The GCD Services are required to merge similar I/O regions that are adjacent to each other into a single I/O descriptor, which reduces the number of entries in the GCD I/O space map.

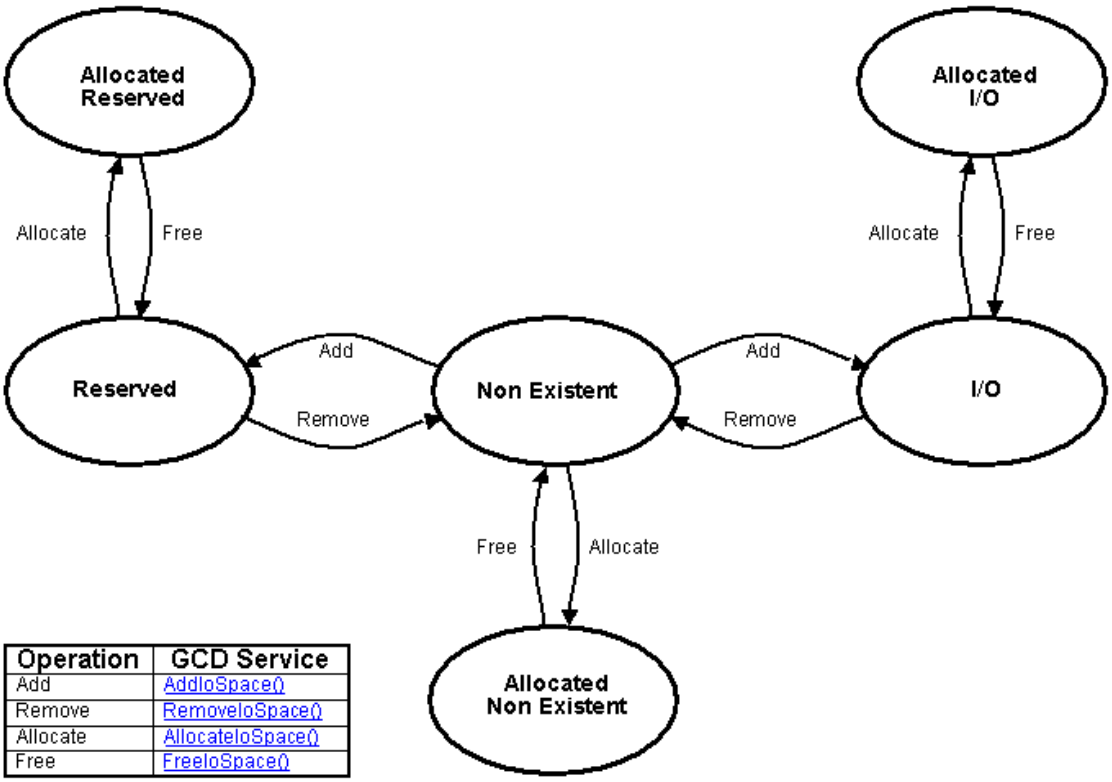


**Figure 7.2. GCD I/O State Transitions**

## Global Coherency Domain Services

The functions that make up Global Coherency Domain (GCD) Services are used during preboot to add, remove, allocate, free, and provide maps of the system memory, memory-mapped I/O, and I/O resources in a platform. These services, used in conjunction with the Memory Allocation Services, provide the ability to manage all the memory and I/O resources in a platform. The table below lists the Global Coherency Domain Services.

**Table 7.1. Global Coherency Domain Services**

| Name | Type | Description |
|---|---|---|
| AddMemorySpace | Boot | This service adds reserved memory, system memory, or memory-mapped I/O resources to the global coherency domain of the processor. |
| AllocateMemorySpace | Boot | This service allocates nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. |
| FreeMemorySpace | Boot | This service frees nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. |
| RemoveMemorySpace | Boot | This service removes reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor. |
| GetMemorySpaceDescriptor | Boot | This service retrieves the descriptor for a memory region containing a specified address. |
| SetMemorySpaceAttributes | Boot | This service modifies the attributes for a memory region in the global coherency domain of the processor. |
| GetMemorySpaceMap | Boot | Returns a map of the memory resources in the global coherency domain of the processor. |
| AddIoSpace | Boot | This service adds reserved I/O, or I/O resources to the global coherency domain of the processor. |
| AllocateIoSpace | Boot | This service allocates nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor. |
| FreeIoSpace | Boot | This service frees nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor. |
| RemoveIoSpace | Boot | This service removes reserved I/O, or I/O resources from the global coherency domain of the processor. |
| GetIoSpaceDescriptor | Boot | This service retrieves the descriptor for an I/O region containing a specified address. |
| GetIoSpaceMap | Boot | Returns a map of the I/O resources in the global coherency domain of the processor. |

# AddMemorySpace()

## Summary

This service adds reserved memory, system memory, or memory-mapped I/O resources to the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
AddMemorySpace (
  IN EFI_GCD_MEMORY_TYPE     GcdMemoryType,
  IN EFI_PHYSICAL_ADDRESS    BaseAddress,
  IN UINT64                  Length,
  IN UINT64                  Capabilities
  );
```

## Parameters

*GcdMemoryType*

The type of memory resource being added. Type **EFI_GCD_MEMORY_TYPE** is defined in "Related Definitions" below. The only types allowed are **EfiGcdMemoryTypeReserved**, **EfiGcdMemoryTypeSystemMemory**, and **EfiGcdMemoryTypeMemoryMappedIo**.

*BaseAddress*

The physical address that is the start address of the memory resource being added. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size, in bytes, of the memory resource that is being added.

*Capabilities*

The bit mask of attributes that the memory resource region supports. The bit mask of available attributes is defined in the **GetMemoryMap()** function description in the *EFI 1.10 Specification.*

## Description

The **AddMemorySpace()** function converts unallocated non-existent memory ranges to a range of reserved memory, a range of system memory, or a range of memory mapped I/O. *BaseAddress* and *Length* specify the memory range, and *GcdMemoryType* specifies the memory type. The bit mask of all supported attributes for the memory range being added is specified by *Capabilities*. If the memory range is successfully added, then **EFI_SUCCESS** is returned.

If the memory range specified by *BaseAddress* and *Length* is of type **EfiGcdMemoryTypeSystemMemory**, then the memory range may be automatically allocated for use by the EFI memory services. If the addition of the memory range specified by

*BaseAddress* and *Length* results in a GCD memory space map containing one or more 4 KB regions of unallocated **EfiGcdMemoryTypeSystemMemory** aligned on 4 KB boundaries, then those regions will always be converted to ranges of allocated **EfiGcdMemoryTypeSystemMemory**. This extra conversion will never be performed for fragments of memory that do not meet the above criteria.

If the GCD memory space map contains adjacent memory regions that only differ in their base address and length fields, then those adjacent memory regions must be merged into a single memory descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If *GcdMemoryType* is not **EfiGcdMemoryTypeReserved**, **EfiGcdMemoryTypeSystemMemory**, or **EfiGcdMemoryTypeMemoryMappedIo**, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the memory range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If any portion of the memory range specified by *BaseAddress* and *Length* is not of type **EfiGcdMemoryTypeNonExistent**, then **EFI_ACCESS_DENIED** is returned.

If any portion of the memory range specified by *BaseAddress* and *Length* was allocated in a prior call to **AllocateMemorySpace()**, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to add the memory resource to the global coherency domain of the processor, then **EFI_OUT_OF_RESOURCES** is returned.

## Related Definitions

```
//*******************************************************
// EFI_GCD_MEMORY_TYPE
//*******************************************************
typedef enum {
  EfiGcdMemoryTypeNonExistent,
  EfiGcdMemoryTypeReserved,
  EfiGcdMemoryTypeSystemMemory,
  EfiGcdMemoryTypeMemoryMappedIo,
  EfiGcdMemoryTypeMaximum
} EFI_GCD_MEMORY_TYPE;
```

**EfiGcdMemoryTypeNonExistent**

A memory region that is visible to the boot processor. However, there are no system components that are currently decoding this memory region.

**EfiGcdMemoryTypeReserved**

A memory region that is visible to the boot processor. This memory region is being decoded by a system component, but the memory region is not considered to be either system memory or memory-mapped I/O.

**EfiGcdMemoryTypeSystemMemory**

A memory region that is visible to the boot processor. A memory controller is currently decoding this memory region and the memory controller is producing a tested system memory region that is available to the memory services.

**EfiGcdMemoryTypeMemoryMappedIo**

A memory region that is visible to the boot processor. This memory region is currently being decoded by a component as memory-mapped I/O that can be used to access I/O devices in the platform.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The memory resource was added to the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *GcdMemoryType* is invalid. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to add the memory resource to the global coherency domain of the processor. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_ACCESS_DENIED | One or more bytes of the memory resource range specified by *BaseAddress* and *Length* conflicts with a memory resource range that was previously added to the global coherency domain of the processor. |
| EFI_ACCESS_DENIED | One or more bytes of the memory resource range specified by *BaseAddress* and *Length* was allocated in a prior call to **AllocateMemorySpace()**. |

## AllocateMemorySpace()

### Summary

This service allocates nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor.

### Prototype

```
EFI_STATUS
AllocateMemorySpace (
  IN     EFI_GCD_ALLOCATE_TYPE    GcdAllocateType,
  IN     EFI_GCD_MEMORY_TYPE      GcdMemoryType,
  IN     UINTN                    Alignment,
  IN     UINT64                   Length,
  IN OUT EFI_PHYSICAL_ADDRESS     *BaseAddress,
  IN     EFI_HANDLE               ImageHandle,
  IN     EFI_HANDLE               DeviceHandle   OPTIONAL
  );
```

### Parameters

*GcdAllocateType*

The type of allocation to perform. Type **EFI_GCD_ALLOCATE_TYPE** is defined in "Related Definitions" below.

*GcdMemoryType*

The type of memory resource being allocated. Type **EFI_GCD_MEMORY_TYPE** is defined in **AddMemorySpace()**. The only types allowed are **EfiGcdMemoryTypeNonExistent**, **EfiGcdMemoryTypeReserved**, **EfiGcdMemoryTypeSystemMemory**, and **EfiGcdMemoryTypeMemoryMappedIo**.

*Alignment*

The log base 2 of the boundary that *BaseAddress* must be aligned on output. For example, a value of 0 means that *BaseAddress* can be aligned on any byte boundary, and a value of 12 means that *BaseAddress* must be aligned on a 4 KB boundary.

*Length*

The size in bytes of the memory resource range that is being allocated.

*BaseAddress*

A pointer to a physical address. On input, the way in which the address is used depends on the value of *Type*. See "Description" below for more information. On output the address is set to the base of the memory resource range that was allocated. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*ImageHandle*

The image handle of the agent that is allocating the memory resource. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*DeviceHandle*

The device handle for which the memory resource is being allocated. If the memory resource is not being allocated for a device that has an associated device handle, then this parameter is optional and may be **NULL**. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Description

The **AllocateMemorySpace()** function searches for a memory range of type *GcdMemoryType* and converts the discovered memory range from the unallocated state to the allocated state. The parameters *GcdAllocateType*, *Alignment*, *Length*, and *BaseAddress* specify the manner in which the GCD memory space map is searched. If a memory range is found that meets the search criteria, then the base address of the memory range is returned in *BaseAddress*, and **EFI_SUCCESS** is returned. *ImageHandle* and *DeviceHandle* are used to convert the memory range from the unallocated state to the allocated state. *ImageHandle* identifies the image that is calling **AllocateMemorySpace()**, and *DeviceHandle* identifies the device that *ImageHandle* is managing that requires the memory range. *DeviceHandle* is optional, because the device that *ImageHandle* is managing might not have an associated device handle. If a memory range meeting the search criteria cannot be found, then **EFI_NOT_FOUND** is returned.

If *GcdAllocateType* is **EfiGcdAllocateAnySearchBottomUp**, then the GCD memory space map is searched from the lowest address up to the highest address looking for unallocated memory ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdMemoryType*.

If *GcdAllocateType* is **EfiGcdAllocateAnySearchTopDown**, then the GCD memory space map is searched from the highest address down to the lowest address looking for unallocated memory ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdMemoryType*.

If *GcdAllocateType* is **EfiGcdAllocateMaxAddressSearchBottomUp**, then the GCD memory space map is searched from the lowest address up to *BaseAddress* looking for unallocated memory ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdMemoryType*.

If *GcdAllocateType* is **EfiGcdAllocateMaxAddressSearchTopDown**, then the GCD memory space map is searched from *BaseAddress* down to the lowest address looking for unallocated memory ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdMemoryType*.

If *GcdAllocateType* is **EfiGcdAllocateAddress**, then the GCD memory space map is checked to see if the memory range starting at *BaseAddress* for *Length* bytes is of type *GcdMemoryType*, unallocated, and begins on a the boundary specified by *Alignment*.

If the GCD memory space map contains adjacent memory regions that only differ in their base address and length fields, then those adjacent memory regions must be merged into a single memory descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If *BaseAddress* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *ImageHandle* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *GcdMemoryType* is not **EfiGcdMemoryTypeNonExistent**, **EfiGcdMemoryTypeReserved**, **EfiGcdMemoryTypeSystem Memory**, or **EfiGcdMemoryTypeMemoryMappedIo**, then **EFI_INVALID_PARAMETER** is returned.

If *GcdAlocateType* is less than zero, or *GcdAllocateType* is greater than or equal to *EfiGcdMaxAllocateType* then **EFI_INVALID_PARAMETER** is returned.

If there are not enough system resources available to allocate the memory range, then **EFI_OUT_OF_RESOURCES** is returned.

## Related Definitions

```
//*******************************************************
// EFI_GCD_ALLOCATE_TYPE
//*******************************************************
typedef enum {
  EfiGcdAllocateAnySearchBottomUp,
  EfiGcdAllocateMaxAddressSearchBottomUp,
  EfiGcdAllocateAddress,
  EfiGcdAllocateAnySearchTopDown,
  EfiGcdAllocateMaxAddressSearchTopDown,
  EfiGcdMaxAllocateType
} EFI_GCD_ALLOCATE_TYPE;
```

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The memory resource was allocated from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *GcdAllocateType* is invalid. |
| EFI_INVALID_PARAMETER | *GcdMemoryType* is invalid. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_INVALID_PARAMETER | *BaseAddress* is **NULL**. |
| EFI_INVALID_PARAMETER | *ImageHandle* is **NULL**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to allocate the memory resource from the global coherency domain of the processor. |
| EFI_NOT_FOUND | The memory resource request could not be satisfied. |

# FreeMemorySpace()

## Summary

This service frees nonexistent memory, reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
FreeMemorySpace (
  IN EFI_PHYSICAL_ADDRESS   BaseAddress,
  IN UINT64                 Length
  );
```

## Parameters

*BaseAddress*

The physical address that is the start address of the memory resource being freed. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the memory resource range that is being freed.

## Description

The **FreeMemorySpace()** function converts the memory range specified by *BaseAddress* and *Length* from the allocated state to the unallocated state. If this conversion is successful, then **EFI_SUCCESS** is returned.

If the GCD memory space map contains adjacent memory regions that only differ in their base address and length fields, then those adjacent memory regions must be merged into a single memory descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the memory range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If one or more bytes of the memory range specified by *BaseAddress* and *Length* were not allocated on previous calls to **AllocateMemorySpace()**, then **EFI_NOT_FOUND** is returned.

If there are not enough system resources available to free the memory range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The memory resource was freed from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_NOT_FOUND | The memory resource range specified by *BaseAddress* and *Length* was not allocated with previous calls to **AllocateMemorySpace()**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to free the memory resource from the global coherency domain of the processor. |

# RemoveMemorySpace()

## Summary

This service removes reserved memory, system memory, or memory-mapped I/O resources from the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
RemoveMemorySpace (
  IN EFI_PHYSICAL_ADDRESS   BaseAddress,
  IN UINT64                 Length
  );
```

## Parameters

*BaseAddress*

The physical address that is the start address of the memory resource being removed. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the memory resource that is being removed.

## Description

The **RemoveMemorySpace()** function converts the memory range specified by *BaseAddress* and *Length* to the memory type **EfiGcdMemoryTypeNonExistent**. If this conversion is successful, then **EFI_SUCCESS** is returned.

If the GCD memory space map contains adjacent memory regions that only differ in their base address and length fields, then those adjacent memory regions must be merged into a single memory descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the memory range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If one or more bytes of the memory range specified by *BaseAddress* and *Length* were not added to the GCD memory space map with previous calls to **AddMemorySpace()**, then **EFI_NOT_FOUND** is returned.

If one or more bytes of the memory range specified by *BaseAddress* and *Length* were allocated from the GCD memory space map with previous calls to **AllocateMemorySpace()**, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to remove the memory range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The memory resource was removed from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_NOT_FOUND | One or more bytes of the memory resource range specified by *BaseAddress* and *Length* was not added with previous calls to **AddMemorySpace()**. |
| EFI_ACCESS_DENIED | One or more bytes of the memory resource range specified by *BaseAddress* and *Length* has been allocated with **AllocateMemorySpace()**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to remove the memory resource from the global coherency domain of the processor. |

# GetMemorySpaceDescriptor()

## Summary

This service retrieves the descriptor for a memory region containing a specified address.

## Prototype

```
EFI_STATUS
GetMemorySpaceDescriptor (
  IN  EFI_PHYSICAL_ADDRESS             BaseAddress,
  OUT EFI_GCD_MEMORY_SPACE_DESCRIPTOR  *Descriptor
  );
```

## Parameters

*BaseAddress*

The physical address that is the start address of a memory region. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Descriptor*

A pointer to a caller allocated descriptor. On return, the descriptor describes the memory region containing *BaseAddress*. Type **EFI_GCD_MEMORY_SPACE_DESCRIPTOR** is defined in "Related Definitions" below.

## Description

The **GetMemorySpaceDescriptor()** function retrieves the descriptor for the memory region that contains the address specified by *BaseAddress*. If a memory region containing *BaseAddress* is found, then the descriptor for that memory region is returned in the caller allocated structure *Descriptor*, and **EFI_SUCCESS** is returned.

If *Descriptor* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If a memory region containing *BaseAddress* is not present in the GCD memory space map, then **EFI_NOT_FOUND** is returned.

## Related Definitions

```
//*******************************************************
// EFI_GCD_MEMORY_SPACE_DESCRIPTOR
//*******************************************************
typedef struct {
  EFI_PHYSICAL_ADDRESS  BaseAddress;
  UINT64                Length;
  UINT64                Capabilities;
  UINT64                Attributes;
  EFI_GCD_MEMORY_TYPE   GcdMemoryType;
  EFI_HANDLE            ImageHandle;
  EFI_HANDLE            DeviceHandle;
} EFI_GCD_MEMORY_SPACE_DESCRIPTOR;
```

*BaseAddress*

The physical address of the first byte in the memory region. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The number of bytes in the memory region.

*Capabilities*

The bit mask of attributes that the memory region is capable of supporting. The bit mask of available attributes is defined in the **GetMemoryMap()** function description in the *EFI 1.10 Specification.*

*Attributes*

The bit mask of attributes that the memory region is currently using. The bit mask of available attributes is defined in **GetMemoryMap()**.

*GcdMemoryType*

Type of the memory region. Type **EFI_GCD_MEMORY_TYPE** is defined in the **AddMemorySpace()** function description.

*ImageHandle*

The image handle of the agent that allocated the memory resource described by *PhysicalStart* and *NumberOfBytes*. If this field is **NULL**, then the memory resource is not currently allocated. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*DeviceHandle*

The device handle for which the memory resource has been allocated. If *ImageHandle* is **NULL**, then the memory resource is not currently allocated. If this field is **NULL**, then the memory resource is not associated with a device that is described by a device handle. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The descriptor for the memory resource region containing *BaseAddress* was returned in *Descriptor*. |
| EFI_INVALID_PARAMETER | *Descriptor* is **NULL**. |
| EFI_NOT_FOUND | A memory resource range containing *BaseAddress* was not found. |

## SetMemorySpaceAttributes()

### Summary

This service modifies the attributes for a memory region in the global coherency domain of the processor.

### Prototype

```
EFI_STATUS
SetMemorySpaceAttributes (
  IN EFI_PHYSICAL_ADDRESS          BaseAddress,
  IN UINT64                        Length,
  IN UINT64                        Attributes
  );
```

### Parameters

*BaseAddress*

The physical address that is the start address of a memory region. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the memory region.

*Attributes*

The bit mask of attributes to set for the memory region. The bit mask of available attributes is defined in the **GetMemoryMap()** function description in the *EFI 1.10 Specification.*

### Description

The **SetMemorySpaceAttributes()** function modifies the attributes for the memory region specified by *BaseAddress* and *Length* from their current attributes to the attributes specified by *Attributes*. If this modification of attributes succeeds, then **EFI_SUCCESS** is returned.

If the GCD memory space map contains adjacent memory regions that only differ in their base address and length fields, then those adjacent memory regions must be merged into a single memory descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the memory range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If the attributes specified by *Attributes* are not supported for the memory region specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned. The *Attributes* bit mask must be a proper subset of the capabilities bit mask for the specified memory region. The capabilities bit mask is specified when a memory region is added with **AddMemorySpace()** and can be retrieved with **GetMemorySpaceDescriptor()** or **GetMemorySpaceMap()**.

If the attributes for one or more bytes of the memory range specified by *BaseAddress* and *Length* cannot be modified because the current system policy does not allow them to be modified, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to modify the attributes of the memory range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The attributes were set for the memory region. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_UNSUPPORTED | The bit mask of attributes is not support for the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_ACCESS_DENIED | The attributes for the memory resource range specified by *BaseAddress* and *Length* cannot be modified. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to modify the attributes of the memory resource range. |

## GetMemorySpaceMap()

### Summary

Returns a map of the memory resources in the global coherency domain of the processor.

### Prototype

```
EFI_STATUS
GetMemorySpaceMap (
  OUT UINTN                                *NumberOfDescriptors,
  OUT EFI_GCD_MEMORY_SPACE_DESCRIPTOR      **MemorySpaceMap
  );
```

### Parameters

*NumberOfDescriptors*

A pointer to number of descriptors returned in the *MemorySpaceMap* buffer. This parameter is ignored on input, and is set to the number of descriptors in the *MemorySpaceMap* buffer on output.

*MemorySpaceMap*

A pointer to the array of **EFI_GCD_MEMORY_SPACE_DESCRIPTOR**s. Type **EFI_GCD_MEMORY_SPACE_DESCRIPTOR** is defined in **GetMemorySpaceDescriptor()**. This buffer is allocated with **AllocatePool()**, so it is the caller's responsibility to free this buffer with a call to **FreePool()**. The number of descriptors in *MemorySpaceMap* is returned in *NumberOfDescriptors*. See the *EFI 1.10 Specification* for definitions of **AllocatePool()** and **FreePool()**.

### Description

The **GetMemorySpaceMap()** function retrieves the entire GCD memory space map. If there are no errors retrieving the GCD memory space map, then the number of descriptors in the GCD memory space map is returned in *NumberOfDescriptors*, the array of descriptors from the GCD memory space map is allocated with **AllocatePool()**, the descriptors are transferred into *MemorySpaceMap*, and **EFI_SUCCESS** is returned.

If *NumberOfDescriptors* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *MemorySpaceMap* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If there are not enough resources to allocate *MemorySpaceMap*, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The memory space map was returned in the *MemorySpaceMap* buffer, and the number of descriptors in *MemorySpaceMap* was returned in *NumberOfDescriptors*. |
| EFI_INVALID_PARAMETER | *NumberOfDescriptors* is **NULL**. |
| EFI_INVALID_PARAMETER | *MemorySpaceMap* is **NULL**. |
| EFI_OUT_OF_RESOURCES | There are not enough resources to allocate *MemorySpaceMap*. |

## AddIoSpace()

### Summary

This service adds reserved I/O, or I/O resources to the global coherency domain of the processor.

### Prototype

```
EFI_STATUS
AddIoSpace (
  IN EFI_GCD_IO_TYPE         GcdIoType,
  IN EFI_PHYSICAL_ADDRESS    BaseAddress,
  IN UINT64                  Length
  );
```

### Parameters

*GcdIoType*

The type of I/O resource being added. Type **EFI_GCD_IO_TYPE** is defined in "Related Definitions" below. The only types allowed are **EfiGcdIoTypeReserved** and **EfiGcdIoTypeIo**.

*BaseAddress*

The physical address that is the start address of the I/O resource being added. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the I/O resource that is being added.

### Description

The **AddIoSpace()** function converts unallocated non-existent I/O ranges to a range of reserved I/O, or a range of I/O. *BaseAddress* and *Length* specify the I/O range, and *GcdIoType* specifies the I/O type. If the I/O range is successfully added, then **EFI_SUCCESS** is returned.

If the GCD I/O space map contains adjacent I/O regions that only differ in their base address and length fields, then those adjacent I/O regions must be merged into a single I/O descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If *GcdIoType* is not **EfiGcdIoTypeReserved** or **EfiGcdIoTypeIo**, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the I/O range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If any portion of the I/O range specified by *BaseAddress* and *Length* is not of type **EfiGcdIoTypeNonExistent**, then **EFI_ACCESS_DENIED** is returned.

If any portion of the I/O range specified by *BaseAddress* and *Length* was allocated in a prior call to **AllocateIoSpace()**, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to add the I/O resource to the global coherency domain of the processor, then **EFI_OUT_OF_RESOURCES** is returned.

## Related Definitions

```
//*******************************************************
// EFI_GCD_IO_TYPE
//*******************************************************
typedef enum {
  EfiGcdIoTypeNonExistent,
  EfiGcdIoTypeReserved,
  EfiGcdIoTypeIo,
  EfiGcdIoTypeMaximum
} EFI_GCD_IO_TYPE;
```

**EfiGcdIoTypeNonExistent**

An I/O region that is visible to the boot processor. However, there are no system components that are currently decoding this I/O region.

**EfiGcdIoTypeReserved**

An I/O region that is visible to the boot processor. This I/O region is currently being decoded by a system component, but the I/O region cannot be used to access I/O devices.

**EfiGcdIoTypeIo**

An I/O region that is visible to the boot processor. This I/O region is currently being decoded by a system component that is producing I/O ports that can be used to access I/O devices.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The I/O resource was added to the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *GcdIoType* is invalid. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to add the I/O resource to the global coherency domain of the processor. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the I/O resource range specified by *BaseAddress* and *Length*. |
| EFI_ACCESS_DENIED | One or more bytes of the I/O resource range specified by *BaseAddress* and *Length* conflicts with an I/O resource range that was previously added to the global coherency domain of the processor. |
| EFI_ACCESS_DENIED | One or more bytes of the I/O resource range specified by *BaseAddress* and *Length* was allocated in a prior call to **AllocateIoSpace()**. |

## AllocateIoSpace()

### Summary

This service allocates nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor.

### Prototype

```
EFI_STATUS
AllocateIoSpace (
  IN     EFI_GCD_ALLOCATE_TYPE   AllocateType,
  IN     EFI_GCD_IO_TYPE         GcdIoType,
  IN     UINTN                   Alignment,
  IN     UINT64                  Length,
  IN OUT EFI_PHYSICAL_ADDRESS    *BaseAddress,
  IN     EFI_HANDLE              ImageHandle,
  IN     EFI_HANDLE              DeviceHandle   OPTIONAL
  );
```

### Parameters

*GcdAllocateType*

The type of allocation to perform. Type **EFI_GCD_ALLOCATE_TYPE** is defined in **AllocateMemorySpace()**.

*GcdIoType*

The type of I/O resource being allocated. Type **EFI_GCD_IO_TYPE** is defined in **AddIoSpace()**. The only types allowed are **EfiGcdIoTypeNonExistent**, **EfiGcdIoTypeReserved**, and **EfiGcdIoTypeIo**.

*Alignment*

The log base 2 of the boundary that *BaseAddress* must be aligned on output. For example, a value of 0 means that *BaseAddress* can be aligned on any byte boundary, and a value of 12 means that *BaseAddress* must be aligned on a 4 KB boundary.

*Length*

The size in bytes of the I/O resource range that is being allocated.

*BaseAddress*

A pointer to a physical address. On input, the way in which the address is used depends on the value of *Type*. See "Description" below for more information. On output the address is set to the base of the I/O resource range that was allocated. Type **EFI_PHYSICAL_ADDRESS** is defined in **AllocatePages()** in the *EFI 1.10 Specification.*

*ImageHandle*

The image handle of the agent that is allocating the I/O resource. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*DeviceHandle*

The device handle for which the I/O resource is being allocated. If the I/O resource is not being allocated for a device that has an associated device handle, then this parameter is optional and may be **NULL**. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Description

The **AllocateIoSpace()** function searches for an I/O range of type *GcdIoType* and converts the discovered I/O range from the unallocated state to the allocated state. The parameters *GcdAllocateType*, *Alignment*, *Length*, and *BaseAddress* specify the manner in which the GCD I/O space map is searched. If an I/O range is found that meets the search criteria, then the base address of the I/O range is returned in *BaseAddress*, and **EFI_SUCCESS** is returned. *ImageHandle* and *DeviceHandle* are used to convert the I/O range from the unallocated state to the allocated state. *ImageHandle* identifies the image that is calling **AllocateIoSpace()**, and *DeviceHandle* identifies the device that *ImageHandle* is managing that requires the I/O range. *DeviceHandle* is optional, because the device that *ImageHandle* is managing might not have an associated device handle. If an I/O range meeting the search criteria cannot be found, then **EFI_NOT_FOUND** is returned.

If *GcdAllocateType* is **EfiGcdAllocateAnySearchBottomUp**, then the GCD I/O space map is searched from the lowest address up to the highest address looking for unallocated I/O ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdIoType*.

If *GcdAllocateType* is **EfiGcdAllocateAnySearchTopDown**, then the GCD I/O space map is searched from the highest address down to the lowest address looking for unallocated I/O ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdIoType*.

If *GcdAllocateType* is **EfiGcdAllocateMaxAddressSearchBottomUp**, then the GCD I/O space map is searched from the lowest address up to *BaseAddress* looking for unallocated I/O ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdIoType*.

If *GcdAllocateType* is **EfiGcdAllocateMaxAddressSearchTopDown**, then the GCD I/O space map is searched from *BaseAddress* down to the lowest address looking for unallocated I/O ranges of *Length* bytes beginning on a boundary specified by *Alignment* that matches *GcdIoType*.

If *GcdAllocateType* is **EfiGcdAllocateAddress**, then the GCD I/O space map is checked to see if the I/O range starting at *BaseAddress* for *Length* bytes is of type *GcdIoType*, unallocated, and begins on a the boundary specified by *Alignment*.

If the GCD I/O space map contains adjacent I/O regions that only differ in their base address and length fields, then those adjacent I/O regions must be merged into a single I/O descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If *BaseAddress* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *ImageHandle* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *GcdIoType* is not **EfiGcdIoTypeNonExistent**, **EfiGcdIoTypeReserved**, or **EfiGcdIoTypeIo**, then **EFI_INVALID_PARAMETER** is returned.

If *GcdAlocateType* is less than zero, or *GcdAllocateType* is greater than or equal to *EfiGcdMaxAllocateType* then **EFI_INVALID_PARAMETER** is returned.

If there are not enough system resources available to allocate the I/O range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The I/O resource was allocated from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *GcdAllocateType* is invalid. |
| EFI_INVALID_PARAMETER | *GcdIoType* is invalid. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_INVALID_PARAMETER | *BaseAddress* is **NULL**. |
| EFI_INVALID_PARAMETER | *ImageHandle* is **NULL**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to allocate the I/O resource from the global coherency domain of the processor. |
| EFI_NOT_FOUND | The I/O resource request could not be satisfied. |

# FreeIoSpace()

## Summary

This service frees nonexistent I/O, reserved I/O, or I/O resources from the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
FreeIoSpace (
  IN EFI_PHYSICAL_ADDRESS   BaseAddress,
  IN UINT64                 Length
  );
```

## Parameters

*BaseAddress*

The physical address that is the start address of the I/O resource being freed. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the I/O resource range that is being freed.

## Description

The **FreeIoSpace()** function converts the I/O range specified by *BaseAddress* and *Length* from the allocated state to the unallocated state. If this conversion is successful, then **EFI_SUCCESS** is returned.

If the GCD I/O space map contains adjacent I/O regions that only differ in their base address and length fields, then those adjacent I/O regions must be merged into a single I/O descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the I/O range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If one or more bytes of the I/O range specified by *BaseAddress* and *Length* were not allocated on previous calls to **AllocateIoSpace()**, then **EFI_NOT_FOUND** is returned.

If there are not enough system resources available to free the I/O range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The I/O resource was freed from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the I/O resource range specified by *BaseAddress* and *Length*. |
| EFI_NOT_FOUND | The I/O resource range specified by *BaseAddress* and *Length* was not allocated with previous calls to **AllocateIoSpace()**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to free the I/O resource from the global coherency domain of the processor. |

# RemoveIoSpace()

## Summary

This service removes reserved I/O, or I/O resources from the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
RemoveIoSpace (
  IN EFI_PHYSICAL_ADDRESS   BaseAddress,
  IN UINT64                 Length
  );
```

## Parameters

*BaseAddress*

A pointer to a physical address that is the start address of the I/O resource being removed. Type **EFI_PHYSICAL_ADDRESS** is defined in **AllocatePages()** in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the I/O resource that is being removed.

## Description

The **RemoveIoSpace()** function converts the I/O range specified by *BaseAddress* and *Length* to the I/O type **EfiGcdIoTypeNonExistent**. If this conversion is successful, then **EFI_SUCCESS** is returned.

If the GCD I/O space map contains adjacent I/O regions that only differ in their base address and length fields, then those adjacent I/O regions must be merged into a single I/O descriptor.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the I/O range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If one or more bytes of the I/O range specified by *BaseAddress* and *Length* were not added to the GCD I/O space map with previous calls to **AddIoSpace()**, then **EFI_NOT_FOUND** is returned.

If one or more bytes of the I/O range specified by *BaseAddress* and *Length* were allocated from the GCD I/O space map with previous calls to **AllocateIoSpace()**, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to remove the I/O range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The I/O resource was removed from the global coherency domain of the processor. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the I/O resource range specified by *BaseAddress* and *Length*. |
| EFI_NOT_FOUND | One or more bytes of the I/O resource range specified by *BaseAddress* and *Length* was not added with previous calls to **AddIoSpace()**. |
| EFI_ACCESS_DENIED | One or more bytes of the I/O resource range specified by *BaseAddress* and *Length* has been allocated with **AllocateIoSpace()**. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to remove the I/O resource from the global coherency domain of the processor. |

# GetIoSpaceDescriptor()

## Summary

This service retrieves the descriptor for an I/O region containing a specified address.

## Prototype

```
EFI_STATUS
GetIoSpaceDescriptor (
  IN  EFI_PHYSICAL_ADDRESS          BaseAddress,
  OUT EFI_GCD_IO_SPACE_DESCRIPTOR   *Descriptor
  );
```

## Parameters

*BaseAddress*

The physical address that is the start address of an I/O region. Type **EFI_PHYSICAL_ADDRESS** is defined in **AllocatePages()** in the *EFI 1.10 Specification.*

*Descriptor*

A pointer to a caller allocated descriptor. On return, the descriptor describes the I/O region containing *BaseAddress*. Type **EFI_GCD_IO_SPACE_DESCRIPTOR** is defined in "Related Definitions" below.

## Description

The **GetIoSpaceDescriptor()** function retrieves the descriptor for the I/O region that contains the address specified by *BaseAddress*. If an I/O region containing *BaseAddress* is found, then the descriptor for that I/O region is returned in the caller allocated structure *Descriptor*, and **EFI_SUCCESS** is returned.

If *Descriptor* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If an I/O region containing *BaseAddress* is not present in the GCD I/O space map, then **EFI_NOT_FOUND** is returned.

## Related Definitions

```
//*******************************************************
// EFI_GCD_IO_SPACE_DESCRIPTOR
//*******************************************************
typedef struct {
  EFI_PHYSICAL_ADDRESS  BaseAddress;
  UINT64                Length;
  EFI_GCD_IO_TYPE       GcdIoType;
  EFI_HANDLE            ImageHandle;
  EFI_HANDLE            DeviceHandle;
} EFI_GCD_IO_SPACE_DESCRIPTOR;
```

*BaseAddress*

Physical address of the first byte in the I/O region. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

Number of bytes in the I/O region.

*GcdIoType*

Type of the I/O region. Type **EFI_GCD_IO_TYPE** is defined in the **AddIoSpace()** function description.

*ImageHandle*

The image handle of the agent that allocated the I/O resource described by *PhysicalStart* and *NumberOfBytes*. If this field is **NULL**, then the I/O resource is not currently allocated. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*DeviceHandle*

The device handle for which the I/O resource has been allocated. If *ImageHandle* is **NULL**, then the I/O resource is not currently allocated. If this field is **NULL**, then the I/O resource is not associated with a device that is described by a device handle. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The descriptor for the I/O resource region containing *BaseAddress* was returned in *Descriptor*. |
| EFI_INVALID_PARAMETER | *Descriptor* is **NULL**. |
| EFI_NOT_FOUND | An I/O resource range containing *BaseAddress* was not found. |

# GetIoSpaceMap()

## Summary

Returns a map of the I/O resources in the global coherency domain of the processor.

## Prototype

```
EFI_STATUS
GetIoSpaceMap (
  OUT UINTN                           *NumberOfDescriptors,
  OUT EFI_GCD_IO_SPACE_DESCRIPTOR     **IoSpaceMap
  );
```

## Parameters

*NumberOfDescriptors*

A pointer to number of descriptors returned in the *IoSpaceMap* buffer. This parameter is ignored on input, and is set to the number of descriptors in the *IoSpaceMap* buffer on output.

*IoSpaceMap*

A pointer to the array of **EFI_GCD_IO_SPACE_DESCRIPTOR**s. Type **EFI_GCD_IO_SPACE_DESCRIPTOR** is defined in **GetIoSpaceDescriptor()**. This buffer is allocated with **AllocatePool()**, so it is the caller's responsibility to free this buffer with a call to **FreePool()**. The number of descriptors in *IoSpaceMap* is returned in *NumberOfDescriptors*.

## Description

The **GetIoSpaceMap()** function retrieves the entire GCD I/O space map. If there are no errors retrieving the GCD I/O space map, then the number of descriptors in the GCD I/O space map is returned in *NumberOfDescriptors*, the array of descriptors from the GCD I/O space map is allocated with **AllocatePool()**, the descriptors are transferred into *IoSpaceMap*, and **EFI_SUCCESS** is returned.

If *NumberOfDescriptors* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If *IoSpaceMap* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If there are not enough resources to allocate *IoSpaceMap*, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The I/O space map was returned in the *IoSpaceMap* buffer, and the number of descriptors in *IoSpaceMap* was returned in *NumberOfDescriptors*. |
| EFI_INVALID_PARAMETER | *NumberOfDescriptors* is **NULL**. |
| EFI_INVALID_PARAMETER | *IoSpaceMap* is **NULL**. |
| EFI_OUT_OF_RESOURCES | There are not enough resources to allocate *IoSpaceMap*. |

# Dispatcher Services

## Dispatcher Services

The functions that make up the Dispatcher Services are used during preboot to schedule drivers for execution. A driver may optionally have the Schedule On Request (SOR) flag set in the driver's dependency expression. Drivers with this bit set will not be loaded and invoked until they are explicitly requested to do so. Files loaded from firmware volumes may be placed in the untrusted state by the Security Architectural Protocol. The services in this section provide this ability to clear the SOR flag in a DXE driver's dependency expression and the ability to promote a file from a firmware volume from the untrusted to the trusted state. The table below lists the Dispatcher Services.

**Table 7.2. Dispatcher Services**

| Name | Type | Description |
|---|---|---|
| Dispatch | Boot | Loads and executed DXE drivers from firmware volumes. |
| Schedule | Boot | Clears the Schedule on Request (SOR) flag for a component that is stored in a firmware volume. |
| Trust | Boot | Changes the state of a file stored in a firmware volume from the untrusted state to the trusted state. |
| ProcessFirmwareVolume | Boot | Creates a firmware volume handle for a firmware volume that is present in system memory. |

# Dispatch()

## Summary

Loads and executes DXE drivers from firmware volumes.

## Prototype

```
EFI_STATUS
Dispatch (
  VOID
  );
```

## Description

The **Dispatch()** function searches for DXE drivers in firmware volumes that have been installed since the last time the **Dispatch()** service was called. It then evaluates the dependency expressions of all the DXE drivers and loads and executes those DXE drivers whose dependency expression evaluate to **TRUE**. This service must interact with the Security Architectural Protocol to authenticate DXE drivers before they are executed. This process is continued until no more DXE drivers can be executed. If one or more DXE drivers are executed, then **EFI_SUCCESS** is returned. If no DXE drivers are executed, **EFI_NOT_FOUND** is returned.

If an attempt is made to invoke the DXE Dispatcher recursively, then no action is performed by the **Dispatch()** service, and **EFI_ALREADY_STARTED** is returned. In this case, because the DXE Dispatcher is already running, it is not necessary to invoke it again. All the DXE drivers that can be dispatched will be dispatched.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | One or more DXE driver were dispatched. |
| EFI_NOT_FOUND | No DXE drivers were dispatched. |
| EFI_ALREADY_STARTED | An attempt is being made to start the DXE Dispatcher recursively. Thus no action was taken. |

# Schedule()

## Summary

Clears the Schedule on Request (SOR) flag for a component that is stored in a firmware volume.

## Prototype

```
EFI_STATUS
Schedule (
  IN EFI_HANLDLE  FirmwareVolumeHandle,
  IN EFI_GUID     *FileName
  );
```

## Parameters

*FirmwareVolumeHandle*

The handle of the firmware volume that contains the file specified by *FileName*. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*FileName*

A pointer to the name of the file in a firmware volume. This is the file that should have its SOR bit cleared. Type **EFI_GUID** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Description

The **Schedule()** function searches the dispatcher queues for the driver specified by *FirmwareVolumeHandle* and *FileName*. If this driver cannot be found, then **EFI_NOT_FOUND** is returned. If the driver is found, and its Schedule On Request (SOR) flag is not set in its dependency expression, then **EFI_NOT_FOUND** is returned. If the driver is found, and its SOR bit is set in its dependency expression, then the SOR flag is cleared, and **EFI_SUCCESS** is returned. After the SOR flag is cleared, the driver will be dispatched if the remaining portions of its dependency expression are satisfied. This service does not automatically invoke the DXE Dispatcher. Instead, the **Dispatch()** service must be used to invoke the DXE Dispatcher.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The DXE driver was found and its SOR bit was cleared. |
| EFI_NOT_FOUND | The DXE driver does not exist, or the DXE driver exists and its SOR bit is not set. |

## Trust()

### Summary

Promotes a file stored in a firmware volume from the untrusted to the trusted state. Only the Security Architectural Protocol can place a file in the untrusted state. A platform specific component may choose to use this service to promote a previously untrusted file to the trusted state.

### Prototype

```
EFI_STATUS
Trust (
  IN EFI_HANLDLE  FirmwareVolumeHandle,
  IN EFI_GUID     *FileName
  );
```

### Parameters

*FirmwareVolumeHandle*

The handle of the firmware volume that contains the file specified by *FileName*. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

*FileName*

A pointer to the name of the file in a firmware volume. This is the file that should be promoted from the untrusted state to the trusted state. Type **EFI_GUID** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

### Description

The **Trust()** function promotes the file specified by *FirmwareVolumeHandle* and *FileName* from the untrusted state to the trusted state. If this file is not found in the queue of untrusted files, then **EFI_NOT_FOUND** is returned. If the driver is found, and its state is changed to trusted and **EFI_SUCCESS** is returned. This service does not automatically invoke the DXE Dispatcher. Instead, the **Dispatch()** service must be used to invoke the DXE Dispatcher.

### Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The file was found in the untrusted state, and it was promoted to the trusted state. |
| EFI_NOT_FOUND | The file was not found in the untrusted state. |

# ProcessFirmwareVolume()

## Summary

Creates a firmware volume handle for a firmware volume that is present in system memory.

## Prototype

```
typedef
EFI_STATUS
ProcessFirmwareVolume (
  IN  VOID        *FirmwareVolumeHeader,
  IN  UINTN       Size,
  OUT EFI_HANDLE  *FirmwareVolumeHandle
  );
```

## Parameters

*FirmwareVolumeHeader*

A pointer to the header of the firmware volume.

*Size*

The size, in bytes, of the firmware volume.

*FirmwareVolumeHandle*

On output, a pointer to the created handle. This service will install the **EFI_FIRMWARE_VOLUME_PROTOCOL** and **EFI_DEVICE_PATH_PROTOCOL** for the of the firmware volume that is described by *FirmwareVolumeHeader* and *Size*. Type **EFI_HANDLE** is defined in **InstallProtocolInterface()** in the *EFI 1.10 Specification.*

## Description

The **ProcessFirmwareVolume()** function examines the contents of the buffer specified by *FirmwareVolumeHeader* and *Size*. If the buffer contains a valid firmware volume, then a new handle is created, and the **EFI_FIRMWARE_VOLUME_PROTOCOL** and a memory-mapped **EFI_DEVICE_PATH_PROTOCOL** are installed onto the new handle. The new handle is returned in *FirmwareVolumeHandle*.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The **EFI_FIRMWARE_VOLUME_PROTOCOL** and **EFI_DEVICE_PATH_PROTOCOL** were installed onto *FirmwareVolumeHandle* for the firmware volume described by *FirmwareVolumeHeader* and *Size*. |
| EFI_VOLUME_CORRUPTED | The firmware volume described by *FirmwareVolumeHeader* and *Size* is corrupted. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources available to produce the **EFI_FIRMWARE_VOLUME_PROTOCOL** and **EFI_DEVICE_PATH_PROTOCOL** for the firmware volume described by *FirmwareVolumeHeader* and *Size*. |

# 8
# Protocols - Device Path Protocol

## Introduction

This book contains the definition of the Device Path Protocol and the information needed to construct and manage device paths in the EFI environment. A device path is constructed and used by the firmware to convey the location of important devices, such as the boot device and console, consistent with the software-visible topology of the system. See the *EFI 1.10 Specification* for details on the Device Path Protocol.

This DXE CIS uses all the device path nodes from the *EFI 1.10 Specification* and adds one device path node type that describes files stored in firmware volumes:

- Firmware Volume File Path Media Device Path

This device path node is used by the updated EFI Boot Service **`LoadImage()`** to load EFI images from firmware volumes. This new capability is used by the DXE Dispatcher to load DXE drivers from firmware volumes.

## Firmware Volume File Path Media Device Path

The table below lists details on the Firmware Volume File Path Media Device Path.

**Table 8.1. Firmware Volume File Path Media Device Path**

| Mnemonic | Byte Offset | Byte Length | Description |
|---|---|---|---|
| Type | 0 | 1 | Type 4 – Media Device Path. |
| Sub-Type | 1 | 1 | Sub-Type 6 – Firmware Volume File Path. |
| Length | 2 | 2 | Length of this structure in bytes. Length is 20 bytes. |
| NameGuid | 4 | 16 | The file name of a file stored in a firmware volume. All file names in firmware volumes are GUIDs. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Specification* for more details. |

# 9
# DXE Foundation

## Introduction

The DXE Foundation is designed to be completely portable with no processor, chipset, or platform dependencies. This lack of dependencies is accomplished by designing in several features:

- The DXE Foundation depends only upon a **HOB list** for its initial state.
  This means that the DXE Foundation does not depend on any services from a previous phase, so all the prior phases can be unloaded once the HOB list is passed to the DXE Foundation.
- The DXE Foundation does not contain any hard-coded addresses.
  This means that the DXE Foundation can be loaded anywhere in physical memory, and it can function correctly no matter where physical memory or where Firmware Volumes (FVs) are located in the processor's physical address space.
- The DXE Foundation does not contain any processor-specific, chipset-specific, or platform-specific information.
  Instead, the DXE Foundation is abstracted from the system hardware through a set of DXE Architectural Protocol interfaces. These architectural protocol interfaces are produced by a set of DXE drivers that are invoked by the DXE Dispatcher.

The DXE Foundation must produce the EFI System Table and its associated set of EFI Boot Services and EFI Runtime Services. The DXE Foundation also contains the DXE Dispatcher whose main purpose is to discover and execute DXE drivers stored in FVs. The execution order of DXE drivers is determined by a combination of the optional *a priori* file and the set of dependency expressions that are associated with the DXE drivers. The FV file format allows dependency expressions to be packaged with the executable DXE driver image. DXE drivers utilize a PE/COFF image format, so the DXE Dispatcher must also contain a PE/COFF loader to load and execute DXE drivers.

## Hand-Off Block (HOB) List

The Hand-Off Block (HOB) list contains all the information that the DXE Foundation requires to produce its memory-based services. The HOB list contains the following:

- Information on the boot mode and the memory that was allocated in the previous phase.
- A description of the system memory that was initialized by the previous phase along with information about the firmware devices that were discovered in the previous phase.

The firmware device information includes the system memory locations of the firmware devices and system memory locations of the firmware volumes that are contained within those firmware devices. The firmware volumes may contain DXE drivers, and the DXE Dispatcher is responsible for loading and executing the DXE drivers that are discovered in those firmware volumes.

- The I/O resources and memory-mapped I/O resources that were discovered in the previous phase.

The HOB list must be treated as a read-only data structure. It conveys the state of the system at the time the DXE Foundation is started. The DXE Foundation and DXE drivers should never modify the contents of the HOB list.

The figure below shows an example HOB list. The first HOB list entry is always the Phase Handoff Information Table (PHIT) HOB that contains the boot mode and a description of the memory regions used by the previous phase. The rest of the HOB list entries can appear in any order. This example shows the various HOB types that are supported. The most important ones to the DXE Foundation are the HOBs that describe system memory and the firmware volumes. A HOB list is terminated by an end of list HOB. There is one additional HOB type that is not shown. This is a GUIDed HOB that allows a module from the previous phase to pass private data to a DXE driver. Only the DXE driver that recognizes the GUID value in the GUIDed HOB will be able to understand the data in the GUIDed HOB. The DXE Foundation does not consume any GUIDed HOBs. The HOB entries are all designed to be position independent. This allows the DXE Foundation to relocate the HOB list to a different location if the DXE Foundation does not like where the previous phase placed the HOB list in memory.

See HOB Translations later in this chapter for more information on HOB types.

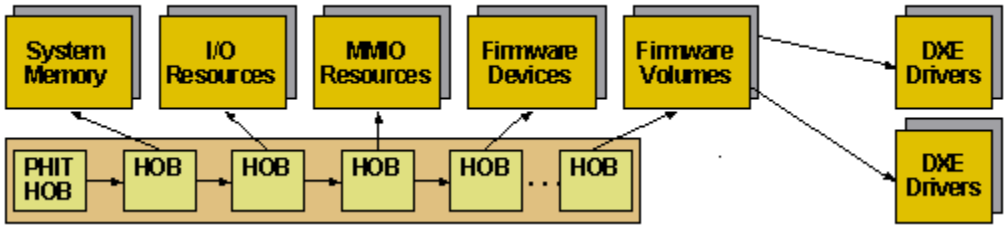


**Figure 9.1. HOB List**

## DXE Foundation Data Structures

The DXE Foundation produces the EFI System Table, and the EFI System Table is consumed by every DXE driver and executable image invoked by the DXE Dispatcher and BDS. It contains all the information required for these components to utilize the services provided by the DXE Foundation and the services provided by any previously loaded DXE driver. The figure below shows the various components that are available through the EFI System Table.

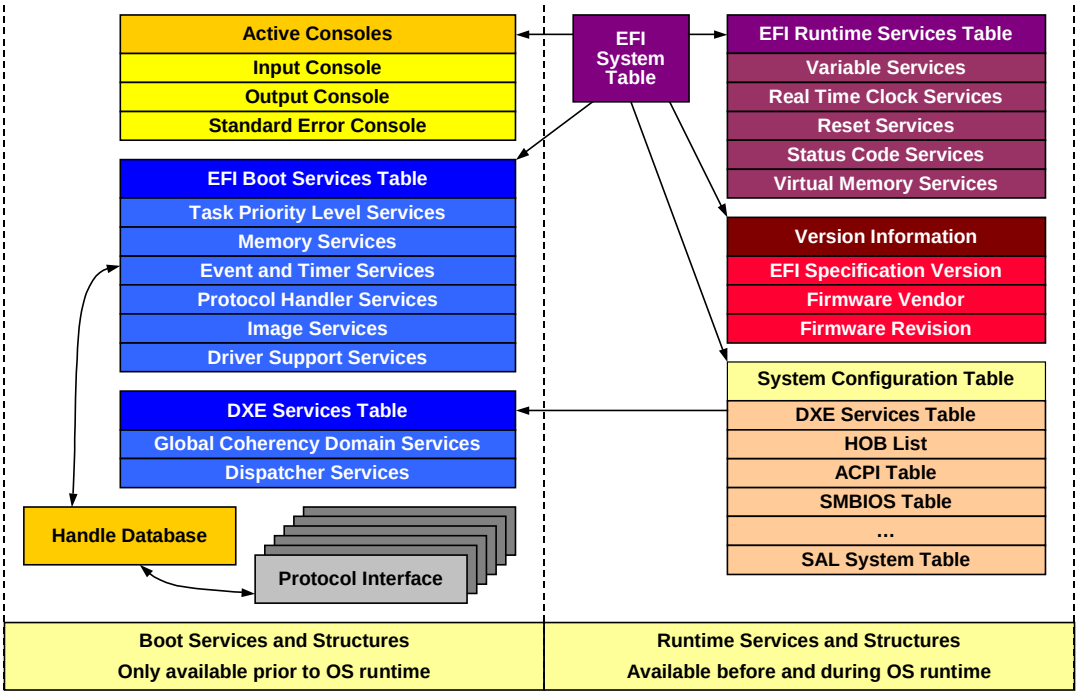


**Figure 9.2. EFI System Table and Related Components**

The DXE Foundation produces the EFI Boot Services, EFI Runtime Services, and DXE Services with the aide of the DXE Architectural Protocols. The EFI System Table also provides access to all the active console devices in the platform and the set of EFI Configuration Tables. The EFI Configuration Tables are an extensible list of tables that describe the configuration of the platform. Today, this includes pointers to tables such as DXE Services, the HOB list, ACPI table, SMBIOS table, and the SAL System Table. This list may be expanded in the future as new table types are defined. Also, through the use of the Protocol Handle Services in the EFI Boot Services Table, any executable image can access the handle database and any of the protocol interfaces that have been registered by DXE drivers.

When the transition to the OS runtime is performed, the handle database, active consoles, EFI Boot Services, DXE Services, and services provided by boot service DXE drivers are terminated. This

frees up memory for use by the OS. This only leaves the EFI System Table, EFI Runtime Services Table, and the EFI Configuration Tables available in the OS runtime environment. There is also the option of converting all of the EFI Runtime Services from a physical address space to an OS-specific virtual address space. This address space conversion may be performed only once.

## Required DXE Foundation Components

The following figure shows the components that a DXE Foundation must contain. A detailed description of these components is listed below.

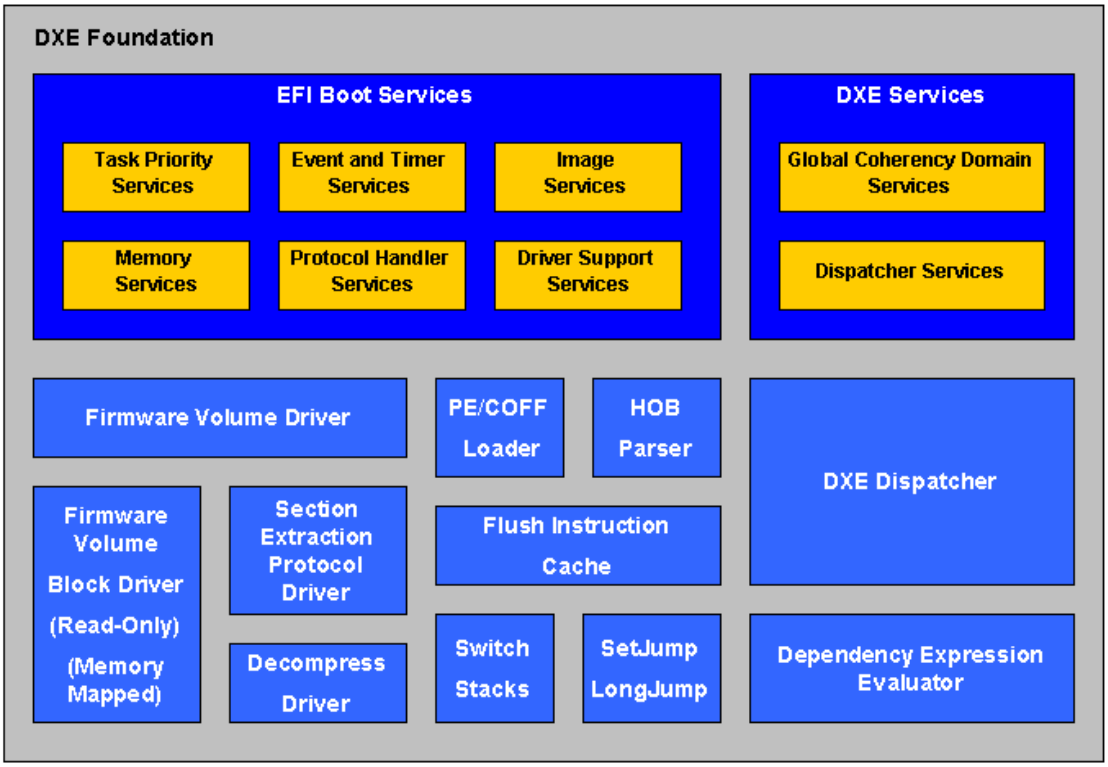


**Figure 9.3. DXE Foundation Components**

A DXE Foundation must have the following components:

- An implementation of the EFI Boot Services. EFI Boot Services Dependencies describes which services can be made available based on the HOB list alone and which services depend on the presence of architectural protocols.
- An implementation of the DXE Services. DXE Services Dependencies describes which services can be made available based on the HOB list alone and which services depend on the presence of architectural protocols.

- A HOB Parser that consumes the HOB list specified by ***HobStart*** and initializes the EFI memory map, GCD memory space map, and GCD I/O space map. See HOB Translations for details on the translation from HOBs to the maps maintained by the DXE Foundation
- An implementation if a DXE Dispatcher that includes a dependency expression evaluator. See DXE Dispatcher for a detailed description of this component.
- A Firmware Volume Block driver that produces the **EFI_FIRMWARE_VOLUME_BLOCK_PROTOCOL** for every firmware volume described in the HOB list. The firmware volumes described in the HOB list are guaranteed to be accessed memory mapped, so this driver only need to support read-only access to a memory mapped firmware volume. Firmware Volume Block drivers with additional capabilities may be produced by DXE drivers. The Firmware Volume driver uses this component so the DXE Dispatcher can search for *a priori* files and DXE drivers in firmware volumes. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Block Specification* for the definition of the Firmware Volume Block Protocol.
- A Firmware Volume driver that produces the **EFI_FIRMWARE_VOLUME_PROTOCOL** for every firmware volume that described in the HOB list. This component is used by the DXE Dispatcher to search for *a priori* files and DXE drivers in firmware volumes. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Specification* for the definition of the Firmware Volume Protocol.
- An instance of the **EFI_SECTION_EXTRACTION_PROTOCOL** with support for all section types except GUIDed section types. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Specification* for the detailed requirements for this component. This component is required to extract dependency expression sections and PE32 sections from DXE driver files stored in firmware volumes.
- An instance of the **EFI_DECOMPRESS_PROTOCOL**. See the *EFI 1.10 Specification* for the detailed requirements for this component. This component is required by the **EFI_SECTION_EXTRACTION_PROTOCOL** to read compressed sections from DXE drivers stored in firmware volumes. It is expected that most DXE drivers will utilize compressed sections to reduce the size of firmware volumes.
- The DXE Dispatcher uses the Boot Service **StartImage()** to invoke a DXE driver. The Boot Services **StartImage()** and **Exit()** work together to hand control to a DXE driver and return control to the DXE Foundation. Since the Boot Service **Exit()** can be called for anywhere inside a DXE driver, the Boot Service **Exit()** is required to rebalance the stack, so it is in the same state it was in when the Boot Service **Start()** was called. This is typically implemented using the processor-specific functions called **SetJump()** and **LongJump()**. Since the DXE Foundation must use the Boot Services **StartImage()** and **Exit()** to invoke DXE drivers, the routines **SetJump()** and **LongJump()** are required by the DXE Foundation.
- A PE/COFF loader that supports PE32+ image types. This PE/COFF loader is used to implement the EFI Boot Service **LoadImage()**. The DXE Dispatcher uses the Boot Service **LoadImage()** to load DXE drivers into system memory. If the processor that the DXE Foundation is compiled for requires an instruction cache when an image is loaded into system memory, then an instruction cache flush routine is also required in the DXE Foundation.
- The phase that executed prior to DXE will initialize a stack for the DXE Foundation to use. This stack is described in the HOB list. If the size of this stack does not meet the DXE

Foundation's minimum stack size requirement or the stack is not located in memory region that is suitable to the DXE Foundation, then the DXE Foundation will have to allocate a new stack that does meet the minimum size and location requirements. As a result, the DXE Foundation must contain a stack switching routine for the processor type that the DXE Foundation is compiled.

## Handing Control to DXE Dispatcher

The DXE Foundation must complete the following tasks before handing control to the DXE Dispatcher. The order that these tasks are performed is implementation dependent.

- Use the HOB list to initialize the GCD memory space map, the GCD I/O space map, and EFI memory map.
- Allocate the EFI Boot Services Table from **EFI_BOOT_SERVICES_MEMORY** and initialize the services that only require system memory to function correctly. The remaining EFI Boot Services must be filled in with a service that returns **EFI_NOT_AVAILABLE_YET**.
- Allocate the DXE Services Table from **EFI_BOOT_SERVICES_MEMORY** and initialize the services that only require system memory to function correctly. The remaining DXE Services must be filled in with a service that returns **EFI_NOT_AVAILABLE_YET**.
- Allocate the EFI Runtime Services Table from **EFI_RUNTIME_SERVICES_MEMORY** and initialize all the services to a service that returns **EFI_NOT_AVAILABLE_YET**.
- Allocate the EFI System Table from **EFI_RUNTIME_SERVICES_MEMORY** and initialize all the fields.
- Build an image handle and **EFI_LOADED_IMAGE_PROTOCOL** instance for the DXE Foundation itself and add it to the handle database.
- If the HOB list is not in a suitable location in memory, then relocate the HOB list to a more suitable location.
- Add the DXE Services Table to the EFI Configuration Table.
- Add the HOB list to the EFI Configuration Table.
- Create a notification event for each of the DXE Architectural Protocols. These events will be signaled when a DXE driver installs a DXE Architectural Protocol in the handle database. The DXE Foundation must have a notification function associated with each of these events, so the full complement of EFI Boot Services, EFI Runtime Services, and DXE Services can be produced. Each of the notification functions should compute the 32-bit CRC of the EFI Boot Services Table, EFI Runtime Services Table, and the DXE Services Table if the **CalculateCrc32()** Boot Services is available.
- Initialize the Section Extraction Protocol driver that must be built into the DXE Foundation.
- Initialize the Decompress Protocol driver that must be built into the DXE Foundation.
- Produce firmware volume handles for the one or more firmware volumes that are described in the HOB list.

Once these tasks have been completed, the DXE Foundation is ready to load and execute DXE drivers stored in firmware volumes. This execution is done by handing control to the DXE Dispatcher. Once the DXE Dispatcher has finished dispatching all the DXE drivers that it can, control is then passed to the BDS Architectural Protocol. If for some reason, any of the DXE Architectural Protocols have not been produced by the DXE drivers, then the system is in an unusable state and the DXE Foundation must halt. Otherwise, control is handed to the BDS Architectural Protocol. The BDS Architectural Protocol is responsible for transferring control to an operating system or system utility.

# DXE Foundation Entry Point

## DXE Foundation Entry Point

The only parameter passed to the DXE Foundation is a pointer to the HOB list. The DXE Foundation and all the DXE drivers must treat the HOB list as read-only data.

The function **DXE_ENTRY_POINT** is the main entry point to the DXE Foundation.

# DXE_ENTRY_POINT

## Summary

This function is the main entry point to the DXE Foundation.

## Prototype

```
typedef
VOID
(EFIAPI *DXE_ENTRY_POINT) (
  IN  VOID  *HobStart
  );
```

## Parameters

*HobStart*

A pointer to the HOB list.

## Description

This function is the entry point to the DXE Foundation. The PEI phase, which executes just before DXE, is responsible for loading and invoking the DXE Foundation in system memory. The only parameter that is passed to the DXE Foundation is *HobStart*. This parameter is a pointer to the HOB list that describes the system state at the hand-off to the DXE Foundation. At a minimum, this system state must include the following:

- PHIT HOB
- CPU HOB
- Description of system memory
- Description of one or more firmware volumes

The DXE Foundation is also guaranteed that only one processor is running and that the processor is running with interrupts disabled. The implementation of the DXE Foundation must not make any assumptions about where the DXE Foundation will be loaded or where the stack is located. In general, the DXE Foundation should make as few assumptions about the state of the system as possible. This lack of assumptions will allow the DXE Foundation to be portable to the widest variety of system architectures.

# Dependencies

## EFI Boot Services Table

## EFI Boot Services Dependencies

The table below lists all the EFI Boot Services and the components upon which each of these services depend. The topics that follow describe what responsibilities the DXE Foundation has in producing the services that depend on the presence of DXE Architectural Protocols.

**Table 9.1. Boot Service Dependencies**

| Name | Dependency |
|---|---|
| CreateEvent | HOB list |
| CloseEvent | HOB list |
| SignalEvent | HOB list |
| WaitForEvent | HOB list |
| CheckEvent | HOB list |
| SetTimer | Timer Architectural Protocol |
| RaiseTPL | CPU Architectural Protocol |
| RestoreTPL | CPU Architectural Protocol |
| AllocatePages | HOB list |
| FreePages | HOB list |
| GetMemoryMap | HOB list |
| AllocatePool | HOB list |
| FreePool | HOB list |
| InstallProtocolInterface | HOB list |
| UninstallProtocolInterface | HOB list |
| ReinstallProtocolInterface | HOB list |
| RegisterProtocolNotify | HOB list |
| LocateHandle | HOB list |
| HandleProtocol | HOB list |
| LocateDevicePath | HOB list |
| OpenProtocol | HOB list |
| CloseProtocol | HOB list |
| OpenProtocolInformation | HOB list |
| ConnectController | HOB list |
| DisconnectController | HOB list |
| ProtocolsPerHandle | HOB list |
| LocateHandleBuffer | HOB list |
| LocateProtocol | HOB list |
| InstallMultipleProtocolInterfaces | HOB list |

| Name | Dependency |
|---|---|
| UninstallMultipleProtocolInterfaces | HOB list |
| LoadImage | HOB list |
| StartImage | HOB list |
| UnloadImage | HOB list |
| EFI_IMAGE_ENTRY_POINT | HOB list |
| Exit | HOB list |
| ExitBootServices | HOB list |
| SetWatchDogTimer | Watchdog Architectural Protocol |
| Stall | Metronome Architectural Protocol<br>Timer Architectural Protocol |
| CopyMem | HOB list |
| SetMem | HOB list |
| GetNextMonotonicCount | Monotonic Counter Architectural Protocol |
| InstallConfigurationTable | HOB list |
| CalculateCrc32 | Runtime Architectural Protocol |

## SetTimer()

When the DXE Foundation is notified that the **EFI_TIMER_ARCH_PROTOCOL** has been installed, then the Boot Service **SetTimer()** can be made available. The DXE Foundation can use the services of the **EFI_TIMER_ARCH_PROTOCOL** to initialize and hook a heartbeat timer interrupt for the DXE Foundation. The DXE Foundation can use this heartbeat timer interrupt to determine when to signal on-shot and periodic timer events. This service may be called before the **EFI_TIMER_ARCH_PROTOCOL** is installed. However, since a heartbeat timer is not running yet, time is essentially frozen at zero. This means that no periodic or one-shot timer events will fire until the **EFI_TIMER_ARCH_PROTOCOL** is installed.

## RaiseTPL()

The DXE Foundation must produce the Boot Service **RaiseTPL()** when the memory-based services are initialized. The DXE Foundation is guaranteed to be handed control of the platform with interrupts disabled. Until the DXE Foundation installs a heartbeat timer interrupt and turns on interrupts, this Boot Service can be a very simple function that always succeeds. When the DXE Foundation is notified that the **EFI_CPU_ARCH_PROTOCOL** has been installed, then the full version of the Boot Service **RaiseTPL()** can be made available. When an attempt is made to raise the TPL level to **EFI_TPL_HIGH_LEVEL** or higher, then the DXE Foundation should use the services of the **EFI_CPU_ARCH_PROTOCOL** to disable interrupts.

## RestoreTPL()

The DXE Foundation must produce the Boot Service **RestoreTPL()** when the memory-based services are initialized. The DXE Foundation is guaranteed to be handed control of the platform with interrupts disabled. Until the DXE Foundation installs a heartbeat timer interrupt and turns on interrupts, this Boot Service can be a very simple function that always succeeds. When the DXE Foundation is notified that the **EFI_CPU_ARCH_PROTOCOL** has been installed, then the full version of the Boot Service **RestoreTPL()** can be made available. When an attempt is made to restore the TPL level to level below **EFI_TPL_HIGH_LEVEL**, then the DXE Foundation should use the services of the **EFI_CPU_ARCH_PROTOCOL** to enable interrupts.

## SetWatchdogTimer()

When the DXE Foundation is notified that the **EFI_WATCHDOG_ARCH_PROTOCOL** has been installed, then the Boot Service **SetWatchdogTimer()** can be made available. The DXE Foundation can use the services of the **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL** to set the amount of time before the system's watchdog timer will expire.

## Stall()

When the DXE Foundation is notified that the **EFI_METRONOME_ARCH_PROTOCOL** has been installed, the DXE Foundation can produce a very simple version of the Boot Service **Stall()**. The granularity of the Boot Service **Stall()** will be based on the period of the **EFI_METRONOME_ARCH_PROTOCOL**.

When the DXE Foundation is notified that the **EFI_TIMER_ARCH_PROTOCOL** has been installed, the DXE Foundation can possibly produce a more accurate version of the Boot Service **Stall()**. This all depends on the periods of the **EFI_METRONOME_ARCH_PROTOCOL** and the period of the **EFI_TIMER_ARCH_PROTOCOL**. The DXE Foundation should produce the Boot Service **Stall()** using the most accurate time base available.

## GetNextMonotonicCount()

When the DXE Foundation is notified that the **EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL** has been installed, then the Boot Service **GetNextMonotonicCount()** is available. The DXE driver that produces the **EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL** is responsible for directly updating the *GetNextMonotonicCount* field of the EFI Boot Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Boot Services Table.

## CalculateCrc32()

When the DXE Foundation is notified that the **EFI_RUNTIME_ARCH_PROTOCOL** has been installed, then the Boot Service **CalculateCrc32()** is available. The DXE driver that produces the **EFI_RUNTIME_ARCH_PROTOCOL** is responsible for directly updating the *CalculateCrc32* field of the EFI Boot Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Boot Services Table.

# EFI Runtime Services Table

## EFI Runtime Services Dependencies

The table below lists all the EFI Runtime Services and the components upon which each of these services depend. The topics that follow describe what responsibilities the DXE Foundation has in producing the services that depend on the presence of DXE Architectural Protocols.

**Table 9.2. Runtime Service Dependencies**

| Name | Dependency |
|---|---|
| GetVariable | Variable Architectural Protocol |
| GetNextVariableName | Variable Architectural Protocol |
| SetVariable | Variable Architectural Protocol / Variable Write Architectural Protocol |
| GetTime | Real Time Clock Architectural Protocol |
| SetTime | Real Time Clock Architectural Protocol |
| GetWakeupTime | Real Time Clock Architectural Protocol |
| SetWakeupTime | Real Time Clock Architectural Protocol |
| SetVirtualAddressMap | Runtime Architectural Protocol |
| ConvertPointer | Runtime Architectural Protocol |
| ResetSystem | Reset Architectural Protocol |
| GetNextHighMonotonicCount | Monotonic Counter Architectural Protocol |
| ReportStatusCode | Status Code Architectural Protocol |

## GetVariable()

When the DXE Foundation is notified that the **EFI_VARIABLE_ARCH_PROTOCOL** has been installed, then the Runtime Service **GetVariable()** is available. The DXE driver that produces the **EFI_VARIABLE_ARCH_PROTOCOL** is responsible for directly updating the *GetVariable* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## GetNextVariableName()

When the DXE Foundation is notified that the **EFI_VARIABLE_ARCH_PROTOCOL** has been installed, then the Runtime Service **GetNextVariableName()** is available. The DXE driver that produces the **EFI_VARIABLE_ARCH_PROTOCOL** is responsible for directly updating the *GetNextVariableName* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## SetVariable()

When the DXE Foundation is notified that the **EFI_VARIABLE_ARCH_PROTOCOL** has been installed, then the Runtime Service **SetVariable()** is available. The DXE driver that produces the **EFI_VARIABLE_ARCH_PROTOCOL** is responsible for directly updating the *SetVariable* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the

32-bit CRC of the EFI Runtime Services Table. The **EFI_VARIABLE_ARCH_PROTOCOL** is required to provide read-only access to all environment variables and write access to volatile environment variables.

When the DXE Foundation is notified that the **EFI_VARIABLE_WRITE_ARCH_PROTOCOL** has been installed, then write access to nonvolatile environment variables will also be available. If an attempt is made to call this function for a nonvolatile environment variable prior to the installation of **EFI_VARIABLE_WRITE_ARCH_PROTOCOL**, then **EFI_NOT_AVAILABLE_YET** must be returned. This allows for flexibility in the design and implementation of the variables services in a platform such that read access to environment variables can be provided very early in the DXE phase and write access to nonvolatile environment variables can be provided later in the DXE phase.

## GetTime()

When the DXE Foundation is notified that the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** has been installed, then the Runtime Service **GetTime()** is available. The DXE driver that produces the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** is responsible for directly updating the *GetTime* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## SetTime()

When the DXE Foundation is notified that the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** has been installed, then the Runtime Service **SetTime()** is available. The DXE driver that produces the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** is responsible for directly updating the *SetTime* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## GetWakeupTime()

When the DXE Foundation is notified that the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** has been installed, then the Runtime Service **GetWakeupTime()** is available. The DXE driver that produces the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** is responsible for directly updating the *GetWakeupTime* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## SetWakeupTime()

When the DXE Foundation is notified that the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** has been installed, then the Runtime Service **SetWakeupTime()** is available. The DXE driver that produces the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL** is responsible for directly updating the *SetWakeupTime* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## SetVirtualAddressMap()

When the DXE Foundation is notified that the **EFI_RUNTIME_ARCH_PROTOCOL** has been installed, then the Runtime Service **SetVirtualAddressMap()** is available. The DXE driver

that produces the **EFI_RUNTIME_ARCH_PROTOCOL** is responsible for directly updating the *SetVirtualAddressMap* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

### ConvertPointer()

When the DXE Foundation is notified that the **EFI_RUNTIME_ARCH_PROTOCOL** has been installed, then the Runtime Service **ConvertPointer()** is available. The DXE driver that produces the **EFI_RUNTIME_ARCH_PROTOCOL** is responsible for directly updating the *ConvertPointer* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

### ResetSystem()

When the DXE Foundation is notified that the **EFI_RESET_ARCH_PROTOCOL** has been installed, then the Runtime Service **ResetSystem()** is available. The DXE driver that produces the **EFI_RESET_ARCH_PROTOCOL** is responsible for directly updating the *Reset* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

### GetNextHighMonotonicCount()

When the DXE Foundation is notified that the **EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL** has been installed, then the Runtime Service **GetNextHighMonotonicCount()** is available. The DXE driver that produces the **EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL** is responsible for directly updating the *GetNextHighMonotonicCount* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

### ReportStatusCode()

When the DXE Foundation is notified that the **EFI_STATUS_CODE_ARCH_PROTOCOL** has been installed, then the Runtime Service **ReportStatusCode()** is available. The DXE driver that produces the **EFI_STATUS_CODE_ARCH_PROTOCOL** is responsible for directly updating the *ReportStatusCode* field of the EFI Runtime Services Table. The DXE Foundation is only responsible for updating the 32-bit CRC of the EFI Runtime Services Table.

## DXE Services Table

## DXE Services Dependencies

The table below lists all the DXE Services and the components upon which each of these services depend. The topics that follow describe what responsibilities the DXE Foundation has in producing the services that depend on the presence of DXE Architectural Protocols.

**Table 9.3. DXE Service Dependencies**

| Name | Dependency |
|---|---|
| AddMemorySpace | HOB list |
| AllocateMemorySpace | HOB list |
| FreeMemorySpace | HOB list |
| RemoveMemorySpace | HOB list |
| GetMemorySpaceDescriptor | CPU Architectural Protocol |
| SetMemorySpaceAttributes | CPU Architectural Protocol |
| GetMemorySpaceMap | CPU Architectural Protocol |
| AddIoSpace | HOB list |
| AllocateIoSpace | HOB list |
| FreeIoSpace | HOB list |
| RemoveIoSpace | HOB list |
| GetIoSpaceDescriptor | HOB list |
| GetIoSpaceMap | HOB list |
| Schedule | HOB list |

## GetMemorySpaceDescriptor()

When the DXE Foundation is notified that the **EFI_CPU_ARCH_PROTOCOL** has been installed, then the DXE Service **GetMemorySpaceDescriptor()** is fully functional. This function is made available when the memory-based services are initialized. However, the *Attributes* field of the **EFI_GCD_MEMORY_SPACE_DESCRIPTOR** is not valid until the **EFI_CPU_ARCH_PROTOCOL** is installed.

## SetMemorySpaceAttributes()

When the DXE Foundation is notified that the **EFI_CPU_ARCH_PROTOCOL** has been installed, then the DXE Service **SetMemorySpaceAttributes()** can be made available. The DXE Foundation can then use the **SetMemoryAttributes()** service of the **EFI_CPU_ARCH_PROTOCOL** to implement the DXE Service **SetMemorySpaceAttributes()**.

## GetMemorySpaceMap()

When the DXE Foundation is notified that the **EFI_CPU_ARCH_PROTOCOL** has been installed, then the DXE Service **GetMemorySpaceMap()** is fully functional. This function is made available when the memory-based services are initialized. However, the *Attributes* field of the array of **EFI_GCD_MEMORY_SPACE_DESCRIPTOR**s is not valid until the **EFI_CPU_ARCH_PROTOCOL** is installed.

# HOB Translations

## HOB Translations Overview

The following topics describe how the DXE Foundation should interpret the contents of the HOB list to initialize the GCD memory space map, GCD I/O space map, and EFI memory map. After all of the HOBs have been parsed, the Boot Service **`GetMemoryMap()`** and the DXE Services **`GetMemorySpaceMap()`** and **`GetIoSpaceMap()`** should reflect the memory resources, I/O resources, and logical memory allocations described in the HOB list.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for detailed information on HOBs.

## PHIT HOB

The Phase Handoff Information Table (PHIT) HOB describes a region of tested system memory. This region of memory contains the following:

- HOB list
- Some amount of free memory
- Potentially some logical memory allocations

The PHIT HOB is used by the DXE Foundation to determine the size of the HOB list so that the DXE Foundation can relocate the HOB list to a new location in system memory. The base address of the HOB list is passed to the DXE Foundation in the parameter *`HobStart`*, and the PHIT HOB field *`EfiFreeMemoryBottom`* specifies the end of the HOB list.

Since the PHIT HOB may contain some of amount of free memory, the DXE Foundation may use this free memory region in its early initialization phase until the full complement of EFI memory services are available.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for the definition of this HOB type.

## CPU HOB

The CPU HOB contains the field *`SizeOfMemorySpaceMap`* . This field is used to initialize the GCD memory space map. The *`SizeOfMemorySpaceMap`* field defines the number of address bits that the processor can use to address memory resources. The DXE Foundation must create the primordial GCD memory space map entry of type **`EfiGcdMemoryTypeNonExistent`** for the region from 0 to **`(1 <<`** *`SizeOfMemorySpaceMap`***`)`**. All future GCD memory space operations must be performed within this memory region.

The CPU HOB also contains the field *`SizeOfIoSpaceMap`* . This field is used to initialize the GCD I/O space map. The *`SizeOfIoSpaceMap`* field defines the number of address bits that the processor can use to address I/O resources. The DXE Foundation must create the primordial GCD I/O space map entry of type **`EfiGcdIoTypeNonExistent`** for the region from 0 to **`(1 <<`** *`SizeOfIoSpaceMap`***`)`**. All future GCD I/O space operations must be performed within this I/O region.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for the definition of this HOB type.

## Resource Descriptor HOBs

The DXE Foundation must traverse the HOB list looking for Resource Descriptor HOBs. These HOBs describe memory and I/O resources that are visible to the processor. All of the resource ranges described in these HOBs must fall in the memory and I/O ranges initialized in the GCD maps based on the contents of the CPU HOB. The DXE Foundation will use the DXE Services **AddMemorySpace()** and **AddIoSpace()** to register these memory and I/O resources in the GCD maps.

The *Owner* field of the Resource Descriptor HOB is ignored by the DXE Foundation. The *ResourceType* field and *ResourceAttribute* fields are used to determine the GCD memory type or GCD I/O type of the resource. The table below shows this mapping. The resource range is specified by the *PhysicalStart* and *ResourceLength* fields of the Resource Descriptor HOB.

The *ResourceAttribute* field also contains the caching capabilities of memory regions. If a memory region is being added to the GCD memory space map, then the *ResourceAttribute* field will be used to initialize the supported caching capabilities. The *ResourceAttribute* field is also be used to further qualify memory regions. For example, a system memory region cannot be added to the EFI memory map if it is read protected. However, it is legal to add a firmware device memory region that is write-protected if the firmware device is a ROM.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for the definition of this HOB type.

**Table 9.4. Resource Descriptor HOB to GCD Type Mapping**

| Resource Descriptor HOB | | GCD Map | |
|---|---|---|---|
| Resource Type | Attributes | Memory Type | I/O Type |
| System Memory | Present | Reserved | |
| System Memory | Present AND Initialized | Reserved | |
| System Memory | Present AND Initialized AND Tested | System Memory | |
| Memory-Mapped I/O | | Memory Mapped I/O | |
| Firmware Device | | Memory Mapped I/O | |
| Memory-Mapped I/O Port | | Reserved | |
| Memory Reserved | | Reserved | |
| I/O | | | I/O |
| I/O Reserved | | | Reserved |

## Firmware Volume HOBs

The DXE Foundation must traverse the HOB list for Firmware Volume HOBs. When the DXE Foundation discovers a Firmware Volume HOB, a new handle must be created in the handle database, and the **EFI_FIRMWARE_VOLUME_BLOCK_PROTOCOL** and the **EFI_FIRMWARE_VOLUME_PROTOCOL** must be installed on that handle. The *BaseAddress* and *Length* fields of the Firmware Volume HOB specific the memory range that the firmware

volume consumes. The DXE Service **AllocateMemorySpace()** is used to allocate the memory regions described in the Firmware Volume HOBs to the DXE Foundation. The EFI Boot Service **InstallProtocolInterface()** is used to create new handles and install protocol interfaces.

See the following specifications for code definitions:

- Firmware Volume HOB type: *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification*
- Firmware Volume Block Protocol: *Intel® Platform Innovation Framework for EFI Firmware Volume Block Specification*
- Firmware Volume Protocol: *Intel® Platform Innovation Framework for EFI Firmware Volume Specification*

## Memory Allocation HOBs

Memory Allocation HOBs describe logical memory allocations that occurred prior to the DXE phase. The DXE Foundation must parse the HOB list for this HOB type. When a HOB of this type is discovered, the GCD memory space map must be updated with a call to the DXE Service **AllocateMemorySpace()**. In addition, the EFI memory map must be updated with logical allocation described by the *MemoryType*, *MemoryBaseAddress*, and *MemoryLength* fields of the Memory Allocation HOB.

Once the DXE Foundation has parsed all of the Memory Allocation HOBs, all of the unallocated system memory regions in the GCD memory space map must be allocated to the DXE Foundation with the DXE Service **AllocateMemorySpace()**. In addition, those same memory regions must be added to the EFI memory map so those memory regions can be allocated and freed using the Boot Services **AllocatePages()**, **AllocatePool()**, **FreePages()**, and **FreePool()**.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for the definition of this HOB type.

## GUID Extension HOBs

The DXE Foundation does not consume any GUID Extension HOBs. The HOB parser in the DXE Foundation will skip HOBs of this type. GUID Extension HOBs contain private data that is being passed from the previous execution phase to a specific DXE driver. DXE drivers may choose to parse the HOB list for GUID Extension HOBs.

See the *Intel® Platform Innovation Framework for EFI Hand-Off Block (HOB) Specification* for the definition of this HOB type.

# 10
# DXE Dispatcher

## Introduction

After the DXE Foundation is initialized, control is handed to the DXE Dispatcher. The DXE Dispatcher examines every firmware volume that is present in the system. Firmware volumes are either declared by HOBs, or they are declared by DXE drivers. For the DXE Dispatcher to run, at least one firmware volume must be declared by a HOB.

The DXE Dispatcher is responsible for loading and invoking DXE drivers found in firmware volumes. Some DXE drivers may depend on the services produced by other DXE drivers, so the DXE Dispatcher is also required to execute the DXE drivers in the correct order. The DXE drivers may also be produced by a variety of different vendors, so the DXE drivers must describe the services they depend upon. The DXE dispatcher must evaluate these dependencies to determine a valid order to execute the DXE drivers. Some vendors may wish to specify a fixed execution order for some or all of the DXE drivers in a firmware volume, so the DXE dispatcher must support this requirement.

In addition, the DXE Dispatcher must support the ability to load “emergency patch” drivers. These drivers would be added to the firmware volume to address an issue that was not know at the time the original firmware was built. These DXE drivers would be loaded just before or just after an existing DXE driver.

Finally, the DXE Dispatcher must be flexible enough to support a variety of platform specific security policies for loading and executing DXE drivers from firmware volumes. Some platforms may choose to run DXE drivers with no security checks, and others may choose to check the validity of a firmware volume before it is used, and other may choose to check the validity of every DXE driver in a firmware volume before it is executed.

## Requirements

The DXE Dispatcher must meet the following requirement:

- **Support fixed execution order of DXE drivers.** This fixed execution order is specified in an *a priori* file in the firmware volume.
- **Determine DXE driver execution order based on each driver's dependencies.** A DXE driver that is stored in a firmware volume may optionally contain a dependency expression section. This section specifies the protocols that the DXE driver requires to execute.
- **Support "emergency patch" DXE drivers.** The dependency expressions are flexible enough to describe the protocols that a DXE drivers may require. In addition, the dependency expression can declare that the DXE driver is to be loaded and executed immediately before or immediately after a different DXE driver.
- **Support platform specific security policies for DXE driver execution.** The DXE Dispatcher is required to use the services of the Security Architecture Protocol every time a firmware volume is discovered and every time a DXE driver is loaded.

When a new firmware volume is discovered, it is first authenticated with the Security Architectural Protocol. The Security Architectural Protocol provides the platform-specific policy for validating all firmware volumes. Then, a search is made for the *a priori* file. The *a priori* file has a fixed file name, and it contains the list of DXE drivers that should be loaded and executed first. There can be at most one *a priori* file per firmware volume, and it is legal to have zero *a priori* files in a firmware volume. Once the DXE drivers from the *a priori* file have been loaded and executed, the dependency expressions of the remaining DXE drivers in the firmware volumes are evaluated to determine the order that they will be loaded and executed. The *a priori* file provides a strongly ordered list of DXE drivers that are not required to use dependency expressions. The dependency expressions provide a weakly ordered execution of the remaining DXE drivers. Before each DXE driver is executed, it must be authenticated through the Security Architectural Protocol. The Security Architectural Protocol provides the platform-specific policy for validating all DXE drivers.

Control is transferred from the DXE Dispatcher to the BDS Architectural Protocol after the DXE drivers in the *a priori* file and all the DXE drivers whose dependency expressions evaluate to **TRUE** have been loaded and executed. The BDS Architectural Protocol is responsible for establishing the console devices and attempting the boot of operating systems. As the console devices are established and access to boot devices is established, additional firmware volumes may be discovered. If the BDS Architectural Protocol is unable to start a console device or gain access to a boot device, it will reinvoke the DXE Dispatcher. This will allow the DXE Dispatcher to load and execute DXE drivers from firmware volumes that have been discovered since the last time the DXE Dispatcher was invoked. Once the DXE Dispatcher has loaded and executed all the DXE drivers it can, control is once again returned to the BDS Architectural Protocol to continue the OS boot process.

## The a priori File

The *a priori* file is a special file that may be present in a firmware volume. The rule is that there may be at most one *a priori* file per firmware volume present in a platform. The *a priori* file has a known GUID file name, so the DXE Dispatcher can always find the *a priori* file if it is present. Every time the DXE Dispatcher discovers a firmware volume, it first looks for the *a priori* file. The *a priori* file contains the list of DXE drivers from that firmware volume that should be loaded and executed before any other DXE drivers are discovered. The DXE drivers listed in the *a priori* file are executed in the order that they appear. If any of those DXE drivers have an associated dependency expression, then those dependency expressions are ignored. The *a priori* file provides a deterministic execution order of DXE drivers. DXE drivers that are executed solely based on their dependency expression are weakly ordered. This means that the execution order is not completely deterministic between boots or between platforms. There are cases where a deterministic execution order is required. One example would be to list the DXE drivers required to debug the rest of the DXE phase in the *a priori* file. These DXE drivers that provide debug services may have been loaded much later if only their dependency expressions were considered. By loading them earlier, more of the DXE Foundation and DXE drivers can be debugged. Another example is to use the *a priori* file to eliminate the need for dependency expressions. Some embedded platforms may only require a few DXE drivers with a highly deterministic execution order. The *a priori* file can provide this ordering, and none of the DXE drivers would require dependency expressions. The dependency expressions do have some amount of size overhead, so this method may reduce the size of firmware images. The main purpose of the *a priori* file is to provide a greater degree of flexibility in the firmware design of a platform.

See the next topic for the GUID definition of the *a priori* file, which is the file name that is stored in a firmware volume.

The *a priori* file contains the file names of DXE drivers that are stored in the same firmware volume as the *a priori* file. File names in firmware volumes are GUIDs, so the *a priori* file is simply a list of byte-packed values of type **`EFI_GUID`**. Type **`EFI_GUID`** is defined in the *EFI 1.10 Specification*. The DXE Dispatcher reads the list of **`EFI_GUID`**s from the *a priori* file. Each **`EFI_GUID`** is used to load and execute the DXE driver with that GUID file name. If the DXE driver specified by the GUID file name is not found in the firmware volume, then the file is skipped. If the *a priori* file is not en even multiple of **`EFI_GUID`**s in length, then the DXE driver specified by the last **`EFI_GUID`** in the *a priori* file is skipped.

After all of the DXE drivers listed in the *a priori* file have been loaded and executed, the DXE Dispatcher searches the firmware volume for any additional DXE drivers and executed them according to their dependency expressions.

### EFI_APRIORI_GUID

The following GUID definition is the file name of the *a priori* file that is stored in a firmware volume. This file must be of type **EFI_FV_FILETYPE_FREEFORM** and must contain a single section of type **EFI_SECTION_RAW**. See the following specifications for details on firmware volumes, firmware file types, and firmware file section types:

- *Intel® Platform Innovation Framework for EFI Firmware File System Specification*
- *Intel® Platform Innovation Framework for EFI Firmware Volume Specification*

#### GUID

```
#define EFI_APRIORI_GUID  \

{0xfc510ee7,0xffdc,0x11d4,0xbd,0x41,0x0,0x80,0xc7,0x3c,0x88,0x81}
```

## Dependency Expressions

### Dependency Expressions Overview

A DXE driver is stored in a firmware volume as a file with one or more sections. One of the sections must be a PE32+ image. If a DXE driver has a dependency expression, then it is stored in a *dependency section.* A DXE driver may contain additional sections for compression and security wrappers. The DXE Dispatcher can identify the DXE drivers by their file type. In addition, the DXE Dispatcher can look up the dependency expression for a DXE driver by looking for a dependency section in a DXE driver file. The dependency section contains a section header followed by the actual dependency expression that is composed of a packed byte stream of opcodes and operands.

Dependency expressions stored in dependency sections are designed to be small to conserve space. In addition, they are designed to be simple and quick to evaluate to reduce execution overhead. These two goals are met by designing a small, stack based, instruction set to encode the dependency expressions. The DXE Dispatcher must implement an interpreter for this instruction set in order to evaluate dependency expressions. The instruction set is defined in the following topics.

See Dependency Expression Grammar for an example BNF grammar for a dependency expression compiler. There are many possible methods of specifying the dependency expression for a DXE driver. Dependency Expression Grammar demonstrates one possible design for a tool that can be used to help build DXE driver images.

### Dependency Expression Instruction Set

The following topics describe each of the dependency expression opcodes in detail. Information includes a description of the instruction functionality, binary encoding, and any limitations or unique behaviors of the instruction.

Several of the opcodes require a GUID operand. The GUID operand is a 16-byte value that matches the type **EFI_GUID** that is described in the *EFI 1.10 Specification*. These GUIDs represent protocols that are produced by DXE drivers and the file names of DXE drivers stored in firmware volumes. A dependency expression is a packed byte stream of opcodes and operands. As

a result, some of the GUID operands will not be aligned on natural boundaries. Care must be taken on processor architectures that do allow unaligned accesses.

The dependency expression is stored in a packed byte stream using postfix notation. As a dependency expression is evaluated, the operands are pushed onto a stack. Operands are popped off the stack to perform an operation. After the last operation is performed, the value on the top of the stack represents the evaluation of the entire dependency expression. If a push operation causes a stack overflow, then the entire dependency expression evaluates to **FALSE**. If a pop operation causes a stack underflow, then the entire dependency expression evaluates to **FALSE**. Reasonable implementations of a dependency expression evaluator should not make arbitrary assumptions about the maximum stack size it will support. Instead, it should be designed to grow the dependency expression stack as required. In addition, DXE drivers that contain dependency expressions should make an effort to keep their dependency expressions as small as possible to help reduce the size of the DXE driver.

All opcodes are 8-bit values, and if an invalid opcode is encountered, then the entire dependency expression evaluates to **FALSE**.

If an END opcode is not present in a dependency expression, then the entire dependency expression evaluates to **FALSE**.

If an instruction encoding extends beyond the end of the dependency section, then the entire dependency expression evaluates to **FALSE**.

The final evaluation of the dependency expression results in either a **TRUE** or **FALSE** result.

The table below is a summary of the opcodes that are used to build dependency expressions. The following topics describe each of these instructions in detail.

**Table 10.1. Dependency Expression Opcode Summary**

| Opcode | Description |
|---|---|
| 0x00 | BEFORE <File Name GUID> |
| 0x01 | AFTER <File Name GUID> |
| 0x02 | PUSH <Protocol GUID> |
| 0x03 | AND |
| 0x04 | OR |
| 0x05 | NOT |
| 0x06 | TRUE |
| 0x07 | FALSE |
| 0x08 | END |
| 0x09 | SOR |

# BEFORE

## SYNTAX:

BEFORE <File Name GUID>

## DESCRIPTION:

This opcode tells the DXE Dispatcher that the DXE driver that is associated with this dependency expression must be dispatched just before the DXE driver with the file name specified by GUID. This means that as soon as the dependency expression for the DXE driver specified by GUID evaluates to **TRUE**, then this DXE driver must be placed in the dispatch queue just before the DXE driver with the file name specified by GUID.

## OPERATION:

None.

The following table defines the BEFORE instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x00 |
| 1..16 | A 16-byte GUID that represents the file name of a different DXE driver. The format is the same at type **EFI_GUID**. |

## BEHAVIORS AND RESTRICTIONS:

If this opcode is present in a dependency expression, it must be the first and only opcode in the expression. If is appears in any other location in the dependency expression, then the dependency expression is evaluated to **FALSE**.

## AFTER

### SYNTAX:

AFTER <File Name GUID>

### DESCRIPTION:

This opcode tells the DXE Dispatcher that the DXE driver that is associated with this dependency expression must be dispatched just after the DXE driver with the file name specified by GUID. This means that as soon as the dependency expression for the DXE driver specified by GUID evaluates to **TRUE**, then this DXE driver must be placed in the dispatch queue just after the DXE Driver with the file name specified by GUID.

### OPERATION:

None.

The following table defines the AFTER instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x01 |
| 1..16 | A 16-byte GUID that represents the file name of a different DXE driver. The format is the same at type **EFI_GUID**. |

### BEHAVIORS AND RESTRICTIONS:

If this opcode is present in a dependency expression, it must be the first and only opcode in the expression. If is appears in any other location in the dependency expression, then the dependency expression is evaluated to **FALSE**.

# PUSH

## SYNTAX:

PUSH <Protocol GUID>

## DESCRIPTION:

Pushes a Boolean value onto the stack. If the GUID is present in the handle database, then a **TRUE** is pushed onto the stack. If the GUID is not present in the handle database, then a **FALSE** is pushed onto the stack. The test for the GUID in the handle database may be performed with the Boot Service **LocateProtocol()**.

## OPERATION:

```
Status = gBS->LocateProtocol (GUID, NULL, &Interface);
if (EFI_ERROR (Status)) {
  PUSH FALSE;
} Else {
  PUSH TRUE;
}
```

The following table defines the PUSH instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x02 |
| 1..16 | A 16-byte GUID that represents a protocol that is produced by a different DXE driver. The format is the same at type **EFI_GUID.** |

## BEHAVIORS AND RESTRICTIONS:

None.

## AND

### SYNTAX:

AND

### DESCRIPTION:

Pops two Boolean operands off the stack, performs a Boolean AND operation between the two operands, and pushes the result back onto the stack.

### OPERATION:

Operand1 <= POP Boolean stack element

Operand2 <= POP Boolean stack element

Result <= Operand1 AND Operand2

PUSH Result

The following table defines the AND instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x03. |

### BEHAVIORS AND RESTRICTIONS:

None.

# OR

## SYNTAX:

OR

## DESCRIPTION:

Pops two Boolean operands off the stack, performs a Boolean OR operation between the two operands, and pushes the result back onto the stack.

## OPERATION:

Operand1 <= POP Boolean stack element

Operand2 <= POP Boolean stack element

Result <= Operand1 OR Operand2

PUSH Result

The following table defines the OR instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x04. |

## BEHAVIORS AND RESTRICTIONS:

None.

## NOT

### SYNTAX:

NOT

### DESCRIPTION:

Pops a Boolean operands off the stack, performs a Boolean NOT operation on the operand, and pushes the result back onto the stack.

### OPERATION:

Operand <= POP Boolean stack element

Result <= NOT Operand1

PUSH Result

The following table defines the NOT instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x05. |

### BEHAVIORS AND RESTRICTIONS:

None.

# TRUE

## SYNTAX:

TRUE

## DESCRIPTION:

Pushes a Boolean **TRUE** onto the stack.

## OPERATION:

PUSH **TRUE**

The following table defines the TRUE instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x06. |

## BEHAVIORS AND RESTRICTIONS:

None.

# FALSE

## SYNTAX:

FALSE

## DESCRIPTION:

Pushes a Boolean **FALSE** onto the stack.

## OPERATION:

PUSH **FALSE**

The following table defines the FALSE instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x07. |

## BEHAVIORS AND RESTRICTIONS:

None.

# END

## SYNTAX:

END

## DESCRIPTION:

Pops the final result of the dependency expression evaluation off the stack and exits the dependency expression evaluator.

## OPERATION:

POP Result

RETURN Result

The following table defines the END instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x08. |

## BEHAVIORS AND RESTRICTIONS:

This opcode must be the last one in a dependency expression.

## SOR

### SYNTAX:

SOR

### DESCRIPTION:

Indicates that the DXE driver is to remain on the Schedule on Request (SOR) queue until the DXE Service **Schedule()** is called for this DXE. The dependency expression evaluator treats this operation like a No Operation (NOP).

### OPERATION:

None.

The following table defines the SOR instruction encoding.

| BYTE | DESCRIPTION |
|---|---|
| 0 | 0x09. |

### BEHAVIORS AND RESTRICTIONS:

- If this instruction is present in a dependency expression, it must be the first instruction in the expression. If it appears in any other location in the dependency expression, then the dependency expression is evaluated to **FALSE**.
- This instruction must be followed by a valid dependency expression. If this instruction is the last instruction or it is followed immediately by an END instruction, then the dependency expression is evaluated to **FALSE**.

## Dependency Expression with No Dependencies

A DXE driver that does not have any dependencies must have a dependency expression that evaluates to **TRUE** with no dependencies on any protocol GUIDs or file name GUIDs. The DXE Dispatcher will queue all the DXE drivers of this type immediately after the *a priori* file has been processed.

The following code example shows the dependency expression for a DXE driver that does not have any dependencies using the BNF grammar listed in Dependency Expression Grammar. This is followed by the 2-byte dependency expression that is encoded using the instruction set described in Dependency Expression Instruction Set.

```
//
// Source
//
TRUE
END

//
// Opcodes, Operands, and Binary Encoding
//
ADDR   BINARY                      MNEMONIC
====   ========================
=============================================
0x00 : 06                          TRUE
0x01 : 08                          END
```

## Empty Dependency Expressions

If a DXE driver file does not contain a dependency section, then the DXE driver has an empty dependency expression. The DXE Foundation must support DXE driver and EFI drivers that conform to the *EFI 1.10 Specification*. These EFI drivers assume that all the EFI Boot Services and EFI Runtime Services are available. If an EFI driver is added to a firmware volume, then the EFI driver will have an empty dependency expression, and it should not be loaded and executed by the DXE Dispatcher until all the EFI Boot Services and EFI Runtime Services are available. The DXE Foundation cannot guarantee that this condition is true until all of the DXE Architectural Protocols have been installed.

From the DXE Dispatcher's perspective, DXE drivers without dependency expressions cannot be loaded until all of the DXE Architectural Protocols have been installed. This is equivalent to an implied dependency expression of all the GUIDs of the architectural protocols ANDed together. This implied dependency expression is shown below. The use of empty dependency expressions may also save space, because DXE drivers that require all the EFI Boot Services and EFI Runtime Services to be present can simply remove the dependency section from the DXE driver file.

The code example below shows the dependency expression that is implied by an empty dependency expression using the BNF grammar listed in Dependency Expression Grammar. It also shows the dependency expression after it has been encoded using the instruction set described in Dependency Expression Instruction Set. This fairly complex dependency expression is encoded into a dependency expression that is 216 bytes long. Typical dependency expressions will contain 2 or 3 terms, so those dependency expressions will typically be less than 60 bytes long.

```
//
// Source
//
EFI_BDS_ARCH_PROTOCOL_GUID                  AND
EFI_CPU_ARCH_PROTOCOL_GUID                  AND
EFI_METRONOME_ARCH_PROTOCOL_GUID            AND
EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL_GUID    AND
EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL_GUID      AND
EFI_RESET_ARCH_PROTOCOL_GUID                AND
EFI_RUNTIME_ARCH_PROTOCOL_GUID              AND
EFI_SECURITY_ARCH_PROTOCOL_GUID             AND
EFI_STATUS_CODE_ARCH_PROTOCOL_GUID          AND
EFI_TIMER_ARCH_PROTOCOL_GUID                AND
EFI_VARIABLE_ARCH_PROTOCOL_GUID             AND
EFI_VARIABLE_WRITE_ARCH_PROTOCOL_GUID       AND
EFI_WATCHDOG_TIMER_ARCH_PROTOCOL_GUID
END

//
// Opcodes, Operands, and Binary Encoding
//
ADDR   BINARY                       MNEMONIC
====   ======================
=============================================
0x00 : 02                           PUSH
0x01 : F6 3F 5E 66 CC 46 d4 11      EFI_BDS_ARCH_PROTOCOL_GUID
       9A 38 00 90 27 3F C1 4D
0x11 : 02                           PUSH
0x12 : B1 CC BA 26 42 6F D4 11      EFI_CPU_ARCH_PROTOCOL_GUID
       BC E7 00 80 C7 3C 88 81
0x22 : 03                           AND
0x23 : 02                           PUSH
0x24 : B2 CC BA 26 42 6F d4 11      EFI_METRONOME_ARCH_PROTOCOL_GUID
       BC E7 00 80 C7 3C 88 81
0x34 : 02                           PUSH
0x35 : 72 70 A9 1D DC BD 30 4B
EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL_GUID
       99 F1 72 A0 B5 6F FF 2A
0x45 : 03                           AND
0x46 : 03                           AND
0x47 : 02                           PUSH
0x48 : 87 AC CF 27 CC 46 d4 11
EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL_GUID
       9A 38 00 90 27 3F C1 4D
```

```
0x58 : 02                               PUSH
0x59 : 88 AC CF 27 CC 46 d4 11          EFI_RESET_ARCH_PROTOCOL_GUID
       9A 38 00 90 27 3F C1 4D
0x69 : 03                               AND
0x6A : 02                               PUSH
0x6B : 53 82 d0 96 83 84 d4 11          EFI_RUNTIME_ARCH_PROTOCOL_GUID
       BC F1 00 80 C7 3C 88 81
0x7B : 02                               PUSH
0x7C : E3 23 64 A4 17 46 f1 49          EFI_SECURITY_ARCH_PROTOCOL_GUID
       B9 FF D1 BF A9 11 58 39
0x8C : 03                               AND
0x8D : 03                               AND
0x8E : 03                               AND
0x8F : 02                               PUSH
0x90 : A3 3E 8E D9 39 6F E4 4B          EFI_STATUS_CODE_ARCH_PROTOCOL_GUID
       82 CE 5A 89 0C CB 2C 95
0xA0 : 02                               PUSH
0xA1 : B3 CC BA 26 42 6F D4 11          EFI_TIMER_ARCH_PROTOCOL_GUID
       BC E7 00 80 C7 3C 88 81
0xB1 : 03                               AND
0xB2 : 02                               PUSH
0xB3 : E2 68 56 1E 81 84 D4 11          EFI_VARIABLE_ARCH_PROTOCOL_GUID
       BC F1 00 80 C7 3C 88 81
0xC3 : 02                               PUSH
0xC4 : 18 F8 41 64 62 63 44 4E          EFI_VARIABLE_WRITE_ARCH_PROTOCOL_GUID
       B5 70 7D BA 31 DD 24 53
0xD4 : 03                               AND
0xD5 : 03                               AND
0xD6 : 03                               AND
0xD7 : 02                               PUSH
0xD8 : F5 3F 5E 66 CC 46 d4 11          EFI_WATCHDOG_TIMER_ARCH_PROTOCOL_GUID
       9A 38 00 90 27 3F C1 4D
0xE8 : 03                               AND
0xE9 : 08                               END
```

## Dependency Expression Reverse Polish Notation (RPN)

The actual equations will be presented by the DXE driver in a simple-to-evaluate form, namely postfix.

The following is a BNF encoding of this grammar. See Dependency Expression Instruction Set for definitions of the dependency expressions.

```
<statement> ::= SOR <expression> END |
                BEFORE <guid> END |
                AFTER <guid> END |
                <expression> END

<expression> ::= PUSH <guid> |
                 TRUE |
                 FALSE |
                 <expression> NOT |
                 <expression> <expression> OR |
                 <expression> <expression> AND
```

# DXE Dispatcher State Machine

## DXE Dispatcher State Machine

The DXE Dispatcher is responsible for tracking the state of a DXE driver from the time that the DXE driver is discovered in a firmware volume until the DXE Foundation is terminated with a call to **`ExitBootServices()`**. During this time, each DXE driver may be in one of several different states. The state machine that the DXE Dispatcher must use to track a DXE driver is shown in the figure below.

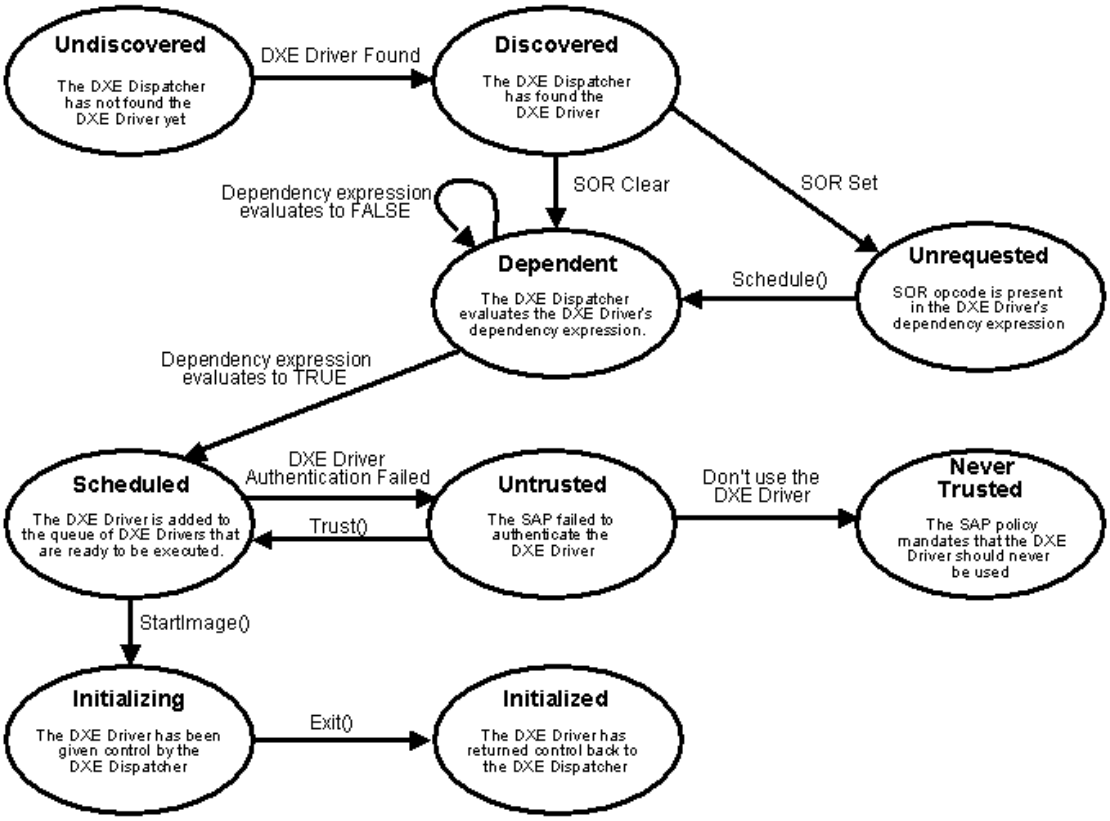


**Figure 10.1. DXE Driver States**

A DXE driver starts in the “Undiscovered” state, which means that the DXE driver is in a firmware volume that the DXE Dispatcher does not know about yet. When the DXE Dispatcher discovers a new firmware volume, any DXE drivers from that firmware volume listed in the *a priori* file are immediately loaded and executed. DXE drivers listed in the *a priori* file are immediately promoted to the “Scheduled” state. The firmware volume is then searched for DXE drivers that are not listed in the *a priori* file. Any DXE drivers found are promoted from the “Undiscovered” to the “Discovered” state. The dependency expression for each DXE driver is evaluated. If the SOR opcode is present in a DXE driver’s dependency expression, then the DXE driver is placed in the “Unrequested” state. If the SOR opcode is not present in the DXE driver’s dependency expression,

then the DXE driver is placed in the “Dependent” state. Once a DXE driver is in the "Unrequested" state, it may only be promoted to the “Dependent” state with a call to the DXE Service **Schedule()**.

Once a DXE Driver is in the “Dependent” state, the DXE Dispatcher will evaluate the DXE driver’s dependency expression. If the DXE driver does not have a dependency expression, then a dependency expression of all the architectural protocols ANDed together is assumed for that DXE driver. If the dependency expression evaluates to **FALSE**, then the DXE driver stays in the “Dependent” state. If the dependency expression never evaluates to **TRUE**, then it will never leave the “Dependent” state. If the dependency expression evaluates to **TRUE**, then the DXE driver will be promoted to the “Scheduled” state.

A DXE driver that is prompted to the “Scheduled” state is added to the end of the queue of other DXE drivers that have been promoted to the “Scheduled” state. When the DXE driver has reached the head of the queue, the DXE Dispatcher must use the services of the Security Authentication Protocol (SAP) to check the authentication status of the DXE Driver. If the Security Authentication Protocol deems that the DXE Driver violates the security policy of the platform, then the DXE Driver is placed in the “Untrusted” state. The Security Authentication Protocol can also tell the DXE Dispatcher that the DXE driver should never be executed and be placed in the “Never Trusted” state. If a DXE driver is placed in the “Untrusted” state, it can only be promoted back to the “Scheduled” state with a call to the DXE Service **Trust()**.

Once a DXE driver has reached the head of the scheduled queue, and the DXE driver has passed the authentication checks of the Security Authentication Protocol, the DXE driver is loaded into memory with the Boot Service **LoadImage()**. Control is then passed from the DXE Dispatcher to the DXE driver with the Boot Service **StartImage()**. When **StartImage()** is called for a DXE driver, that DXE driver is promoted to the “Initializing” state. The DXE driver returns control to the DXE Dispatcher through the Boot Service **Exit()**. When a DXE driver has returned control to the DXE Dispatcher, the DXE driver is in the terminal state called “Initialized.”

The DXE Dispatcher is responsible for draining the queue of DXE drivers in the “Scheduled” state until the queue is empty. Once the queue is empty, then DXE Dispatcher must evaluate all the DXE drivers in the "Dependent” state to see if any of them need to be promoted to the “Scheduled” state. These evaluations need to be performed every time one or more DXE drivers have been promoted to the “Initialized” state, because those DXE drivers may have produced protocol interfaces for which the DXE drivers in the "Dependent” state are waiting.

## Example Orderings

The order that DXE drivers are loaded and executed by the DXE Dispatcher is a mix of strong and weak orderings. The strong orderings are specified through *a priori* files, and the weak orderings are specified by dependency expressions in DXE drivers. The figure below shows the contents of a sample firmware volume that contains the following:

- DXE Foundation image
- DXE driver images
- An *a priori* file

The order that these images appear in the firmware volume is arbitrary. The DXE Foundation and the DXE Dispatcher must not make any assumptions about the locations of files in firmware volumes. The *a priori* file contains the GUID file names of the DXE drivers that are to be loaded and executed first. The dependency expressions and the protocols that each DXE driver produces is shown next to each DXE driver image in the firmware volume.

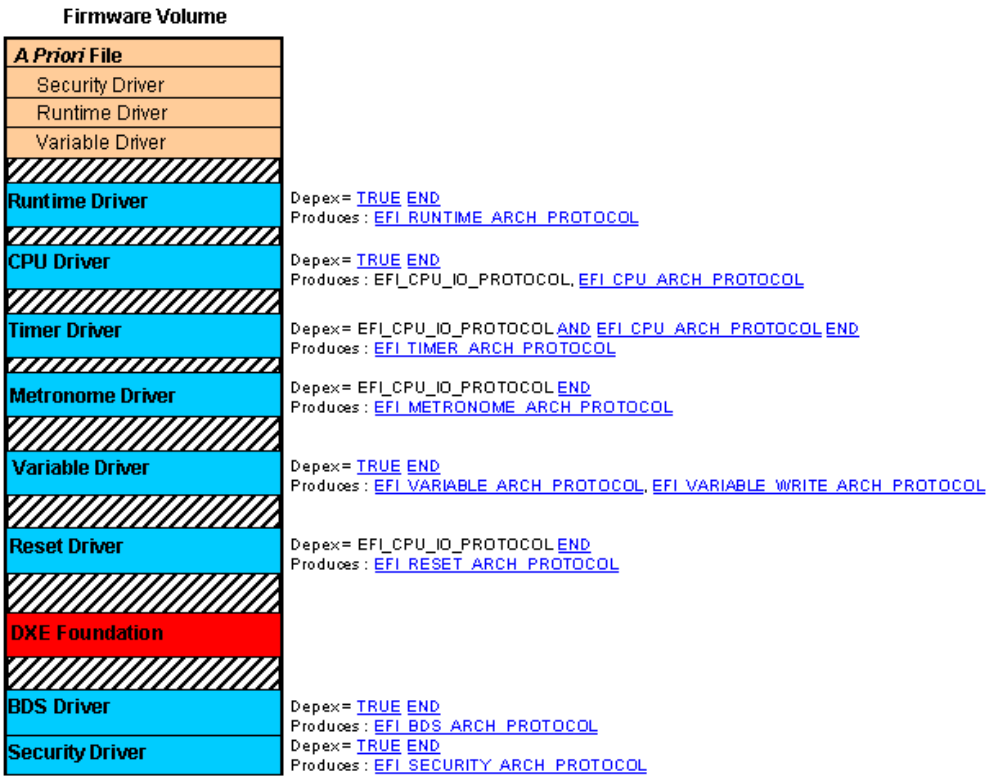


**Figure 10.2. Sample Firmware Volume**

Based on the contents of the firmware volume in the figure above, the Security Driver, Runtime Driver, and Variable Driver will always be executed first. This is an example of a strongly ordered dispatch due to the *a priori* file. The DXE Dispatcher will then evaluate the dependency

expressions of the remaining DXE drivers to determine the order that they will be executed. Based on the dependency expressions and the protocols that each DXE driver produces, there are 30 valid orderings from which the DXE Dispatcher may choose. The BDS Driver and CPU Driver tie for the next drivers to be scheduled, because their dependency expressions are simply TRUE. A dependency expression of TRUE means that the DXE driver does not require any other protocol interfaces to be executed. The DXE Dispatcher may choose either one of these drivers to be scheduled first. The Timer Driver, Metronome Driver, and Reset Driver all depend on the protocols produced by the CPU Driver. Once the CPU Driver has been loaded and executed, the Timer Driver, Metronome Driver, and Reset Driver may be scheduled in any order. The table below shows all 30 possible orderings from the sample firmware volume in the figure above. Each ordering is listed from left to right across the table. A reasonable implementation of a DXE Dispatcher would consistently produce the same ordering for a given system configuration. If the configuration of the system is changed in any way (including a order of files stored in a firmware volume), then a different dispatch ordering may be generated, but this new ordering should be consistent until the next system configuration change.

**Table 10.2. DXE Dispatcher Orderings**

| | Dispatch Order | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **1** | **2** | **3** | **4** | **5** | **6** | **7** | **8** |
| 1 | Security | Runtime | Variable | BDS | CPU | Timer | Metronome | Reset |
| 2 | Security | Runtime | Variable | BDS | CPU | Timer | Reset | Metronome |
| 3 | Security | Runtime | Variable | BDS | CPU | Metronome | Timer | Reset |
| 4 | Security | Runtime | Variable | BDS | CPU | Metronome | Reset | Timer |
| 5 | Security | Runtime | Variable | BDS | CPU | Reset | Timer | Metronome |
| 6 | Security | Runtime | Variable | BDS | CPU | Reset | Metronome | Timer |
| 7 | Security | Runtime | Variable | CPU | BDS | Timer | Metronome | Reset |
| 8 | Security | Runtime | Variable | CPU | BDS | Timer | Reset | Metronome |
| 9 | Security | Runtime | Variable | CPU | BDS | Metronome | Timer | Reset |
| 10 | Security | Runtime | Variable | CPU | BDS | Metronome | Reset | Timer |
| 11 | Security | Runtime | Variable | CPU | BDS | Reset | Timer | Metronome |
| 12 | Security | Runtime | Variable | CPU | BDS | Reset | Metronome | Timer |
| 13 | Security | Runtime | Variable | CPU | Timer | BDS | Metronome | Reset |
| 14 | Security | Runtime | Variable | CPU | Timer | BDS | Reset | Metronome |
| 15 | Security | Runtime | Variable | CPU | Timer | Metronome | BDS | Reset |
| 16 | Security | Runtime | Variable | CPU | Timer | Metronome | Reset | BDS |
| 17 | Security | Runtime | Variable | CPU | Timer | Reset | BDS | Metronome |
| 18 | Security | Runtime | Variable | CPU | Timer | Reset | Metronome | BDS |
| 19 | Security | Runtime | Variable | CPU | Metronome | Timer | BDS | Reset |
| 20 | Security | Runtime | Variable | CPU | Metronome | Timer | Reset | BDS |

| | Dispatch Order | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| | **1** | **2** | **3** | **4** | **5** | **6** | **7** | **8** |
| 21 | Security | Runtime | Variable | CPU | Metronome | BDS | Timer | Reset |
| 22 | Security | Runtime | Variable | CPU | Metronome | BDS | Reset | Timer |
| 23 | Security | Runtime | Variable | CPU | Metronome | Reset | Timer | BDS |
| 24 | Security | Runtime | Variable | CPU | Metronome | Reset | BDS | Timer |
| 25 | Security | Runtime | Variable | CPU | Reset | Timer | Metronome | BDS |
| 26 | Security | Runtime | Variable | CPU | Reset | Timer | BDS | Metronome |
| 27 | Security | Runtime | Variable | CPU | Reset | Metronome | Timer | BDS |
| 28 | Security | Runtime | Variable | CPU | Reset | Metronome | BDS | Timer |
| 29 | Security | Runtime | Variable | CPU | Reset | BDS | Timer | Metronome |
| 30 | Security | Runtime | Variable | CPU | Reset | BDS | Metronome | Timer |

## Security Considerations

The DXE Dispatcher is required to use the services of the Security Architectural Protocol every time a firmware volume is discovered and before each DXE driver is executed. Because the Security Architectural Protocol is produced by a DXE driver, there will be at least one firmware volume discovered, and one or more DXE drivers loaded and executed before the Security Architectural Protocol is installed. The DXE Dispatcher should not attempt to use the services of the Security Architectural Protocol until the Security Architectural Protocol is installed. If a platform requires the Security Architectural Protocol to be present very early in the DXE phase, then the *a priori* file may be used to specify the name of the DXE driver that produces the Security Architectural Protocol.

The Security Architectural Protocol provides a service to evaluate the authentication status of a file. This service can also be used to evaluate the authenticate status of a firmware volume. If the authentication status is good, then no action is taken. If there is a problem with the firmware volume's authentication status, then the Security Architectural Protocol may perform a platform specific action. One option is to force the DXE Dispatcher to ignore the firmware volume so no DXE drivers will be loaded and executed from it. Another is to log the fact that the DXE Dispatcher is going to start dispatching DXE driver from a firmware volume with a questionable authentication status.

The Security Architectural Protocol can also be used to evaluate the authentication status of each DXE driver discovered in a firmware volume. If the authentication status is good, then no action is taken. If there is a problem with the DXE driver's authentication status, then the Security Architectural Protocol may take a platform-specific action. One possibility is to force the DXE driver into the "Untrusted" state, so it will not be considered for dispatch until the Boot Service **`Trust()`** is called for that DXE driver. Another possibility is to have the DXE Dispatcher place the DXE driver in the "Never Trusted" state, so it will never be loaded or executed. Another option is to log the fact that a DXE driver with a questionable authentication status is about to be loaded and executed.

# 11
# DXE Drivers

## Introduction

The DXE architecture provides a rich set of extensible services that provides for wide variety of different system firmware designs. The DXE Foundation provides the generic services required to locate and execute DXE drivers. The DXE drivers are the components that actually initialize the platform and provide the services required to boot an EFI-compliant operating system or a set of EFI-compliant system utilities. There are many possible firmware implementations for any given platform. Because the DXE Foundation has fixed functionality, all the added value and flexibility in a firmware design is embodied in the implementation and organization of DXE drivers.

There are two basic classes of DXE drivers:

- Early DXE Drivers
- DXE Drivers that follow the EFI Driver Model

Additional classifications of DXE drivers are also possible.

All DXE drivers may consume the EFI Boot Services, EFI Runtime Services, and DXE Services to perform their functions. DXE drivers must use dependency expressions to guarantee that the services and protocol interfaces they require are available before they are executed. See the following topics for the DXE Architectural Protocols upon which the services depend:

- EFI Boot Services Dependencies
- EFI Runtime Services Dependencies
- DXE Services Dependencies

## Classes of DXE Drivers

### Basic

### Early DXE Drivers

The first class of DXE drivers are those that execute very early in the DXE phase. The execution order of these DXE drivers depends on the following:

- The presence and contents of an *a priori* file
- The evaluation of dependency expressions

These early DXE drivers will typically contain basic services, processor initialization code, chipset initialization code, and platform initialization code. These early drivers will also typically produce the DXE Architectural Protocols that are required for the DXE Foundation to produces its full complement of EFI Boot Services and EFI Runtime Services. To support the fastest possible boot time, as much initialization should be deferred to the DXE drivers that follow EFI Driver Model described in the *EFI 1.10 Specification*.

The early DXE drivers need to be aware that not all of the EFI Boot Services, EFI Runtime Services, and DXE Services may be available when they execute because not all of the DXE Architectural Protocols may be been registered yet.

## DXE Drivers That Follow the EFI Driver Model

The second class of DXE drivers are those that follow the EFI Driver Model in the *EFI 1.10 Specification.* These drivers do not touch any hardware resources when they initialize. Instead, they register a Driver Binding Protocol interface in the handle database. The set of Driver Binding Protocols are used by the Boot Device Selection (BDS) phase to connect the drivers to the devices that are required to establish consoles and provide access to boot devices. The DXE drivers that follow the EFI Driver Model ultimately provide software abstractions for console devices and boot devices, but only when they are explicitly asked to do so.

The DXE drivers that follow the EFI Driver Model do not need to be concerned with dependency expressions. These drivers simply register the Driver Binding Protocol in the handle database when they are executed, and this operation can be performed without the use of any DXE Architectural Protocols. DXE drivers with empty dependency expressions will not be dispatched by the DXE Dispatcher until all of the DXE Architectural Protocols have been installed.

## Additional

## Additional Classifications

DXE drivers can also be classified as the following:

- Boot service drivers
- Runtime drivers

Boot service drivers provide services that are available until the **`ExitBootServices()`** function is called. When **`ExitBootServices()`** is called, all the memory used by boot service drivers is released for use by an operating system.

Runtime drivers provide services that are available before and after **`ExitBootServices()`** is called, including the time that an operating system is running. All of the services in the EFI Runtime Services Table are produced by runtime drivers.

The DXE Foundation is considered a boot service component, so the DXE Foundation is also released when **`ExitBootServices()`** is called. As a result, runtime drivers may not use any of the EFI Boot Services, DXE Services, or services produced by boot service drivers after **`ExitBootServices()`** is called.

# 12
# DXE Architectural Protocols

## Introduction

The DXE Foundation is abstracted from the platform hardware through a set of architectural protocols. These protocols function just like other protocols in every way. The only difference is that these architectural protocols are the protocols that the DXE Foundation itself consumes to produce the EFI Boot Services, EFI Runtime Services, and DXE Services. DXE drivers that are loaded from firmware volumes produce the DXE Architectural Protocols. This means that the DXE Foundation must have enough services to load and start DXE drivers before even a single DXE driver is executed.

The DXE Foundation is passed a HOB list that must contain a description of some amount of system memory and at least one firmware volume. The system memory descriptors in the HOB list are used to initialize the EFI services that require only memory to function correctly. The system is also guaranteed to be running on only one processor in flat physical mode with interrupts disabled. The firmware volume is passed to the DXE Dispatcher, and the DXE Dispatcher must contain a read-only firmware file system driver to search for the *a priori* file and any DXE drivers in the firmware volumes. When a driver is discovered that needs to be loaded and executed, the DXE Dispatcher will use a PE/COFF loader to load and invoke the DXE driver. The early DXE drivers will produce the DXE Architectural Protocols, so the DXE Foundation can produce the full complement of EFI Boot Services and EFI Runtime Services.

The figure below shows the HOB list being passed to the DXE Foundation.

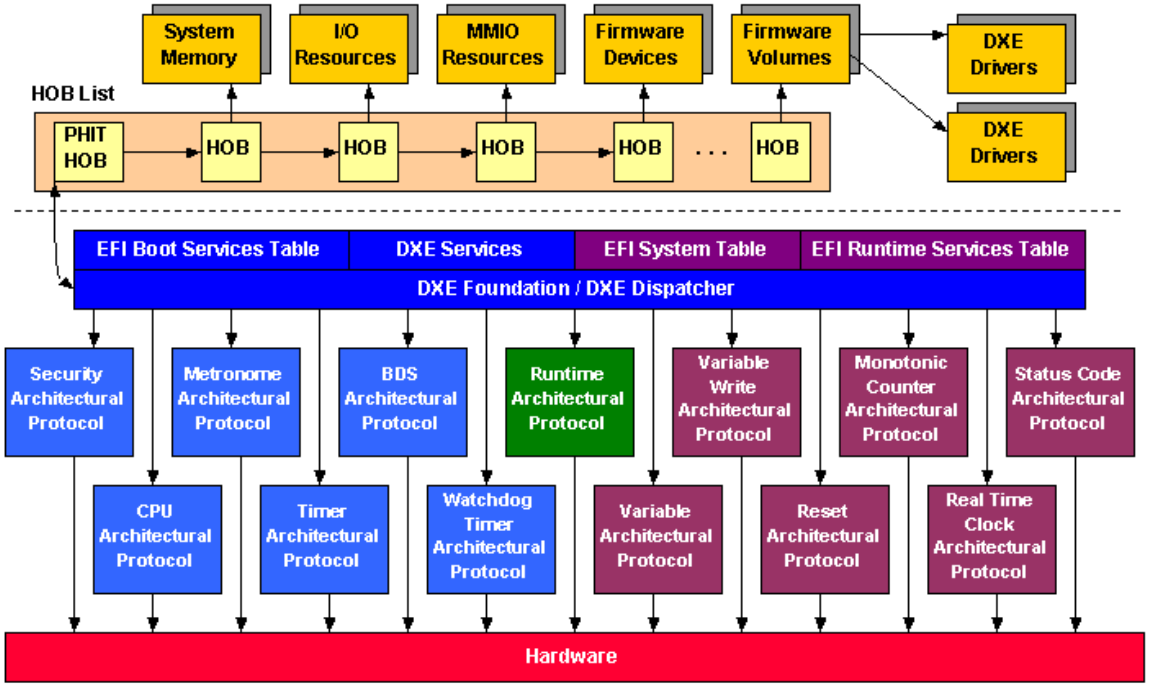


**Figure 12.1. DXE Architectural Protocols**

The DXE Foundation consumes the services of the DXE Architectural Protocols and produces the following:

- EFI System Table
- EFI Boot Services Table
- EFI Runtime Services Table
- DXE Services Table

The EFI Boot Services Table and DXE Services Table are allocated from EFI boot services memory, which means that the EFI Boot Services Table and DXE Services Table are freed when the OS runtime phase is entered. The EFI System Table and EFI Runtime Services Table are allocated from EFI runtime services memory, and they persist into the OS runtime phase.

The DXE Architectural Protocols shown on the left of the figure are used to produce the EFI Boot Services and DXE Services. The DXE Foundation and these protocols will be freed when the system transitions to the OS runtime phase. The DXE Architectural Protocols shown on the right are used to produce the EFI Runtime Services. These services will persist in the OS runtime phase. The Runtime Architectural Protocol in the middle is unique. This protocol provides the services that are required to transition the runtime services from physical mode to virtual mode under the direction of an OS. Once this transition is complete, the services of the Runtime Architectural Protocol can no longer be used. The following topics describe all of the DXE Architectural Protocols in detail.

# Boot Device Selection (BDS) Architectural Protocol

## EFI_BDS_ARCH_PROTOCOL

### Summary

Transfers control from the DXE phase to an operating system or system utility. This protocol must be produced by a boot service or runtime DXE driver and may only be consumed by the DXE Foundation.

### GUID

```
#define EFI_BDS_ARCH_PROTOCOL_GUID  \

{0x665E3FF6,0x46CC,0x11d4,0x9A,0x38,0x00,0x90,0x27,0x3F,0xC1,0x4D
}
```

### Protocol Interface Structure

```
typedef struct {
  EFI_BDS_ENTRY  Entry;
} EFI_BDS_ARCH_PROTOCOL;
```

### Parameters

*Entry*

The entry point to BDS. See the **Entry()** function description. This call does not take any parameters, and the return value can be ignored. If it returns, then the dispatcher must be invoked again, if it never returns, then an operating system or a system utility have been invoked.

### Description

The **EFI_BDS_ARCH_PROTOCOL** transfers control from DXE to an operating system or a system utility. If there are not enough drivers initialized when this protocol is used to access the required boot device(s), then this protocol should add drivers to the dispatch queue and return control back to the dispatcher. Once the required boot devices are available, then the boot device can be used to load and invoke an OS or a system utility.

## EFI_BDS_ARCH_PROTOCOL.Entry()

### Summary

Performs Boot Device Selection (BDS) and transfers control from the DXE Foundation to the selected boot device. The implementation of the boot policy must follow the rules outlined in the Boot Manager chapter of the *EFI 1.10 Specification*. This boot policy allows for flexibility, so the platform vendor will typically customize the implementation of this service.

### Prototype

```
typedef
VOID
(EFIAPI *EFI_BDS_ENTRY) (
  IN EFI_BDS_ARCH_PROTOCOL  *This
  );
```

### Parameters

*This*

The **EFI_BDS_ARCH_PROTOCOL** instance.

### Description

This function uses policy data from the platform to determine what operating system or system utility should be loaded and invoked. This function call also optionally uses the user's input to determine the operating system or system utility to be loaded and invoked. When the DXE Foundation has dispatched all the drivers on the dispatch queue, this function is called. This function will attempt to connect the boot devices required to load and invoke the selected operating system or system utility. During this process, additional firmware volumes may be discovered that may contain addition DXE drivers that can be dispatched by the DXE Foundation. If a boot device cannot be fully connected, this function calls the DXE Service **Dispatch()** to allow the DXE drivers from any newly discovered firmware volumes to be dispatched. Then the boot device connection can be attempted again. If the same boot device connection operation fails twice in a row, then that boot device has failed, and should be skipped. This function should never return.

# CPU Architectural Protocol

## EFI_CPU_ARCH_PROTOCOL

### Summary

Abstracts the processor services that are required to implement some of the DXE services. This protocol must be produced by a boot service or runtime DXE driver and may only be consumed by the DXE Foundation and DXE drivers that produce architectural protocols.

### GUID

```
#define EFI_CPU_ARCH_PROTOCOL_GUID  \

{0x26baccb1,0x6f42,0x11d4,0xbc,0xe7,0x0,0x80,0xc7,0x3c,0x88,0x81}
```

### Protocol Interface Structure

```
typedef struct _EFI_CPU_ARCH_PROTOCOL {
  EFI_CPU_FLUSH_DATA_CACHE              FlushDataCache;
  EFI_CPU_ENABLE_INTERRUPT              EnableInterrupt;
  EFI_CPU_DISABLE_INTERRUPT             DisableInterrupt;
  EFI_CPU_GET_INTERRUPT_STATE           GetInterruptState;
  EFI_CPU_INIT                          Init;
  EFI_CPU_REGISTER_INTERRUPT_HANDLER    RegisterInterruptHandler;
  EFI_CPU_GET_TIMER_VALUE               GetTimerValue;
  EFI_CPU_SET_ATTRIBUTES                SetMemoryAttributes;
  UINT32                                NumberOfTimers;
  UINT32                                DmaBufferAlignment;
} EFI_CPU_ARCH_PROTOCOL;
```

### Parameters

*FlushDataCache*

Flushes a range of the processor's data cache. See the **FlushDataCache()** function description. If the processor does not contain a data cache, or the data cache is fully coherent, then this function can just return **EFI_SUCCESS**. If the processor does not support flushing a range of addresses from the data cache, then the entire data cache must be flushed. This function is used by the root bridge I/O abstractions to flush data caches for DMA operations.

*EnableInterrupt*

Enables interrupt processing by the processor. See the **EnableInterrupt()** function description. This function is used by the Boot Service **RaiseTPL()** and **RestoreTPL()**.

*DisableInterrupt*

Disables interrupt processing by the processor. See the **DisableInterrupt()** function description. This function is used by the Boot Service **RaiseTPL()** and **RestoreTPL()**.

*GetInterruptState*

Retrieves the processor's current interrupt state. See the **GetInterruptState()** function description.

*Init*

Generates an INIT on the processor. See the **Init()** function description. This function may be used by the **EFI_RESET** Protocol depending upon a specified boot path. If a processor cannot programmatically generate an INIT without help from external hardware, then this function returns **EFI_UNSUPPORTED**.

*RegisterInterruptHandler*

Associates an interrupt service routine with one of the processor's interrupt vectors. See the **RegisterInterruptHandler()** function description. This function is typically used by the **EFI_TIMER_ARCH_PROTOCOL** to hook the timer interrupt in a system. It can also be used by the debugger to hook exception vectors.

*GetTimerValue*

Returns the value of one of the processor's internal timers. See the **GetTimerValue()** function description.

*SetMemoryAttributes*

Attempts to set the attributes of a memory region. See the **SetMemoryAttributes()** function description.

*NumberOfTimers*

The number of timers that are available in a processor. The value in this field is a constant that must not be modified after the CPU Architectural Protocol is installed. All consumers must treat this as a read-only field.

*DmaBufferAlignment*

The size, in bytes, of the alignment required for DMA buffer allocations. This is typically the size of the largest data cache line in the platform. This value can be determined by looking at the data cache line sizes of all the caches present in the platform, and returning the largest. This is used by the root bridge I/O abstraction protocols to guarantee that no two DMA buffers ever share the same cache line. The value in this field is a constant that must not be modified after the CPU Architectural Protocol is installed. All consumers must treat this as a read-only field.

## Description

The **EFI_CPU_ARCH_PROTOCOL** is used to abstract processor-specific functions from the DXE Foundation. This includes flushing caches, enabling and disabling interrupts, hooking interrupt vectors and exception vectors, reading internal processor timers, resetting the processor, and determining the processor frequency.

The GCD memory space map is initialized by the DXE Foundation based on the contents of the HOB list. The HOB list contains the capabilities of the different memory regions, but it does not contain their current attributes. The DXE driver that produces the **EFI_CPU_ARCH_PROTOCOL** is responsible for maintaining the current attributes of the memory regions visible to the processor.

This means that the DXE driver that produces the **EFI_CPU_ARCH_PROTOCOL** must seed the GCD memory space map with the initial state of the attributes for all the memory regions visible to the processor. The DXE Service **SetMemorySpaceAttributes()** allows the attributes of a memory range to be modified. The **SetMemorySpaceAttributes()** DXE Service is implemented using the **SetMemoryAttributes()** service of the **EFI_CPU_ARCH_PROTOCOL**.

To initialize the state of the attributes in the GCD memory space map, the DXE driver that produces the **EFI_CPU_ARCH_PROTOCOL** must call the DXE Service **SetMemorySpaceAttributes()** for all the different memory regions visible to the processor passing in the current attributes. This, in turn, will call back to the **SetMemoryAttributes()** service of the **EFI_CPU_ARCH_PROTOCOL**, and all of these calls must return **EFI_SUCCESS**, since the DXE Foundation is only requesting that the attributes of the memory region be set to their current settings. This will force the current attributes in the GCD memory space map to be set to these current settings. After this initialization is complete, the next call to the DXE Service **GetMemorySpaceMap()** will correctly show the current attributes of all the memory regions. In addition, any future calls to the DXE Service **SetMemorySpaceAttributes()** will in turn call the **EFI_CPU_ARCH_PROTOCOL** so see of those attributes can be modified, and if they can, the GCD memory space map will be updated accordingly.

## EFI_CPU_ARCH_PROTOCOL.FlushDataCache()

### Summary

Flushes a range of the processor's data cache. If the processor does not contain a data cache, or the data cache is fully coherent, then this function can just return **EFI_SUCCESS**. If the processor does not support flushing a range of addresses from the data cache, then the entire data cache must be flushed. This function is used by the root bridge I/O abstractions to flush caches for DMA operations.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_FLUSH_DATA_CACHE) (
  IN EFI_CPU_ARCH_PROTOCOL   *This,
  IN EFI_PHYSICAL_ADDRESS    Start,
  IN UINT64                  Length,
  IN EFI_CPU_FLUSH_TYPE      FlushType
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*Start*

The beginning physical address to flush from the processor's data cache.

*Length*

The number of bytes to flush from the processor's data cache. This function may flush more bytes than *Length* specifies depending upon the granularity of the flush operation that the processor supports.

*FlushType*

Specifies the type of flush operation to perform. Type **EFI_CPU_FLUSH_TYPE** is defined in "Related Definitions" below.

### Description

This function flushes the range of addresses from *Start* to *Start+Length* from the processor's data cache. If *Start* is not aligned to a cache line boundary, then the bytes before *Start* to the preceding cache line boundary are also flushed. If *Start+Length* is not aligned to a cache line boundary, then the bytes past *Start+Length* to the end of the next cache line boundary are also flushed. If the address range is flushed, then **EFI_SUCCESS** is returned. If the address range cannot be flushed, then **EFI_DEVICE_ERROR** is returned. If the processor does not support the flush type specified by *FlushType*, then **EFI_UNSUPPORTED** is returned. The *FlushType* of *EfiCpuFlushTypeWriteBackInvalidate* must be supported. If the data cache is fully coherent with all DMA operations, then this function can just return **EFI_SUCCESS**. If the

processor does not support flushing a range of the data cache, then the entire data cache can be flushed.

## Related Definitions

```
typedef enum {
  EfiCpuFlushTypeWriteBackInvalidate,
  EfiCpuFlushTypeWriteBack,
  EfiCpuFlushTypeInvalidate,
  EfiCpuMaxFlushType
} EFI_CPU_FLUSH_TYPE;
```

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The address range from *Start* to *Start+Length* was flushed from the processor's data cache. |
| EFI_UNSUPPORTED | The processor does not support the cache flush type specified by *FlushType*. |
| EFI_DEVICE_ERROR | The address range from *Start* to *Start+Length* could not be flushed from the processor's data cache. |

## EFI_CPU_ARCH_PROTOCOL.EnableInterrupt()

### Summary

Enables interrupt processing by the processor. This function is used to implement the Boot Services **RaiseTPL()** and **RestoreTPL()**.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_ENABLE_INTERRUPT) (
  IN EFI_CPU_ARCH_PROTOCOL  *This
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

### Description

This function enables interrupt processing by the processor. If interrupts are enabled, then **EFI_SUCCESS** is returned. Otherwise, **EFI_DEVICE_ERROR** is returned.

### Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | Interrupts are enabled on the processor. |
| EFI_DEVICE_ERROR | Interrupts could not be enabled on the processor. |

# EFI_CPU_ARCH_PROTOCOL.DisableInterrupt()

## Summary

Disables interrupt processing by the processor. This function is used to implement the Boot Services **RaiseTPL()** and **RestoreTPL()**.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_DISABLE_INTERRUPT) (
  IN EFI_CPU_ARCH_PROTOCOL  *This
  );
```

## Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

## Description

This function disables interrupt processing by the processor. If interrupts are disabled, then **EFI_SUCCESS** is returned. Otherwise, **EFI_DEVICE_ERROR** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | Interrupts are disabled on the processor. |
| EFI_DEVICE_ERROR | Interrupts could not be disabled on the processor. |

## EFI_CPU_ARCH_PROTOCOL.GetInterruptState()

### Summary

Retrieves the processor's current interrupt state.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_GET_INTERRUPT_STATE) (
  IN  EFI_CPU_ARCH_PROTOCOL   *This,
  OUT BOOLEAN                 *State
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*State*

A pointer to the processor's current interrupt state. Set to **TRUE** if interrupts are enabled and **FALSE** if interrupts are disabled.

### Description

This function retrieves the processor's current interrupt state a returns it in *State*. If interrupts are currently enabled, then **TRUE** is returned. If interrupts are currently disabled, then **FALSE** is returned. If *State* is **NULL**, then **EFI_INVALID_PARAMETER** is returned. Otherwise, **EFI_SUCCESS** is returned.

### Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The processor's current interrupt state was returned in *State*. |
| EFI_INVALID_PARAMETER | *State* is **NULL**. |

# EFI_CPU_ARCH_PROTOCOL.Init()

## Summary

Generates an INIT on the processor.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_INIT) (
  IN EFI_CPU_ARCH_PROTOCOL   *This,
  IN EFI_CPU_INIT_TYPE       InitType
  );
```

## Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*InitType*

The type of processor INIT to perform. Type **EFI_CPU_INIT_TYPE** is defined in "Related Definitions" below.

## Description

This function generates an INIT on the processor. If this function succeeds, then the processor will be reset, and control will not be returned to the caller. If *InitType* is not supported by this processor, or the processor cannot programmatically generate an INIT without help from external hardware, then **EFI_UNSUPPORTED** is returned. If an error occurs attempting to generate an INIT, then **EFI_DEVICE_ERROR** is returned.

## Related Definitions

```
typedef enum {
  EfiCpuInit,
  EfiCpuMaxInitType
} EFI_CPU_INIT_TYPE;
```

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The processor INIT was performed. This return code should never be seen. |
| EFI_UNSUPPORTED | The processor INIT operation specified by *InitType* is not supported by this processor. |
| EFI_DEVICE_ERROR | The processor INIT failed. |

## EFI_CPU_ARCH_PROTOCOL. RegisterInterruptHandler()

### Summary

Registers a function to be called from the processor interrupt handler.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_REGISTER_INTERRUPT_HANDLER) (
  IN EFI_CPU_ARCH_PROTOCOL       *This,
  IN EFI_EXCEPTION_TYPE          InterruptType,
  IN EFI_CPU_INTERRUPT_HANDLER   InterruptHandler
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*InterruptType*

Defines which interrupt or exception to hook. Type **EFI_EXCEPTION_TYPE** and the valid values for this parameter are defined in **EFI_DEBUG_SUPPORT_PROTOCOL** of the *EFI 1.10 Specification.*

*InterruptHandler*

A pointer to a function of type **EFI_CPU_INTERRUPT_HANDLER** that is called when a processor interrupt occurs. If this parameter is **NULL**, then the handler will be uninstalled. Type **EFI_CPU_INTERRUPT_HANDLER** is defined in "Related Definitions" below.

### Description

The **RegisterInterruptHandler()** function registers and enables the handler specified by *InterruptHandler* for a processor interrupt or exception type specified by *InterruptType*. If *InterruptHandler* is **NULL**, then the handler for the processor interrupt or exception type specified by *InterruptType* is uninstalled. The installed handler is called once for each processor interrupt or exception.

If the interrupt handler is successfully installed or uninstalled, then **EFI_SUCCESS** is returned.

If *InterruptHandler* is not **NULL**, and a handler for *InterruptType* has already been installed, then **EFI_ALREADY_STARTED** is returned.

If *InterruptHandler* is **NULL**, and a handler for *InterruptType* has not been installed, then **EFI_INVALID_PARAMETER** is returned.

If *InterruptType* is not supported, then **EFI_UNSUPPORTED** is returned.

The **EFI_CPU_ARCH_PROTOCOL** implementation of this function must handle saving and restoring system context to the system context record around calls to the interrupt handler. It must

also perform the necessary steps to return to the context that was interrupted by the interrupt. No chaining of interrupt handlers is allowed.

## Related Definitions

```
typedef
VOID
(*EFI_CPU_INTERRUPT_HANDLER) (
  IN EFI_EXCEPTION_TYPE  InterruptType,
  IN EFI_SYSTEM_CONTEXT  SystemContext
  );
```

*InterruptType*

Defines the type of interrupt or exception that occurred on the processor. This parameter is processor architecture specific. The type **EFI_EXCEPTION_TYPE** and the valid values for this parameter are defined in **EFI_DEBUG_SUPPORT_PROTOCOL** of the *EFI 1.10 Specification.*

*SystemContext*

A pointer to the processor context when the interrupt occurred on the processor. Type **EFI_SYSTEM_CONTEXT** is defined in the **EFI_DEBUG_SUPPORT_PROTOCOL** of the *EFI 1.10 Specification.*

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The handler for the processor interrupt was successfully installed or uninstalled. |
| EFI_ALREADY_STARTED | *InterruptHandler* is not **NULL**, and a handler for *InterruptType* was previously installed. |
| EFI_INVALID_PARAMETER | *InterruptHandler* is **NULL**, and a handler for *InterruptType* was not previously installed. |
| EFI_UNSUPPORTED | The interrupt specified by *InterruptType* is not supported. |

## EFI_CPU_ARCH_PROTOCOL.GetTimerValue()

### Summary

Returns a timer value from one of the processor's internal timers.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_GET_TIMER_VALUE) (
  IN  EFI_CPU_ARCH_PROTOCOL  *This,
  IN  UINT32                 TimerIndex,
  OUT UINT64                 *TimerValue,
  OUT UINT64                 *TimerPeriod  OPTIONAL
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*TimerIndex*

Specifies which processor timer is to be returned in *TimerValue*. This parameter must be between 0 and *NumberOfTimers*-**1**.

*TimerValue*

Pointer to the returned timer value.

*TimerPeriod*

A pointer to the amount of time that passes in femtoseconds ($10^{-15}$) for each increment of *TimerValue*. If *TimerValue* does not increment at a predictable rate, then 0 is returned. The amount of time that has passed between two calls to **GetTimerValue()** can be calculated with the formula (*TimerValue2* – *TimerValue1*) * *TimerPeriod*. This parameter is optional and may be **NULL**.

### Description

This function reads the processor timer specified by *TimerIndex* and returns it in *TimerValue*. If *TimerValue* is **NULL**, then **EFI_INVALID_PARAMETER** is returned. If *TimerPeriod* is not **NULL**, then the amount of time that passes in femtoseconds($10^{-15}$) for each increment if *TimerValue* is returned in *TimerPeriod*. If the timer does not run at a predictable rate, then a *TimerPeriod* of 0 is returned. If *TimerIndex* does not specify a valid timer in this processor, then **EFI_INVALID_PARAMETER** is returned. The valid range for *TimerIndex* is 0..*NumberOfTimers*-1. If the processor does not contain any readable timers, then this function returns **EFI_UNSUPPORTED**. If an error occurs attempting to read one of the processor's timers, then **EFI_DEVICE_ERROR** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The processor timer value specified by *TimerIndex* was returned in *TimerValue*. |
| EFI_INVALID_PARAMETER | *TimerValue* is **NULL**. |
| EFI_INVALID_PARAMETER | *TimerIndex* is not valid. |
| EFI_UNSUPPORTED | The processor does not have any readable timers. |
| EFI_DEVICE_ERROR | An error occurred attempting to read one of the processor's timers. |

## EFI_CPU_ARCH_PROTOCOL.SetMemoryAttributes()

### Summary

Attempts to set the attributes for a memory range.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_CPU_SET_MEMORY_ATTRIBUTES) (
  IN  EFI_CPU_ARCH_PROTOCOL  *This,
  IN  EFI_PHYSICAL_ADDRESS   BaseAddress,
  IN  UINT64                 Length,
  IN  UINT64                 Attributes
  );
```

### Parameters

*This*

The **EFI_CPU_ARCH_PROTOCOL** instance.

*BaseAddress*

The physical address that is the start address of a memory region. Type **EFI_PHYSICAL_ADDRESS** is defined in the **AllocatePages()** function description in the *EFI 1.10 Specification.*

*Length*

The size in bytes of the memory region.

*Attributes*

The bit mask of attributes to set for the memory region. See the EFI Boot Service **GetMemoryMap()** for the set of legal attribute bits.

### Description

This function modifies the attributes for the memory region specified by *BaseAddress* and *Length* from their current attributes to the attributes specified by *Attributes*. If this modification of attributes succeeds, then **EFI_SUCCESS** is returned.

If *Length* is zero, then **EFI_INVALID_PARAMETER** is returned.

If the processor does not support one or more bytes of the memory range specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If the attributes specified by *Attributes* are not supported for the memory region specified by *BaseAddress* and *Length*, then **EFI_UNSUPPORTED** is returned.

If the attributes for one or more bytes of the memory range specified by *BaseAddress* and *Length* cannot be modified because the current system policy does not allow them to be modified, then **EFI_ACCESS_DENIED** is returned.

If there are not enough system resources available to modify the attributes of the memory range, then **EFI_OUT_OF_RESOURCES** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The attributes were set for the memory region. |
| EFI_INVALID_PARAMETER | *Length* is zero. |
| EFI_UNSUPPORTED | The processor does not support one or more bytes of the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_UNSUPPORTED | The bit mask of attributes is not support for the memory resource range specified by *BaseAddress* and *Length*. |
| EFI_ACCESS_DENIED | The attributes for the memory resource range specified by *BaseAddress* and *Length* cannot be modified. |
| EFI_OUT_OF_RESOURCES | There are not enough system resources to modify the attributes of the memory resource range. |

# Metronome Architectural Protocol

## EFI_METRONOME_ARCH_PROTOCOL

### Summary

Used to wait for ticks from a known time source in a platform. This protocol may be used to implement a simple version of the **Stall()** Boot Service. This protocol must be produced by a boot service or runtime DXE driver and may only be consumed by the DXE Foundation and DXE drivers that produce DXE Architectural Protocols.

### GUID

```
#define EFI_METRONOME_ARCH_PROTOCOL_GUID  \

{0x26baccb2,0x6f42,0x11d4,0xbc,0xe7,0x0,0x80,0xc7,0x3c,0x88,0x81}
```

### Protocol Interface Structure

```
typedef struct _EFI_METRONOME_ARCH_PROTOCOL {
  EFI_METRONOME_WAIT_FOR_TICK  WaitForTick;
  UINT32                       TickPeriod;
} EFI_METRONOME_ARCH_PROTOCOL;
```

### Parameters

*WaitForTick*

Waits for a specified number of ticks from a known time source in the platform. See the **WaitForTick()** function description. The actual time passed between entry of this function and the first tick is between 0 and *TickPeriod* 100 nS units. To guarantee that at least *TickPeriod* time has elapsed, wait for two ticks.

*TickPeriod*

The period of platform's known time source in 100 nS units. This value on any platform must be at least 10 uS, and must not exceed 200 uS. The value in this field is a constant that must not be modified after the Metronome architectural protocol is installed. All consumers must treat this as a read-only field.

### Description

This protocol provides access to a known time source in the platform to the DXE Foundation. The DXE Foundation uses this known time source to produce DXE Foundation services that require calibrated delays.

# EFI_METRONOME_ARCH_PROTOCOL.WaitForTick()

## Summary

Waits for a specified number of ticks from a known time source in a platform.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_METRONOME_WAIT_FOR_TICK) (
  IN EFI_METRONOME_ARCH_PROTOCOL  *This,
  IN UINT32                       TickNumber
  );
```

## Parameters

*This*

The **EFI_METRONOME_ARCH_PROTOCOL** instance.

*TickNumber*

Number of ticks to wait.

## Description

The **WaitForTick()** function waits for the number of ticks specified by *TickNumber* from a known time source in the platform. If *TickNumber* of ticks are detected, then **EFI_SUCCESS** is returned. The actual time passed between entry of this function and the first tick is between 0 and *TickPeriod* 100 nS units. If you want to guarantee that at least *TickPeriod* time has elapsed, wait for two ticks. This function waits for a hardware event to determine when a tick occurs. It is possible for interrupt processing, or exception processing to interrupt the execution of the **WaitForTick()** function. Depending on the hardware source for the ticks, it is possible for a tick to be missed. This function cannot guarantee that ticks will not be missed. If a timeout occurs waiting for the specified number of ticks, then **EFI_TIMEOUT** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The wait for the number of ticks specified by *TickNumber* succeeded. |
| EFI_TIMEOUT | A timeout occurred waiting for the specified number of ticks. |

# Monotonic Counter Architectural Protocol

## EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL

### Summary

Provides the services required to access the system's monotonic counter. This protocol must be produced by a runtime DXE driver and may only be consumed by the DXE Foundation and DXE drivers that produce DXE Architectural Protocols.

### GUID

```
#define EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL_GUID  \

{0x1da97072,0xbddc,0x4b30,0x99,0xf1,0x72,0xa0,0xb5,0x6f,0xff,0x2a
}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **`GetNextHighMonotonicCount()`** field of the EFI Runtime Services Table and the **`GetNextMonotonicCount()`** field of the EFI Boot Services Table. See Services - Runtime Services and Services - Boot Services for details on these services. After the field of the EFI Runtime Services Table and the field of the EFI Boot Services Table have been initialized, the driver must install the **`EFI_MONOTONIC_COUNTER_ARCH_PROTOCOL_GUID`** on a new handle with a **`NULL`** interface pointer. The installation of this protocol informs the DXE Foundation that the monotonic counter services are now available and that the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table and the 32-bit CRC of the EFI Boot Services Table.

# Real Time Clock Architectural Protocol

## EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL

### Summary

Provides the services required to access a system's real time clock hardware. This protocol must be produced by a runtime DXE driver and may only be consumed by the DXE Foundation.

### GUID

```
#define EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL_GUID  \

{0x27CFAC87,0x46CC,0x11d4,0x9A,0x38,0x00,0x90,0x27,0x3F,0xC1,0x4D
}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **GetTime()**, **SetTime()**, **GetWakeupTime()**, and **SetWakeupTime()** fields of the EFI Runtime Services Table. See Services - Runtime Services for details on these services. After the four fields of the EFI Runtime Services Table have been initialized, the driver must install the **EFI_REAL_TIME_CLOCK_ARCH_PROTOCOL_GUID** on a new handle with a **NULL** interface pointer. The installation of this protocol informs the DXE Foundation that the real time clock–related services are now available and that the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table.

# Reset Architectural Protocol

## EFI_RESET_ARCH_PROTOCOL

### Summary

Provides the service required to reset a platform. This protocol must be produced by a runtime DXE driver and may only be consumed by the DXE Foundation.

### GUID

```
#define EFI_RESET_ARCH_PROTOCOL_GUID  \

{0x27CFAC88,0x46CC,0x11d4,0x9A,0x38,0x00,0x90,0x27,0x3F,0xC1,0x4D
}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **`ResetSystem()`** field of the EFI Runtime Services Table. See Services - Runtime Services for details on this service. After this field of the EFI Runtime Services Table has been initialized, the driver must install the **`EFI_RESET_ARCH_PROTOCOL_GUID`** on a new handle with a **`NULL`** interface pointer. The installation of this protocol informs the DXE Foundation that the reset system service is now available and that the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table.

# Runtime Architectural Protocol

## Runtime Architectural Protocol

The following topics provide a detailed description of the **EFI_RUNTIME_ARCH_PROTOCOL**. The DXE Foundation contains no runtime code, so all runtime code is contained in DXE Architectural Protocols. This is due to the fact that runtime code must be callable in physical or virtual mode. The Runtime Architectural Protocol contains the EFI runtime services that are callable only in physical mode. The Runtime Architectural Protocol can be thought of as the runtime portion of the DXE Foundation.

The Runtime Architectural Protocol contains support for transition of runtime drivers from physical mode calling to virtual mode calling.

## EFI_RUNTIME_ARCH_PROTOCOL

### Summary

Allows the runtime functionality of the DXE Foundation to be contained in a separate driver. It also provides hooks for the DXE Foundation to export information that is needed at runtime. As such, this protocol allows the DXE Foundation to manage runtime drivers and events. This protocol also implies that the runtime services required to transition to virtual mode, **SetVirtualAddressMap()** and **ConvertPointer()**, have been registered into the EFI Runtime Table in the EFI System Partition. This protocol must be produced by a runtime DXE driver and may only be consumed by the DXE Foundation.

### GUID

```
#define EFI RUNTIME ARCH PROTOCOL GUID  \
   {0xb7dfb4e1, 0x52f, 0x449f, 0x87, 0xbe, 0x98, 0x18, 0xfc, 0x91, 0xb7,
   0x33}
```

### Protocol Interface Structure

```
typedef struct EFI RUNTIME ARCH PROTOCOL {
     EFI LIST ENTRY              ImageHead;
     EFI LIST ENTRY              EventHead;
     UINTN                       MemoryDescriptorSize;
     UINT32                      MemoryDescriptorVersion;
     UINTN                       MemoryMapSize;
     EFI MEMORY DESCRIPTOR       *MemoryMapPhysical;
     EFI MEMORY DESCRIPTOR       *MemoryMapVirtural;
     BOOLEAN                     VirtualMode;
     BOOLEAN                     AtRuntime;
} EFI_RUNTIME_ARCH_PROTOCOL;
```

## Parameters

*ImageHead*

A list of type EFI RUNTIME IMAGE ENTRY where the DXE Foundation inserts items into the list and the Runtime AP consumes the data to implement the SetVertualAddressMap() call.

*EventHead*

A list of type EFI_RUNTIME_EVENT_ENTRY where the DXE Foundation inserts items into the list and the Runtime AP consumes the data to implement the SetVertualAddressMap() call.

*MemoryDescriptioSize*

Size of a memory descriptor that is returned by GetMemoryMap(). This value is updated by the DXE Foundation.

*MemoryDescriptioVersion*

Version of a memory descriptor that is returned by GetMemoryMap(). This value is updated by the DXE Foundation.

*MemoryMapSize*

Size of a memory map in bytes contained in *MemoryMapPhysical* and *MemoryMapVirtual*. This value is updated by the DXE Foundation when memory for *MemoryMapPhysical* gets allocated.

*MemoryMapPhysical*

Pointer to a runtime buffer that contains a copy of the memory map returned via GetMemoryMap(). The memory must be allocated by the DXE Foundation so that it is accounted for in the memory map.

*MemoryMapVirtual*

Pointer to *MemoryMapPhysical* that is updated to virtual mode after SetVirtualAddressMap(). The DXE Foundation updates this value when it updates *MemoryMapPhysical* with the same physical address. The Runtime AP is responsible for converting *MemoryMapVirtural* to a virtual pointer.

*VirtualMode*

Boolean that is TRUE if the SetVirtualAddressMap() has been called. This field is set by the Runtime AP. When the *VirtualMode* is TRUE *MemoryMapVirtural* pointer contains the virtual address of the *MemoryMapPhysical*.

*AtRuntime*

Boolean that is TRUE if ExitBootServices() has been called. This field is set by the Runtime AP.

## Related Definitions

```
//*******************************************************
// EFI_RUNTIME_IMAGE_ENTRY
//*******************************************************


typedef struct {
     VOID                    *ImageBase;
     UINT64                  ImageSize;
     VOID                    *RelocationData;
     EFI_HANDLE              Handle;
     EFI_LIST_ENTRY          Link;
} EFI_RUNTIME_IMAGE_ENTRY;
```

*ImageBase*

Start of image that has been loaded in memory. It is a pointer to either the DOS header of PE header of the image. Type EFI PHYSICAL ADDRESS is defined in the AllocatePages() function description of the EFI 1.10 Specification.

*ImageSize*

Size in bytes of the image represented by *ImageBase.*

*RelocationData*

Information about the fix-ups that were performed on *ImageBase* when it was loaded into memory. This information is needed when the virtual mode fix-ups are reapplied so that data that has been programmatically updated will not be fixed up. If code updates a global variable the code is responsible for fixing up the variable for virtual mode.

*Handle*

The *ImageHandle* passed into *ImageBase* when it was loaded. See EFI_IMAGE_ENTRY_POINT for the definition of *ImageHandle*.

*Link*

Entry for this node in the EFI_RUNTIME_ARCHITECTURE_PROTOCOL.*ImageHead* list.

```
//*******************************************************
// EFI_RUNTIME_EVENT_ENTRY
//*******************************************************


typedef struct {
UINT32                  Type;
EFI_TPL                 NotifyTpl;
EFI_EVENT_NOTIFY        NotifyFunction;
VOID                    *NotifyContext;
EFI_EVENT               *Event;
EFI_LIST_ENTRY          Link;

} EFI_RUNTIME_EVENT_ENTRY;
```

*Type*

Same as Type passed into CreateEvent().

*NotifyTpl*

The same as *NotifyTpl* passed into CreateEvent(). Type EFI_TPL is defined in the RaiseTPL() in the EFI 1.10 Specification.

*NotifyFunction*

The same as *NotifyFunction* passed into CreateEvent(). Type EFI_EVENT_NOTIFY is defined in the CreateEvent() function description.

*NotifyContext*

The same as *NotifyContext* passed into CreateEvent().

*Event*

The EFI_EVENT returned by CreateEvent(). Event must be in runtime memory. Type EFI_EVENT is defined in the CreateEvent() function description.

*Link*

Entry for this node in the EFI_RUNTIME_ARCHITECTURE_PROTOCOL.EventHead list.

## Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **`SetVirtualAddressMap()`** and **`ConvertPointer()`** fields of the EFI Runtime Services Table and the **`CalculateCrc32()`** field of the EFI Boot Services Table. See Services - Runtime Services and Services - Boot Services for details on these services. After the two fields of the EFI Runtime Services Table and the one field of the EFI Boot Services Table have been initialized, the driver must install the **`EFI_RUNTIME_ARCH_PROTOCOL_GUID`** on a new handle with an **`EFI_RUNTIME_ARCH_PROTOCOL`** interface pointer. The installation of this protocol informs the DXE Foundation that the virtual memory services and the 32-bit CRC services are now available, and the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table and the 32-bit CRC of the EFI Boot Services Table.

All runtime DXE Foundation services are provided by the **`EFI_RUNTIME_ARCH_PROTOCOL`**. This includes the support for registering runtime images that must be fixed up again when a transition is made from physical mode to virtual mode. This protocol also supports all events that are defined to fire at runtime. This protocol also contains a CRC-32 function that will be used by the DXE Foundation as a boot service. The **`EFI_RUNTIME_ARCH_PROTOCOL`** needs the CRC-32 function when a transition is made from physical mode to virtual mode and the EFI System Table and EFI Runtime Table are fixed up with virtual pointers.

# Security Architectural Protocol

## EFI_SECURITY_ARCH_PROTOCOL

### Summary

Abstracts security-specific functions from the DXE Foundation. This protocol must be produced by a boot service or runtime DXE driver and may only be consumed by the DXE Foundation and any other DXE drivers that need to validate the authentication of files.

### GUID

```
#define EFI_SECURITY_ARCH_PROTOCOL_GUID  \

{0xA46423E3,0x4617,0x49f1,0xB9,0xFF,0xD1,0xBF,0xA9,0x11,0x58,0x39
}
```

### Protocol Interface Structure

```
typedef struct _EFI_SECURITY_ARCH_PROTOCOL {
  EFI_SECURITY_FILE_AUTHENTICATION_STATE
                                FileAuthenticationState;
} EFI_SECURITY_ARCH_PROTOCOL;
```

### Parameters

*FileAuthenticationState*

This service is called upon fault with respect to the authentication of a section of a file. See the **FileAuthenticationState()** function description.

### Description

The **EFI_SECURITY_ARCH_PROTOCOL** is used to abstract platform-specific policy from the DXE Foundation. This includes locking flash upon failure to authenticate, attestation logging, and other exception operations.

The driver that produces the **EFI_SECURITY_ARCH_PROTOCOL** may also optionally install the **EFI_SECURITY_POLICY_PROTOCOL_GUID** onto a new handle with a **NULL** interface. The existence of this GUID in the protocol database means that the GUIDed Section Extraction Protocol should authenticate the contents of an Authentication Section. The expectation is that the GUIDed Section Extraction protocol will look for the existence of the **EFI_SECURITY_POLICY_PROTOCOL_GUID** in the protocol database. If it exists, then the publication thereof is taken as an injunction to attempt an authentication of any section wrapped in an Authentication Section. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Specification* for details on the GUIDed Section Extraction Protocol and Authentication Sections.

## Additional GUID Definitions

```
#define EFI_SECURITY_POLICY_PROTOCOL_GUID  \

{0x78E4D245,0xCD4D,0x4a05,0xA2,0xBA,0x47,0x43,0xE8,0x6C,0xFC,0xAB
}
```

# EFI_SECURITY_ARCH_PROTOCOL. FileAuthenticationState()

## Summary

The DXE Foundation uses this service to check the authentication status of a file. This allows the system to execute a platform-specific policy in response the different authentication status values.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_SECURITY_FILE_AUTHENTICATION_STATE) (
  IN  EFI_SECURITY_ARCH_PROTOCOL   *This,
  IN  UINT32                       AuthenticationStatus,
  IN  EFI_DEVICE_PATH_PROTOCOL     *File
  );
```

## Parameters

*This*

The **EFI_SECURITY_ARCH_PROTOCOL** instance.

*AuthenticationStatus*

The authentication type returned from the Section Extraction Protocol. See the *Intel® Platform Innovation Framework for EFI Firmware Volume Specification* for details on this type.

*File*

A pointer to the device path of the file that is being dispatched. This will optionally be used for logging. Type **EFI_DEVICE_PATH_PROTOCOL** is defined Chapter 8 of the *EFI 1.10 Specification.*

## Description

The **EFI_SECURITY_ARCH_PROTOCOL** (SAP) is used to abstract platform-specific policy from the DXE Foundation response to an attempt to use a file that returns a given status for the authentication check from the section extraction protocol.

The possible responses in a given SAP implementation may include locking flash upon failure to authenticate, attestation logging for all signed drivers, and other exception operations. The *File* parameter allows for possible logging within the SAP of the driver.

If *File* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

If the file specified by *File* with an authentication status specified by *AuthenticationStatus* is safe for the DXE Foundation to use, then **EFI_SUCCESS** is returned.

If the file specified by *File* with an authentication status specified by *AuthenticationStatus* is not safe for the DXE Foundation to use under any circumstances, then **EFI_ACCESS_DENIED** is returned.

If the file specified by *File* with an authentication status specified by *AuthenticationStatus* is not safe for the DXE Foundation to use right now, but it might be possible to use it at a future time, then **EFI_SECURITY_VIOLATION** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The file specified by *File* did authenticate, and the platform policy dictates that the DXE Foundation may use *File*. |
| EFI_INVALID_PARAMETER | *File* is **NULL**. |
| EFI_SECURITY_VIOLATION | The file specified by *File* did not authenticate, and the platform policy dictates that *File* should be placed in the untrusted state. A file may be promoted from the untrusted to the trusted state at a future time with a call to the **Trust()** DXE Service. |
| EFI_ACCESS_DENIED | The file specified by *File* did not authenticate, and the platform policy dictates that *File* should not be used for any purpose. |

# Status Code Architectural Protocol

## EFI_STATUS_CODE_ARCH_PROTOCOL

### Summary

Provides the service required to report a status code to the platform firmware. This protocol must be produced by a runtime DXE driver and may be consumed only by the DXE Foundation.

### GUID

```
#define EFI_STATUS_CODE_ARCH_PROTOCOL_GUID  \

{0xd98e3ea3,0x6f39,0x4be4,0x82,0xce,0x5a,0x89,0xc,0xcb,0x2c,0x95}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **ReportStatusCode()** field of the EFI Runtime Services Table. See Service - Runtime Services for details on this service. After this field of the EFI Runtime Services Table has been initialized, the driver must install the **EFI_STATUS_CODE_ARCH_PROTOCOL_GUID** on a new handle with a **NULL** interface pointer. The installation of this protocol informs the DXE Foundation that the status code service is now available, and the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table.

# Timer Architectural Protocol

## EFI_TIMER_ARCH_PROTOCOL

### Summary

Used to set up a periodic timer interrupt using a platform specific timer, and a processor-specific interrupt vector. This protocol enables the use of the **SetTimer()** Boot Service. This protocol must be produce by a boot service or runtime DXE driver and may only be consumed by the DXE Foundation or DXE drivers that produce other DXE Architectural Protocols.

### GUID

```
#define EFI_TIMER_ARCH_PROTOCOL_GUID  \

{0x26baccb3,0x6f42,0x11d4,0xbc,0xe7,0x0,0x80,0xc7,0x3c,0x88,0x81}
```

### Protocol Interface Structure

```
typedef struct _EFI_TIMER_ARCH_PROTOCOL {
  EFI_TIMER_REGISTER_HANDLER          RegisterHandler;
  EFI_TIMER_SET_TIMER_PERIOD          SetTimerPeriod;
  EFI_TIMER_GET_TIMER_PERIOD          GetTimerPeriod;
  EFI_TIMER_GENERATE_SOFT_INTERRUPT   GenerateSoftInterrupt;
} EFI_TIMER_ARCH_PROTOCOL;
```

### Parameters

*RegisterHandler*

Registers a handler that will be called each time the timer interrupt fires. See the **RegisterHandler()** function description. *TimerPeriod* defines the minimum time between timer interrupts, so *TimerPeriod* will also be the minimum time between calls to the registered handler.

*SetTimerPeriod*

Sets the period of the timer interrupt in 100 nS units. See the **SetTimerPeriod()** function description. This function is optional and may return **EFI_UNSUPPORTED**. If this function is supported, then the timer period will be rounded up to the nearest supported timer period.

*GetTimerPeriod*

Retrieves the period of the timer interrupt in 100 nS units. See the **GetTimerPeriod()** function description.

*GenerateSoftInterrupt*

Generates a soft timer interrupt that simulates the firing of the timer interrupt. This service can be used to invoke the registered handler if the timer interrupt has been

masked for a period of time. See the **GenerateSoftInterrupt()** function description.

## Description

This protocol provides the services to initialize a periodic timer interrupt and to register a handler that is called each time the timer interrupt fires. It may also provide a service to adjust the rate of the periodic timer interrupt. When a timer interrupt occurs, the handler is passed the amount of time that has passed since the previous timer interrupt.

## EFI_TIMER_ARCH_PROTOCOL.RegisterHandler()

### Summary

Registers a handler that is called each timer the timer interrupt fires.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_TIMER_REGISTER_HANDLER) (
  IN EFI_TIMER_ARCH_PROTOCOL   *This,
  IN EFI_TIMER_NOTIFY          NotifyFunction
  );
```

### Parameters

*This*

The **EFI_TIMER_ARCH_PROTOCOL** instance.

*NotifyFunction*

The function to call when a timer interrupt fires. This function executes at **EFI_TPL_HIGH_LEVEL**. The DXE Foundation will register a handler for the timer interrupt, so it can know how much time has passed. This information is used to signal timer based events. **NULL** will unregister the handler. Type **EFI_TIMER_NOTIFY** is defined in "Related Definitions" below.

### Description

This function registers the handler *NotifyFunction* so it is called every time the timer interrupt fires. It also passes the amount of time since the last handler call to the *NotifyFunction*. If *NotifyFunction* is **NULL**, then the handler is unregistered. If the handler is registered, then **EFI_SUCCESS** is returned. If the processor does not support registering a timer interrupt handler, then **EFI_UNSUPPORTED** is returned. If an attempt is made to register a handler when a handler is already registered, then **EFI_ALREADY_STARTED** is returned. If an attempt is made to unregister a handler when a handler is not registered, then **EFI_INVALID_PARAMETER** is returned. If an error occurs attempting to register the *NotifyFunction* with the timer interrupt, then **EFI_DEVICE_ERROR** is returned.

### Related Definitions

```
typedef
VOID
(EFIAPI *EFI_TIMER_NOTIFY) (
  IN UINT64  Time
  );
```

*Time*

Time since the last timer interrupt in 100 ns units. This will typically be *TimerPeriod*, but if a timer interrupt is missed, and the **EFI_TIMER_ARCH_PROTOCOL** driver can detect missed interrupts, then *Time* will contain the actual amount of time since the last interrupt.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The timer handler was registered. |
| EFI_UNSUPPORTED | The platform does not support timer interrupts. |
| EFI_ALREADY_STARTED | *NotifyFunction* is not **NULL**, and a handler is already registered. |
| EFI_INVALID_PARAMETER | *NotifyFunction* is **NULL**, and a handler was not previously registered. |
| EFI_DEVICE_ERROR | The timer handler could not be registered. |

## EFI_TIMER_ARCH_PROTOCOL.SetTimerPeriod()

### Summary

Sets the rate of the periodic timer interrupt.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_TIMER_SET_TIMER_PERIOD) (
  IN  EFI_TIMER_ARCH_PROTOCOL  *This,
  IN  UINT64                   TimerPeriod
  );
```

### Parameters

*This*

The **EFI_TIMER_ARCH_PROTOCOL** instance.

*TimerPeriod*

The rate to program the timer interrupt in 100 nS units. If the timer hardware is not programmable, then **EFI_UNSUPPORTED** is returned. If the timer is programmable, then the timer period will be rounded up to the nearest timer period that is supported by the timer hardware. If *TimerPeriod* is set to 0, then the timer interrupts will be disabled.

### Description

This function adjusts the period of timer interrupts to the value specified by *TimerPeriod*. If the timer period is updated, then **EFI_SUCCESS** is returned. If the timer hardware is not programmable, then **EFI_UNSUPPORTED** is returned. If an error occurs while attempting to update the timer period, then the timer hardware will be put back in its state prior to this call, and **EFI_DEVICE_ERROR** is returned. If *TimerPeriod* is 0, then the timer interrupt is disabled. This is not the same as disabling the processor's interrupts. Instead, it must either turn off the timer hardware, or it must adjust the interrupt controller so that a processor interrupt is not generated when the timer interrupt fires.

### Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The timer period was changed. |
| EFI_UNSUPPORTED | The platform cannot change the period of the timer interrupt. |
| EFI_DEVICE_ERROR | The timer period could not be changed due to a device error. |

# EFI_TIMER_ARCH_PROTOCOL.GetTimerPeriod()

## Summary

Retrieves the rate of the periodic timer interrupt.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_TIMER_GET_TIMER_PERIOD) (
  IN  EFI_TIMER_ARCH_PROTOCOL   *This,
  OUT UINT64                    *TimerPeriod
  );
```

## Parameters

*This*

The **EFI_TIMER_ARCH_PROTOCOL** instance.

*TimerPeriod*

A pointer to the timer period to retrieve in 100 ns units. If 0 is returned, then the timer is currently disabled.

## Description

This function retrieves the period of timer interrupts in 100 ns units, returns that value in *TimerPeriod*, and returns **EFI_SUCCESS**. If *TimerPeriod* is **NULL**, then **EFI_INVALID_PARAMETER** is returned. If a *TimerPeriod* of 0 is returned, then the timer is currently disabled.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The timer period was returned in *TimerPeriod*. |
| EFI_INVALID_PARAMETER | *TimerPeriod* is **NULL**. |

## EFI_TIMER_ARCH_PROTOCOL. GenerateSoftInterrupt()

### Summary

Generates a soft timer interrupt.

### Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_TIMER_GENERATE_SOFT_INTERRUPT) (
  IN  EFI_TIMER_ARCH_PROTOCOL    *This
  );
```

### Parameters

*This*

The **EFI_TIMER_ARCH_PROTOCOL** instance.

### Description

This function generates a soft timer interrupt. If the platform does not support soft timer interrupts, then **EFI_UNSUPPORTED** is returned. Otherwise, **EFI_SUCCESS** is returned. If a handler has been registered through the **EFI_TIMER_ARCH_PROTOCOL.RegisterHandler()** service, then a soft timer interrupt will be generated. If the timer interrupt is enabled when this service is called, then the registered handler will be invoked. The registered handler should not be able to distinguish a hardware-generated timer interrupt from a software-generated timer interrupt.

### Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The soft timer interrupt was generated. |
| EFI_UNSUPPORTED | The platform does not support the generation of soft timer interrupts. |

# Variable Architectural Protocol

## EFI_VARIABLE_ARCH_PROTOCOL

### Summary

Provides the services required to get and set environment variables. This protocol must be produced by a runtime DXE driver and may be consumed only by the DXE Foundation.

### GUID

```
#define EFI_VARIABLE_ARCH_PROTOCOL_GUID  \

{0x1e5668e2,0x8481,0x11d4,0xbc,0xf1,0x0,0x80,0xc7,0x3c,0x88,0x81}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver is responsible for initializing the **`GetVariable()`**, **`GetNextVariableName()`**, and **`SetVariable()`** fields of the EFI Runtime Services Table. See Services - Runtime Services for details on these services. After the three fields of the EFI Runtime Services Table have been initialized, the driver must install the **`EFI_VARIABLE_ARCH_PROTOCOL_GUID`** on a new handle with a NULL interface pointer. The installation of this protocol informs the DXE Foundation that the read-only and the volatile environment variable related services are now available and that the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table. The full complement of environment variable services are not available until both this protocol and **`EFI_VARIABLE_WRITE_ARCH_PROTOCOL`** are installed. DXE drivers that require read-only access or read/write access to volatile environment variables must have this architectural protocol in their dependency expressions. DXE drivers that require write access to nonvolatile environment variables must have the **`EFI_VARIABLE_WRITE_ARCH_PROTOCOL`** in their dependency expressions.

# Variable Write Architectural Protocol

## EFI_VARIABLE_WRITE_ARCH_PROTOCOL

### Summary

Provides the services required to set nonvolatile environment variables. This protocol must be produced by a runtime DXE driver and may be consumed only by the DXE Foundation.

### GUID

```
#define EFI_VARIABLE_WRITE_ARCH_PROTOCOL_GUID  \
  {0x6441f818,0x6362,0x4e44,0xb5,0x70,0x7d,0xba,0x31,0xdd,0x24,0x
53}
```

### Description

The DXE driver that produces this protocol must be a runtime driver. This driver may update the **`SetVariable()`** field of the EFI Runtime Services Table. See Services - Runtime Services for details on this service. After the EFI Runtime Services Table has been initialized, the driver must install the **`EFI_VARIABLE_WRITE_ARCH_PROTOCOL_GUID`** on a new handle with a NULL interface pointer. The installation of this protocol informs the DXE Foundation that the write services for nonvolatile environment variables are now available and that the DXE Foundation must update the 32-bit CRC of the EFI Runtime Services Table. The full complement of environment variable services are not available until both this protocol and **`EFI_VARIABLE_ARCH_PROTOCOL`** are installed. DXE drivers that require read-only access or read/write access to volatile environment variables must have the **`EFI_VARIABLE_WRITE_ARCH_PROTOCOL`** in their dependency expressions. DXE drivers that require write access to nonvolatile environment variables must have this architectural protocol in their dependency expressions.

# Watchdog Timer Architectural Protocol

## Watchdog Timer Architectural Protocol

This following topics provide a detailed description of the **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL**. This protocol is used to implement the Boot Service **SetWatchdogTimer()**. The watchdog timer may be implemented in software using Boot Services, or it may be implemented with specialized hardware. The protocol provides a service to register a handler when the watchdog timer fires and a service to set the amount of time to wait before the watchdog timer is fired.

## EFI_WATCHDOG_TIMER_ARCH_PROTOCOL

### Summary

Used to program the watchdog timer and optionally register a handler when the watchdog timer fires. This protocol must be produced by a boot service or runtime DXE driver and may be consumed only by the DXE Foundation or DXE drivers that produce other DXE Architectural Protocols. If a platform wishes to perform a platform-specific action when the watchdog timer expires, then the DXE driver that contains the implementation of the **EFI_BDS_ARCH_PROTOCOL** should use this protocol's **RegisterHandler()** service.

### GUID

```
#define EFI_WATCHDOG_TIMER_ARCH_PROTOCOL_GUID  \

{0x665E3FF5,0x46CC,0x11d4,0x9A,0x38,0x00,0x90,0x27,0x3F,0xC1,0x4D
}
```

### Protocol Interface Structure

```
typedef struct _EFI_WATCHDOG_TIMER_ARCH_PROTOCOL {
  EFI_WATCHDOG_TIMER_REGISTER_HANDLER  RegisterHandler;
  EFI_WATCHDOG_TIMER_SET_TIMER_PERIOD  SetTimerPeriod;
  EFI_WATCHDOG_TIMER_GET_TIMER_PERIOD  GetTimerPeriod;
} EFI_WATCHDOG_TIMER_ARCH_PROTOCOL;
```

### Parameters

*RegisterHandler*

Registers a handler that is invoked when the watchdog timer fires. See the **RegisterHandler()** function description.

*SetTimerPeriod*

Sets the amount of time in 100 ns units to wait before the watchdog timer is fired. See the **SetTimerPeriod()** function description. If this function is supported, then the watchdog timer period will be rounded up to the nearest supported watchdog timer period.

*GetTimerPeriod*

Retrieves the amount of time in 100 ns units that the system will wait before the watchdog timer is fired. See the **GetTimerPeriod()** function description.

## Description

This protocol provides the services required to implement the Boot Service **SetWatchdogTimer()**. It provides a service to set the amount of time to wait before firing the watchdog timer, and it also provides a service to register a handler that is invoked when the watchdog timer fires. This protocol can implement the watchdog timer by using the event and timer Boot Services, or it can make use of custom hardware. When the watchdog timer fires, control will be passed to a handler if one has been registered. If no handler has been registered, or the registered handler returns, then the system will be reset by calling the Runtime Service **ResetSystem()**.

# EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. RegisterHandler()

## Summary

Registers a handler that is to be invoked when the watchdog timer fires.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_WATCHDOG_TIMER_REGISTER_HANDLER) (
  IN EFI_WATCHDOG_TIMER_ARCH_PROTOCOL   *This,
  IN EFI_WATCHDOG_TIMER_NOTIFY          NotifyFunction
  );
```

## Parameters

*This*

The **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL** instance.

*NotifyFunction*

The function to call when the watchdog timer fires. If this is **NULL**, then the handler will be unregistered. Type **EFI_WATCHDOG_TIMER_NOTIFY** is defined in "Related Definitions" below.

## Description

This function registers a handler that is to be invoked when the watchdog timer fires. By default, **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL** will call the Runtime Service **ResetSystem()** when the watchdog timer fires. If a *NotifyFunction* is registered, then *NotifyFunction* will be called before the Runtime Service **ResetSystem()** is called. If *NotifyFunction* is **NULL**, then the watchdog handler is unregistered. If a watchdog handler is registered, then **EFI_SUCCESS** is returned. If an attempt is made to register a handler when a handler is already registered, then **EFI_ALREADY_STARTED** is returned. If an attempt is made to uninstall a handler when a handler is not installed, then return **EFI_INVALID_PARAMETER**.

## Related Definitions

```
typedef
VOID
(EFIAPI *EFI_WATCHDOG_TIMER_NOTIFY) (
  IN UINT64  Time
  );
```

*Time*

The time in 100 ns units that has passed since the watchdog timer was armed. For the notify function to be called, this must be greater than *TimerPeriod*.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The watchdog timer handler was registered or unregistered. |
| EFI_ALREADY_STARTED | *NotifyFunction* is not **NULL**, and a handler is already registered. |
| EFI_INVALID_PARAMETER | *NotifyFunction* is **NULL**, and a handler was not previously registered. |

# EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. SetTimerPeriod()

## Summary

Sets the amount of time in the future to fire the watchdog timer.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_WATCHDOG_TIMER_SET_TIMER_PERIOD) (
  IN EFI_WATCHDOG_TIMER_ARCH_PROTOCOL   *This,
  IN UINT64                             TimerPeriod
  );
```

## Parameters

*This*

The **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL** instance.

*TimerPeriod*

The amount of time in 100 nS units to wait before the watchdog timer is fired. If *TimerPeriod* is zero, then the watchdog timer is disabled.

## Description

This function sets the amount of time to wait before firing the watchdog timer to *TimerPeriod* 100 nS units. If *TimerPeriod* is zero, then the watchdog timer is disabled.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The watchdog timer has been programmed to fire in *Time* 100 nS units. |
| EFI_DEVICE_ERROR | A watchdog timer could not be programmed due to a device error. |

# EFI_WATCHDOG_TIMER_ARCH_PROTOCOL. GetTimerPeriod()

## Summary

Retrieves the amount of time in 100 ns units that the system will wait before firing the watchdog timer.

## Prototype

```
typedef
EFI_STATUS
(EFIAPI *EFI_WATCHDOG_TIMER_GET_TIMER_PERIOD) (
  IN  EFI_WATCHDOG_TIMER_ARCH_PROTOCOL  *This,
  OUT UINT64                            *TimerPeriod
  );
```

## Parameters

*This*

The **EFI_WATCHDOG_TIMER_ARCH_PROTOCOL** instance.

*TimerPeriod*

A pointer to the amount of time in 100 nS units that the system will wait before the watchdog timer is fired. If *TimerPeriod* of zero is returned, then the watchdog timer is disabled.

## Description

This function retrieves the amount of time the system will wait before firing the watchdog timer. This period is returned in *TimerPeriod*, and **EFI_SUCCESS** is returned. If *TimerPeriod* is **NULL**, then **EFI_INVALID_PARAMETER** is returned.

## Status Codes Returned

| | |
|---|---|
| EFI_SUCCESS | The amount of time that the system will wait before firing the watchdog timer was returned in *TimerPeriod*. |
| EFI_INVALID_PARAMETER | *TimerPeriod* is **NULL**. |

# 13
# Returned Status Codes

## Returned Status Codes

EFI interfaces return an **EFI_STATUS** code. The topics in this book discuss the following:

- Ranges of EFI_STATUS codes
- Success codes
- Error codes
- Warning codes

Error codes also have their highest bit set, so all error codes have negative values. The range of status codes that have the highest bit set and the next to highest bit clear are reserved for use by EFI. The range of status codes that have both the highest bit set and the next to highest bit set are reserved for use by OEMs.

Success and warning codes have their highest bit clear, so all success and warning codes have positive values. The range of status codes that have both the highest bit clear and the next to highest bit clear are reserved for use by EFI. The range of status code that have the highest bit clear and the next to highest bit set are reserved for use by OEMs.

## EFI_STATUS Codes Ranges

The following table lists the ranges of **EFI_STATUS** codes.

| IA-32 Range | Itanium® Architecture Range | Description |
|---|---|---|
| 0x00000000-0x3fffffff | 0x0000000000000000-0x3fffffffffffffff | Success and warning codes reserved for use by EFI. See EFI_STATUS Success Codes (High Bit Clear) and EFI_STATUS Warning Codes (High Bit Clear) for valid values in this range. |
| 0x40000000-0x7fffffff | 0x4000000000000000-0x7fffffffffffffff | Success and warning codes reserved for use by OEMs. |
| 0x80000000-0xbfffffff | 0x8000000000000000-0xbfffffffffffffff | Error codes reserved for use by EFI. See EFI_STATUS Error Codes (High Bit Set) for valid values for this range. |
| 0xc0000000-0xffffffff | 0xc000000000000000-0xffffffffffffffff | Error codes reserved for use by OEMs. |

## EFI_STATUS Success Codes (High Bit Clear)

The following table lists the success codes for **EFI_STATUS**.

| Mnemonic | Value | Description |
|---|---|---|
| EFI_SUCCESS | 0 | The operation completed successfully. |

# EFI_STATUS Error Codes (High Bit Set)

The following table lists the error codes for **EFI_STATUS**.

| Mnemonic | Value | Description |
|---|---|---|
| EFI_LOAD_ERROR | 1 | The image failed to load. |
| EFI_INVALID_PARAMETER | 2 | A parameter was incorrect. |
| EFI_UNSUPPORTED | 3 | The operation is not supported. |
| EFI_BAD_BUFFER_SIZE | 4 | The buffer was not the proper size for the request. |
| EFI_BUFFER_TOO_SMALL | 5 | The buffer is not large enough to hold the requested data. The required buffer size is returned in the appropriate parameter when this error occurs. |
| EFI_NOT_READY | 6 | There is no data pending upon return. |
| EFI_DEVICE_ERROR | 7 | The physical device reported an error while attempting the operation. |
| EFI_WRITE_PROTECTED | 8 | The device cannot be written to. |
| EFI_OUT_OF_RESOURCES | 9 | A resource has run out. |
| EFI_VOLUME_CORRUPTED | 10 | An inconstancy was detected on the file system causing the operating to fail. |
| EFI_VOLUME_FULL | 11 | There is no more space on the file system. |
| EFI_NO_MEDIA | 12 | The device does not contain any medium to perform the operation. |
| EFI_MEDIA_CHANGED | 13 | The medium in the device has changed since the last access. |
| EFI_NOT_FOUND | 14 | The item was not found. |
| EFI_ACCESS_DENIED | 15 | Access was denied. |
| EFI_NO_RESPONSE | 16 | The server was not found or did not respond to the request. |
| EFI_NO_MAPPING | 17 | A mapping to a device does not exist. |
| EFI_TIMEOUT | 18 | The timeout time expired. |
| EFI_NOT_STARTED | 19 | The protocol has not been started. |
| EFI_ALREADY_STARTED | 20 | The protocol has already been started. |
| EFI_ABORTED | 21 | The operation was aborted. |
| EFI_ICMP_ERROR | 22 | An ICMP error occurred during the network operation. |
| EFI_TFTP_ERROR | 23 | A TFTP error occurred during the network operation. |
| EFI_PROTOCOL_ERROR | 24 | A protocol error occurred during the network operation. |
| EFI_INCOMPATIBLE_VERSION | 25 | The function encountered an internal version that was incompatible with a version requested by the caller. |
| EFI_SECURITY_VIOLATION | 26 | The function was not performed due to a security violation. |
| EFI_CRC_ERROR | 27 | A CRC error was detected. |
| EFI_NOT_AVAILABLE_YET | 28 | The service is not available yet because one of its dependencies has not been satisfied yet. |
| EFI_UNLOAD_IMAGE | 29 | If this value is returned by an EFI image, then the image should be unloaded. |

## EFI_STATUS Warning Codes (High Bit Clear)

The following table lists the warning codes for **EFI_STATUS**.

| Mnemonic | Value | Description |
|---|---|---|
| EFI_WARN_UNKOWN_GLYPH | 1 | The Unicode string contained one or more characters that the device could not render and were skipped. |
| EFI_WARN_DELETE_FAILURE | 2 | The handle was closed, but the file was not deleted. |
| EFI_WARN_WRITE_FAILURE | 3 | The handle was closed, but the data to the file was not flushed properly. |
| EFI_WARN_BUFFER_TOO_SMALL | 4 | The resulting buffer was too small, and the data was truncated to the buffer size. |

# 14
# Dependency Expression Grammar

## Dependency Expression Grammar

This topic contains an example BNF grammar for a DXE driver dependency expression compiler that converts a dependency expression source file into a dependency section of a DXE driver stored in a firmware volume.

## Example Dependency Expression BNF Grammar

```
<depex>    ::= BEFORE <guid>
             | AFTER <guid>
             | SOR <bool>
             | <bool>
<bool>     ::= <bool> AND <term>
             | <bool> OR <term>
             | <term>
<term>     ::= NOT <factor>
             | <factor>
<factor>   ::= <bool>
             | TRUE
             | FALSE
             | GUID
             | END
<guid>     ::= '{' <hex32> ',' <hex16> ',' <hex16> ','
               <hex8> ',' <hex8> ',' <hex8> ',' <hex8> ','
               <hex8> ',' <hex8> ',' <hex8> ',' <hex8> '}'
<hex32>    ::= <hexprefix> <hexvalue>
<hex16>    ::= <hexprefix> <hexvalue>
<hex8>     ::= <hexprefix> <hexvalue>
<hexprefix>::= '0' 'x'
             | '0' 'X'
<hexvalue> ::= <hexdigit> <hexvalue>
             | <hexdigit>
<hexdigit> ::= [0-9]
             | [a-f]
             | [A-F]
```

## Sample Dependency Expressions

The following contains three examples of source statements using the BNF grammar from above along with the opcodes, operands, and binary encoding that a dependency expression compiler would generate from these source statements.

```
//
// Source
//
EFI_CPU_IO_PROTOCOL_GUID AND EFI_CPU_ARCH_PROTOCOL_GUID END

//
// Opcodes, Operands, and Binary Encoding
//
ADDR   BINARY                       MNEMONIC
====   ======================
================================================
0x00 : 02                           PUSH
0x01 : 26 25 73 b0 c8 38 40 4b      EFI_CPU_IO_PROTOCOL_GUID
       88 77 61 c7 b0 6a ac 45
0x11 : 02                           PUSH
0x12 : b1 cc ba 26 42 6f d4 11      EFI_CPU_ARCH_PROTOCOL_GUID
       bc e7 00 80 c7 3c 88 81
0x22 : 03                           AND
0x23 : 08                           END

//
// Source
//
AFTER (EFI_CPU_DRIVER_FILE_NAME_GUID) END

//
// Opcodes, Operands, and Binary Encoding
//
ADDR   BINARY                       MNEMONIC
====   ======================
================================================
0x00 : 01                           AFTER
0x01 : 93 e5 7b 98 43 16 0b 45      EFI_CPU_DRIVER_FILE_NAME_GUID
       be 4f 8f 07 66 6e 36 56
0x11 : 08                           END

//
// Source
//
SOR EFI_CPU_IO_PROTOCOL_GUID END

//
// Opcodes, Operands and Binary Encoding
//
ADDR   BINARY                       MNEMONIC
====   ======================
================================================
0x00 : 09                           SOR
0x01 : 02                           PUSH
0x02 : b1 cc ba 26 42 6f d4 11      EFI_CPU_IO_PROTOCOL_GUID
       bc e7 00 80 c7 3c 88 81
0x12 : 03                           END
```